

Social Rankings in Human-Computer Committees

Moshe Bitan

Bar-Ilan University, Israel

Ya'akov Gal

Ben-Gurion University, Israel

Sarit Kraus* Elad Dokow Amos Azaria

Bar-Ilan University, Israel

Abstract

Despite committees and elections being widespread in the real-world, the design of agents for operating in human-computer committees has received far less attention than the theoretical analysis of voting strategies. We address this gap by providing an agent design that outperforms other voters in groups comprising both people and computer agents. In our setting participants vote by simultaneously submitting a ranking over a set of candidates and the election system uses a social welfare rule to select a ranking that minimizes disagreements with participants' votes. We ran an extensive study in which hundreds of people participated in repeated voting rounds with other people as well as computer agents that differed in how they employ strategic reasoning in their voting behavior. Our results show that over time, people learn to deviate from truthful voting strategies, and use heuristics to guide their play, such as repeating their vote from the previous round. We show that a computer agent using a best response voting strategy was able to outperform people in the game. Our study has implication for agent designers, highlighting the types of strategies that enable agents to succeed in committees comprising both human and computer participants. This is the first work to study the role of computer agents in voting settings involving both human and agent participants.

Introduction

Voting systems have been used by people for centuries as tools for group decision making (Riker and Ordeshook 1968; Cox 1997; Palfrey 2009). More recently, voting and aggregation methods have been utilized by computers for tasks such as aggregating search results from the web (Dwork et al. 2001), collaborative filtering (Pennock et al. 2000) and planning (Ephrati, Rosenschein, and others 1993).

In virtually all electoral systems, participants can affect the result of the election by manipulating their vote, and such strategic voting behavior has been studied from both a theoretical and psychological perspective. As computers become ubiquitous in people's lives, heterogeneous group activities

of computer systems and people are becoming more prevalent. As a result, opportunities arise for computer agents to participate in voting systems, whether as autonomous agents or proxies for individual people. As an example, consider a recent on-line poll for ranking the world's seven wonders.¹ Suppose a user prefers the Golden Gate Bridge to Yellowstone Park. However, the user's most preferred choice—the Big Sur—is in close competition with the Golden Gate Bridge. Based on its beliefs about others' rankings, the proxy agent may reverse this preference in the user's ranking in order to ensure that Big Sur is chosen.

The contribution of this paper is an agent-design that outperforms other voters in mixed network settings involving both human and computer participants. In our setting all participants are assigned a preferred ranking over a set of candidates prior to commencing a series of voting rounds. At each round participants vote by simultaneously submitting a ranking over the set of candidates. The election system adapts the Kemeny-Young method (1978; 1959) that minimizes the sum of conflicts with the votes that are submitted by the participants. The utility of participants is proportional to the extent to which the chosen ranking agrees with their preferences. Such settings are analogous to real-world voting scenarios such as rating grant proposals and ranking applicants for positions in academia, industry or competitions.

We designed a three-player game that implements the voting system described above using a budget allocation analogy. The preferences of participants over the various sectors were chosen such that players could potentially improve their score in the game if they deviated from their truthful vote. We formalized several voting strategies for the game that differ in the extent to which they reason strategically about other's voting behavior. We conducted an extensive empirical study in which hundreds of human subjects played this game repeatedly with other people as well as computer agents that varied in the extent to which they voted strategically. We hypothesized that over time, people would vote less truthfully, and that computer agents using various levels of strategic voting would be able to outperform people.

Our results show that people deviate more from their

*Sarit Kraus is also affiliated with the University of Maryland Institute for Advanced Computer Studies.
Copyright © 2013, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<http://world.n7w.com/>. Other examples exist, like sites for ranking Time Magazine's 100 most influential people.

truthful votes in later rounds than in earlier rounds, but that this deviation does not necessarily result in an improvement in performance. Although people’s behavior was erratic, about 40% of the time their actions corresponded to voting their true preferences or repeating their vote in the previous round. A computer agent using a best-response strategy to people’s voting actions in the previous round was able to outperform people, as well as a baseline agent that consistently voted according to its true preferences. The efficacy of the best-response agent is highlighted by the fact that its performance was not significantly different than that of an oracle strategy that used the optimal vote in retrospect. The significance of this work is in demonstrating the success of using computational voting strategies when interacting with people. It is the first work to study the performance of agents using different voting strategies in mixed networks involving people and other computer agents.

Related Work

Voting systems and their convergence have been studied extensively in computer science (Meir et al. 2010; Reijngoud and Endriss 2012) and economics (Gohar 2012; Dhillon and Lockwood 2004). The most widely used voting rule is the plurality rule, in which each voter has one vote and the winner is determined as the candidate that receives the highest number of votes. Other popular voting rules, such as the Borda rule, allow voters to order the candidates, and the winner is determined by the candidate that receives the most points (relative to its positions in all of the voters’ rankings). However, all voting rules are susceptible to manipulation, that is, self-interested players have an incentive to vote strategically against their true preferences in certain situations (Gibbard 1977; Satterthwaite 1975). Consequently, studies in behavioral economics emerged which studied the effect of these voting rules on people’s voting strategies (Regenwetter and Rykhlevskaia 2007). Specifically, Forsythe et al. (1996) studied the effect of different voting rules on people’s voting strategies in three-candidate elections in which a single candidate is elected and there was full information about voters’ preferences. They showed that people generally diverge from truthful voting, and that over time, they learn to cast votes that are consistent with a single equilibrium. In a follow-up study, Bassi (2008) showed that people invoked different voting strategies depending on the voting rule implemented by the system. In particular, incorporating a simple plurality voting rule led people to adopt more strategic voting than when incorporating the Borda rule which was based on ranking the candidates. Our research extends these studies in two ways. First, we consider more complex settings in which the voting system outputs a ranking over the candidates, rather than a single winning candidate. Such settings occur frequently in the real world, yet people’s behavior in such voting systems has not been studied. We hypothesized that people do not play equilibrium strategies in these settings, and thus computer agents will need to adopt other types of voting strategies in order to succeed. Second, it provides a first study that compares the performance of computational strategies with people’s voting behavior.

There is significant work in economics on the design of voting systems in which agents submit total rankings over candidates (Dokow and Holzman 2010). Also there is scant work about modeling people’s behavior in voting settings, despite the growing literature in human-computer decision-making (Lin and Kraus 2010). A notable exception is the work by Mao et. al (2012) that compared the performance of several voting strategies for aggregating people’s ranking of solutions to optimization problems. They did not study the effect of computer agents using different voting strategies on people’s behavior.

Social Rankings

In this section, we describe how we adapted a popular voting system from the economics literature to be used in committees that include both humans and computer agents. We first provide the following definitions. Let N be a set of agents and C be a set of candidates. A ranking of C is a total order over the set C . Let L denote the set of all possible rankings of C . Each agent i has a preferred ranking $p_i \in L$ over C . A profile $p^N \in L^N$ is the set of preferred rankings for each agent in N . A *vote* of agent i is a ranking $v_i \in L$, and $v^N \in L^N$ denotes a set of votes for all agents in N . A *social welfare* function $f : L^N \rightarrow L$ provides a ranking $f(v^N) \in L$ for any $v^N \in L^N$.

A candidate pair $a, b \in C$ (w.l.o.g) is called an *issue*. Following notation introduced by Wilson (1975) we represent a ranking using a binary vector $\{0, 1\}^K$, where $K = \binom{|C|}{2}$ is the number of issues (all possible pairs in C). There exists a single corresponding entry in the vector for each issue that equals “1” if $a \succ b$ in the ranking and “0” if $b \succ a$. For example, consider a committee with N agents that needs to prioritize the following candidates for a budget: education (e), defense (d) and health (h). The first entry in the vector representing a ranking over the candidates will specify whether $e \succ d$; the second entry will specify whether $d \succ h$; and the third entry will specify whether $h \succ e$. For example, the vector (110) represents the ranking $e \succ d, d \succ h, e \succ h$.

The distance between two vectors v_1 and v_2 , denoted $d(v_1, v_2)$, is the Hamming distance between v_1 and v_2 . We extend this notion to provide a distance metric between a set of vectors v^N and vector v .

$$d(v^N, v) = \sum_{i \in N} d(v_i, v) \quad (1)$$

Social Welfare Rules

Let $f(v^N)$ represent the ranking that is chosen by applying the social welfare rule f to the set of votes v^N . We define the utility for agent i given $f(v^N)$ as reversely proportional to the distance between $f(v^N)$ and the agent’s preferred ranking p_i . We add a constant that is equal to the number of issues K to ensure that utilities are greater or equal to zero.

$$u_i(f(v^N)) = K - d(p_i, f(v^N)) \quad (2)$$

The set of all possible rankings for three candidates is $L = \{(001), (010), (100), (110), (101), (011)\}$. Importantly, any ranking of C can be represented as a vector of order K , but not all vectors of order K are rankings.

Specifically, for the 3-candidate example described above, (111) and (000) are the only vectors describing the cycles $e \succ d \succ h \succ e$ and $h \succ d \succ e \succ h$, respectively. They do not represent valid rankings and therefore are not in L .

A natural candidate for designing social welfare rules for human-computer settings is the majority method: choosing the value that agrees with the majority of agents' votes for each issue. There are several advantages to this rule: It fulfills canonical conditions of voting systems from the social choice literature, namely non-dictatorship, independence of irrelevant alternatives and pareto optimality (May 1952); it is the unique vector that maximizes agents' utilities; it is natural and intuitive to explain to people in the lab. However, the majority method may not produce a valid ranking for some voting profiles. For example, for the voting profile $v^N = \{(110), (011), (101)\}$ the majority method will produce the ranking (111) $\notin L$ which as we have shown above is not a valid ranking.

We therefore need an alternative method for combining agents' votes that preserves as many qualities of the majority method as possible, while still producing a valid voting rule. To this end, we will define the following set:

$$MIN_{v^N} = \{v \in L \mid \forall v' \in L, d(v^N, v) \leq d(v^N, v')\} \quad (3)$$

Intuitively, the set MIN_{v^N} includes those rankings in L that minimize the total distance to agents' votes v^N . For the voting profile v^N given in the above example, the set $MIN_{v^N} = \{(110), (101), (011)\}$ (this is because the distance between each of these rankings and v^N is 4, whereas the distance between the other rankings in L and v^N is 5).

Our Social Welfare Rule

We can now define a social welfare rule for our setting as a function \tilde{f} such that $\tilde{f}(v^N) \in MIN_{v^N}$ for any $v^N \in L^N$. This rule, called Kemeny-Young (1978; 1959), is a primary method for choosing a valid ranking given that agents' submits rankings over candidates. Computing the Kemeny-Young rule is an NP-Hard problem (Dwork et al. 2001) and recent work has proposed algorithms for computing bounds on this computation using search techniques (Conitzer, Davenport, and Kalagnanam 2006).² A particular advantage of using this method is that when the majority method outputs a valid ranking in L , MIN_{v^N} is a singleton and $\tilde{f}(v^N)$ reduces to the majority method. For the case where $|MIN_{v^N}| \geq 2$, we define $\tilde{f}(v^N)$ to equal the ranking in MIN_{v^N} that is first according to lexicographical order.³ For the set of agents' votes v^N in our example this social welfare rule will produce the ranking $\tilde{f}(v^N) = (011)$.

Voting Strategies

In this section we present and formalize several voting strategies. The most intuitive voting strategy for agents is to

²In practice, the computation of a Kemeny-Young rule was feasible for our setting, which included 4 candidates and 3 participants. For 3 participants, the Kemeny-Young rule is equivalent to using the Slater aggregation rule (Conitzer 2006).

³Other possibilities exist, like random. We chose an intuitive deterministic option.

vote according to their preferred rankings. We say that a vote v_i of agent i is *truthful* if v_i is equal to the agent's preferred ranking p_i . To illustrate, we extend the three-candidate example to include an additional candidate t (transportation). This is shown in Table 1, which lists the rankings of three voters over four candidates. (This was one of the preference profiles used in our empirical study that is described in the following section). When all agents vote truthfully, we have $v^N = p^N$ and the chosen ranking $\tilde{f}(v^N)$ assigns utilities 4, 4, 3 to agents 1, 2 and 3 (shown in the right-most column of the table).

It can be shown that for three candidates, no agent can do better than to vote according to its true preferences under this voting rule (Dokow and Holzman 2010). However, this is not the case in general. In fact, even for four candidates, players may be able to improve their outcome by deviating from their truthful vote. The situation in which an agent deviates from its true vote, that is $v_i \neq p_i$, is called *manipulation*. For the social welfare rule \tilde{f} , when $|C| = 4$, there exists a set of preferred rankings for which agents can improve their utility by manipulating their vote (Dokow and Holzman 2010). We illustrate using our example. Suppose agent 1 changes the value of issue (d, e) from 1 to 0 (with the values for all other issues staying the same) and agents 2 and 3 vote truthfully. (This manipulation is shown in parentheses in the first line of Table 1). In this case, the resulting rank \tilde{f} changes to the one shown in the last line of the table. As a result, agent 1 improves its utility to 5, while the other agents' utilities reduce to 3 and 2 (shown in parentheses in the last column of the table).

We now formalize an interesting set of voting behavior that differ in the "sophistication" of agents' reasoning about how other agents. Recall that v^{N-i} denotes the set of votes for all agents other than i . Given the social welfare rule f , and the set of votes v^{N-i} for all agents other than i , we define a set of *best-response* votes for agent i as follows:

$$BR_i(v^{N-i}) = \operatorname{argmax}_{v' \in L} u_i(f(v^{N-i}, v')) \quad (4)$$

Importantly, the best-response vote for agent i depends on the votes of all other agents $N \setminus \{-i\}$. We say that a vote for agent i is *Level-0*, denoted $v_i^{l_0}$ if it is a best-response for agent i given that all other agents vote truthfully, that is, $v_i^{l_0} \in BR_i(p^{N-i})$. The manipulative vote $d \succ e \succ h \succ t$ for agent 1 in the first line of Table 1 is level-0, because it maximizes its utility given that the other agents vote truthfully. Similarly, we say that a vote for agent i is *Level-1*, denoted $v_i^{l_1}$, if it is the best-response vote for i given that the other agents vote level-0, that is, $v_i \in BR_i((v_i^{l_0})^{N-i})$. For example, the level-1 vote for agent 3 is $h \succ d \succ t \succ e$. Lastly, a set of votes $v^N \in L^N$ is a *Nash equilibrium* for a social welfare rule f if-and-only-if for each agent i , it holds that

$$\forall v' \in L, u_i(f(v^{N-i}, v_i)) \geq u_i(f(v^{N-i}, v')) \quad (5)$$

In our example, the case in which agent 1 submits a truthful vote ($e \succ d \succ h \succ t$), agent 2 submits a level-0 vote ($t \succ e \succ d \succ h$), and agent 3 submits a truthful vote

		$e \succ d$	$d \succ h$	$h \succ e$	$e \succ t$	$d \succ t$	$h \succ t$	$u_i(\tilde{f}(v^N))$
$v_1 = p_1 (v_1^{t_0})$	$e \succ d \succ h \succ t \quad (d \succ e \succ h \succ t)$	1 (0)	1	0	1	1	1	4 (5)
$v_2 = p_2$	$e \succ t \succ d \succ h$	1	1	0	1	0	0	4 (3)
$v_3 = p_3$	$h \succ t \succ d \succ e$	0	0	1	0	0	1	3 (2)
$\tilde{f}(v^N)$	$e \succ h \succ t \succ d \quad (d \succ e \succ h \succ t)$	1 (0)	0 (1)	0	1	0 (1)	1	

Table 1: Truthful and strategic voting example for 3 agents and 4 candidates



Figure 1: Snapshot of the Budget Allocation Game

$(h \succ t \succ d \succ e)$ is Nash equilibrium for the social welfare rule \tilde{f} in which the chosen ranking is $t \succ e \succ d \succ h$. This profile incurs utilities of three points for agent 1, five points for agent 2 and two points for agent 3.

Having defined the set of voting strategies above, the natural question to ask is how agents should vote in mixed human-computer settings where the possibility of manipulation may increase participants' performance. We explore this question in the next section.

Empirical Methodology

To study people's voting behavior we designed a "budget allocation" game in which N agents vote to allocate a budget among a set of candidates C . Each agent is assigned a preferred ranking over the four candidates, and this information is common knowledge among all agents. The game comprises a finite number of rounds. In each round, all agents simultaneously submit a set of votes v^N . Each of these votes is a ranking over C (players do not actually propose a split of a monetary budget). The chosen ranking $\tilde{f}(v^N)$ is computed using the process defined in the previous section. Each agent's score is computed using Equation 2. Players can observe each other's votes for past rounds, as well as the chosen ranking and their respective scores. Agents' preferred ranking remain constant across rounds.

We implemented a version of the budget allocation game in which there are three players and four candidates: education, transportation, health and defense. A snapshot of the main game board is shown in Figure 1 from the point of view of "Player 1". The board shows the preferences of the three players in the game, as well as an editable ranking that player can modify and submit as its vote.

Rules of the Game

The budget allocation game is played repeatedly for five rounds. At the onset of the game, each player i is assigned a preference p_i over the candidates C . This information is common knowledge (all players can see each other's preferences as shown in Figure 1), and stays constant throughout the game. At each round, the three players simultaneously submit their votes $v^N = \{v_1, v_2, v_3\}$. The chosen ranking is computed according to \tilde{f} and each player incurs a score that is equal to its utility $u_i(\tilde{f}(v^N))$. Players have three minutes in which to submit their votes at each round (in practice, all subjects took well below 3 minutes to vote). If no vote is submitted, then a default vote is selected as follows. In the first round, the default vote for each player i is simply its preferred vote p_i . The default vote for each consecutive round is the ranking that the player submitted in the previous round. Once all players have submitted their rankings, the chosen ranking and scores are displayed to all of the players. In particular, all players can see each other's choices and incurred utilities. The bottom panel of Figure 1 shows the chosen ranking given that all players voted according to their preferred rankings. As shown by the Figure, the chosen ranking $\tilde{f}(v^N)$ is $e \succ h \succ t \succ d$.

There are two arguments for using this game to study human and computational voting strategies. First, four candidates is the smallest number for which manipulation may be beneficial, as we have described in the previous section. Second, the fact that players vote repeatedly allows them to adapt their voting behavior over time, and reflects real-world settings such as annual budget decisions and recurring elections.

Preference Profiles

As described above, players' scores for each round of voting depend on the extent to which the chosen ranking agrees with their preferred ranking that is assigned to them at the onset of the interactions. In real world voting scenarios, some players may be in better positions than others to affect the voting outcome. To reflect this we defined different power conditions between committee players by varying their assigned preference profile. Specifically, we used two preference profiles in the study that differed in the extent to which they allowed players to affect the voting result by deviating from their truthful vote.

In the profile called "symmetric", the preferred ranking of player 1 was $e \succ d \succ h \succ t$; the preferred ranking of player 2 was $e \succ t \succ d \succ h$; the preferred ranking of player 3 was $h \succ t \succ d \succ e$. These rankings are the same as the ones shown in Table 1, and are manifested in the game board

in Figure 1. If all players vote truthfully (we call this the “naive” voting baseline), player 3 is at a disadvantage, because the chosen ranking will be $e \succ h \succ t \succ d$, incurring a score of 4, 4, and 3 for players 1, 2 and 3, respectively. (the outcome is “symmetric” from the point of view of players 1 and 2). Moreover, the naive voting baseline is not stable, in the sense that player 1 and 2 can improve their score by voting strategically. Specifically, player 1 can improve its score by voting its level-0 strategy of $d \succ e \succ h \succ t$, given that other players vote truthfully. In this case, the scores will be 5, 3 and 2 for players 1, 2 and 3, respectively. In a similar way, player 2 can improve its score over the naive baseline by voting its level-0 strategy of $t \succ e \succ d \succ h$, given that the other players vote truthfully. In this case, the scores will be 4, 5 and 3 for players 1, 2 and 3, respectively. In fact, this voting profile in which player 2 deviates from its truthful vote, while player 1 and player 3 vote truthfully, is the Nash Equilibrium for this preference profile. We also used a profile called “asymmetric” in which the score of player 1 was higher than the score of player 2 and the score of player 3 if all players vote truthfully, but player 1 loses this advantage if player 2 votes level 0 and other agents vote truthfully. In our empirical study, this profile achieved similar results to the symmetric profile, which we omit for brevity.

Empirical Methodology and Results

We recruited 335 human subjects from the U.S. to play the game using Amazon Mechanical Turk. All participants were provided with an identical tutorial of how to play the budget allocation game, and their participation in the study was contingent on passing a quiz which tested their knowledge of the rules of the game. Participants were paid in a manner that was consistent with their performance, measured by accumulating their scores over five rounds of voting.

The subjects were randomly divided into three different groups. The first group consisted of people playing the budget allocation game with other people. The second group consisted of two people playing the game with another computer agent. The third group consisted of one person playing the game with two other computer agents. As a baseline, we also included a fourth group comprising solely computer agents. Each participant (both people and computers) played five rounds of the game. We hypothesized that (1) people will be less likely to vote truthfully, and more likely to play more sophisticated strategies; (2) that computer agents using best-response strategies (Equation 4) would be more successful when playing against people than computer agents that vote truthfully. All reported results in the upcoming section are significant in the $p < 0.05$ range using Analysis of Variance (ANOVA) tests.

Analysis of Human Behavior

We first present an analysis of people’s behavior in the game when playing with other people. In general, people’s strategy significantly deviated from the Nash equilibrium voting strategy. For example, in the symmetric profile, there were only 7 out of 80 rounds played in the 3-person group configuration in which a Nash equilibrium strategy was played,

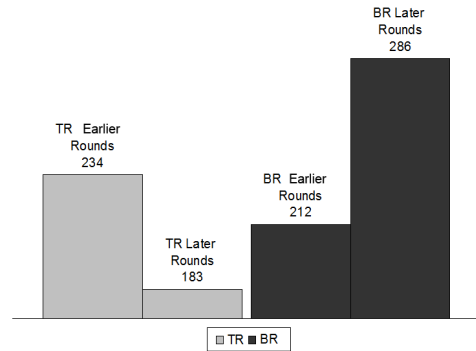


Figure 2: Difference in people’s Naive (TR) and Best-Response votes (BR) between earlier and later rounds in the game.

which is not significantly different than random. As a group, people’s voting behavior was noisy. Out of 80 rounds of the budget allocation game that were played by three people, 64 rounds included a unique set of vote combinations ($v^N = \langle v_1, v_2, v_3 \rangle$) that appeared only once. However, we did find two interesting trends in people’s behavior *as individuals*. Specifically, 40% of people’s votes were “naive” (votes that are truthful and consistent with their preferences), while 44% of people’s votes repeated their vote in the previous round. As we describe later in the section, this behavior was key in explaining the success of our computer agents.

To understand how people’s voting behavior evolved over time, we compared between the number of naive votes (TR) and best-response (BR) votes (best response to the votes of the other participants in the previous round) for different rounds of the game. Figure 2 shows the difference in the average number of naive and best-response votes for each role between rounds 4-5 and rounds 1-2 for games that included three people or two people and one computer agent. As shown in the figure, there was a drop in the number of naive votes for all players between earlier and later rounds in the game, confirming our hypothesis. In addition, the figure also shows an increase in the number of best-response votes between earlier and later rounds in the game. We conjecture that this is because participants learned to be more strategic about their voting behavior. However, this did not lead to improved performance, as there was no significant increase in people’s scores as rounds progressed. A possible reason for this is the large strategy space in each round (64 possible ranking profiles), making it challenging for people to predict others’ rankings when reasoning about how to manipulate. Interestingly, (and not shown by the figure) there was no increase in the number of best-response votes for people playing the role of Player 3 in the symmetric preference profile. We attribute this to the inherent disadvantage of this role in the game, in that it has a limited number of voting strategies that can improve its score (as we described in the previous section).

	PRBR	TR	People
2 People	3.9	3.78	3.40
Agents	3.79	3.12	3.09

Table 2: Performance of computer agents and people for different group structures

Type	Player 1	Player 2	Player 3
People	4.56	3.69	1.28
PRBR	4.87	4.04	2.78
TR	4.33	4.18	2.82

Table 3: Performance for different player roles in the symmetric preference profile

Agent-Design and Performance

We designed two types of computer agents playing deterministic voting strategies. The first agent, called *Previous Round Best Response* (PRBR), used the best-response vote of Equation 4 to rank the candidates, given that all other players repeat their vote in the previous round. That is, $v_i \in BR_i(v^{N-i})$ where v^{N-i} is the other agents’ votes in the prior round. In the first round, it is assumed that v^{N-i} equals p^{N-i} for all agents. The second agent, called *truthful* (TR), provided a baseline voting strategy that ranked all candidates according to its assigned preferences, that is $v_i = p_i$ at each round t given that i is a TR agent. We did not use a level-0 agent despite the fact that people were also likely to vote truthfully. This is because this voting strategy is static and easy to learn by people.

We first compare the performance of these computer agents and people in groups comprising two other people (that is, each game included a person or a computer agent voting with two other people). The first row in Table 2 shows the average performance of people and agents across all roles in the game for both preference profiles. As shown in the table, the PRBR agent was able to outperform the TR agent, and both PRBR and TR agents were able to outperform people. The second row of Table 2 shows the performance of computer agents and people in groups comprising two other computer agents (that is, each game included a person or computer agent voting with two other agents). As shown in the table, the PRBR agent also outperformed people and the TR agent in this additional group configuration. This demonstrates that the success of the best-response strategy was independent of the group structure.

To compare performance for different roles, we present Table 3 which compares performance for each role in groups comprising a computer or person interacting with two other people for the symmetric preference profile. As shown by the Table, the PRBR agent was significantly more successful than people in all player roles. In the role of Player 1, the PRBR agent was significantly more successful than the TR agent. Although the TR agent scored higher than the PRBR agent in both Player 2 and Player 3 roles, this difference was not statistically significant. We conclude that among the two agent strategies we evaluated, the PRBR agent was the best agent-design to play with people in our setting.

Player Role	Prev. round distance	true distance
1	0.86	1.1
2	1.18	1.51
3	1.51	1.70
avg.	1.18	1.43

Table 4: Distances between people’s votes in the game, their votes in the previous round, and their preferred rankings.

To induce an upper-bound on performance in the game, we computed a strategy for an oracle agent that could observe people’s actual votes in the game prior to submitting its own vote. The oracle strategy provides an upper bound for agents’ actual performance in the game. We found no significant difference between the performance of the PRBR agent and the oracle, for the data that we obtained of people’s behavior.

Lastly, we explain the success of the PRBR agent by computing the Hamming distance between people’s votes in consecutive rounds of the game (termed “Prev. round distance”) and the distance between their vote and the truthful vote (termed “true distance”). The distance for different players’ roles is shown in Table 4. (This distance can take values from 0 to 6, the number of possible issues.) As shown by the table, for all roles, participants’ previous round distance was smaller than their true distance. Because of this proximity, playing a BR to people’s votes in the previous round was a successful strategy for the agent.

To conclude, our results have implications for agent designers, suggesting that the PRBR strategy is sufficient towards enabling agents to perform well in voting systems which aggregate people’s rankings over candidates..

Conclusion

This paper described a first study comparing people’s voting strategies to that of computer agents in heterogeneous human-computer committees. In our setting participants vote by simultaneously submitting a ranking over the set of candidates and the election system uses the Kemeny-Young voting system to select a ranking that minimizes disagreements with participants’ votes. Our results show that over time, people learned to deviate from truthful voting strategies, and use more sophisticated voting strategies. A computer agent using a best response voting strategy to people’s actions in the previous round was able to outperform people in the game. In future work, we intend to design computer agents that adapt to people’s play in settings of incomplete information.

Acknowledgments

This work was supported in part by ERC grant #267523, Marie Curie reintegration grant 268362, the Google Inter-university center for Electronic Markets and Auctions, MURI grant number W911NF-08-1-0144 and ARO grants W911NF0910206 and W911NF1110344.

References

- Bassi, A. 2008. Voting systems and strategic manipulation: an experimental study. Technical report, mimeo.
- Conitzer, V.; Davenport, A.; and Kalagnanam, J. 2006. Improved bounds for computing kemeny rankings. In *Proc. of AAAI*.
- Conitzer, V. 2006. Computing slater rankings using similarities among candidates. In *Proc. of AAAI*.
- Cox, G. 1997. *Making votes count: strategic coordination in the world's electoral systems*, volume 7. Cambridge Univ Press.
- Dhillon, A., and Lockwood, B. 2004. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior* 46(1):55–75.
- Dokow, E., and Holzman, R. 2010. Aggregation of binary evaluations. *Journal of Economic Theory* 145(2):495–511.
- Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proceedings of the 10th international conference on World Wide Web*.
- Ephrati, E.; Rosenschein, J.; et al. 1993. Multi-agent planning as a dynamic search for social consensus. In *Proc. of IJCAI*.
- Forsythe, R.; Rietz, T.; Myerson, R.; and Weber, R. 1996. An experimental study of voting rules and polls in three-candidate elections. *International Journal of Game Theory* 25(3):355–383.
- Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica: Journal of the Econometric Society* 665–681.
- Gohar, N. 2012. *Manipulative Voting Dynamics*. Ph.D. Dissertation, University of Liverpool.
- Kemeny, J. 1959. Mathematics without numbers. *Daedalus* 88(4):577–591.
- Lin, R., and Kraus, S. 2010. Can automated agents proficiently negotiate with humans? *Communications of the ACM* 53(1):78–88.
- Mao, A.; Procaccia, A.; and Chen, Y. 2012. Social choice for human computation. In *HCOMP-12: Proc. 4th Human Computation Workshop*.
- May, K. 1952. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society* 680–684.
- Meir, R.; Polukarov, M.; Rosenschein, J.; and Jennings, N. 2010. Convergence to equilibria in plurality voting. In *Proceedings of AAAI*.
- Palfrey, T. 2009. Laboratory experiments in political economy. *Annual Review of Political Science* 12:379–388.
- Pennock, D.; Horvitz, E.; Giles, C.; et al. 2000. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proc. of AAAI*.
- Regenwetter, M., and Rykhlevskaia, E. 2007. A general concept of scoring rules: general definitions, statistical inference, and empirical illustrations. *Social Choice and Welfare* 29(2):211–228.
- Reijngoud, A., and Endriss, U. 2012. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*.
- Riker, W., and Ordeshook, P. 1968. A theory of the calculus of voting. *The American Political Science Review* 62(1):25–42.
- Satterthwaite, M. 1975. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* 10(2):187–217.
- Wilson, R. 1975. On the theory of aggregation. *Journal of Economic Theory* 10(1):89–99.
- Young, H., and Levenglick, A. 1978. A consistent extension of condorcet's election principle. *SIAM Journal on Applied Mathematics* 35(2):285–300.