Argumentation-Based Information Exchange in Prediction Markets

Santi Ontañón[1] and Enric Plaza[2]

[1] CCL, Cognitive Computing Lab Georgia Institute of Technology,
Atlanta, GA 303322/0280
`santi@cc.gatech.edu`
[2] IIIA, Artificial Intelligence Research Institute - CSIC, Spanish Council for Scientific Research
Campus UAB, 08193 Bellaterra, Catalonia, Spain
`enric@iiia.csic.es`

Abstract. The purpose of this paper is to investigate how argumentation processes among a group of agents may affect the outcome of group judgments. In particular we will focus on prediction markets (also called information markets) and we will investigate how the existence of social networks (that allow agents to argue with one another to improve their individual predictions) effect on group judgments. Social networks allow agents to exchange information about the group judgment by arguing about the most likely choice based on their individual experience. We develop an argumentation-based deliberation process by which the agents acquire new and relevant information. Finally, we experimentally assess how different social network connectivity and different data distribution affect group judgment.

1 Introduction

The purpose of this paper is to investigate how argumentation processes among a group of agents may affect the outcome of group judgments. In particular we will focus on prediction markets (also called information markets) and we will investigate how the existence of social networks (that allow agents to argue with one another to improve their individual predictions) effect on group judgments using prediction markets.

There are different ways to aggregate the information held by a group of agents. According to C. R. Sunstein [17] there are three main paradigms to achieve *group judgments*, that is to say a joint decision or prediction based on aggregating the information or preferences of a group of agents (Sunstein deals with human agents, while we will focus only on artificial software agents). One paradigm is using statistical means to aggregate the group information: techniques like plurality voting, Condorcet voting or weighted voting define aggregation functions based on statistical means (i.e. on diminishing the joint error). Human committees, panels and juries use these techniques — and groups of agents also, see for example [11] where learning agents' joint predictions are compared when using plurality voting vs. weighted voting.

A second paradigm is that of deliberation, where arguments in favor or against a joint judgment are exchanged by the member agents of a group. Human public and private institutions traditionally favor deliberative ways of taking decisions, and certain

I. Rahwan and P. Moraitis (Eds.): ArgMas 2008, LNAI 5384, pp. 181–196, 2009.

accounts of democracy are based on the deliberation process. The main feature here is that rough preferences are not considered sufficient to justify a joint judgment, and deliberation provides reasons by an exchange of *arguments* by individuals with different information and diverse perspectives. Agents can also use argumentation to deliberate on joint judgments, as for example in the work reported in [13].

The third paradigm is the one this paper focuses on: *prediction markets*, also known as *information markets*. Prediction markets' goal is to aggregate information based on a *price signal* emitted by the members of a group. The advantage of the price signal is that it encapsulates both the information and the preferences of a number of individuals. In this approach, the task of aggregating information is achieved by *creating a market*, and that market should offer the right *incentives* for the participating people or agents to disclose the information they possess.

The purpose of this paper is to analyze the effect of social network relationships in group judgment —specifically in prediction markets. These social networks allow agents to exchange information about the prediction task domain. We model this information exchange as an argumentation process, where an agent A tells an agent A' its prediction S together with an argument α intended to justify why this prediction is correct. Agent A can agree or disagree with S, and in the case of disagreement A' communicates to A a counterargument or a counterexample that contradicts α. Agent A may keep its original prediction S or change it to some new prediction S' due to the counterarguments and counterexamples A has exchanged with one or more other agents. Social networks establish the different possible graphs of trusted acquaintances with which an agent can soundly exchange information; several simple social networks are tested in order to analyze the impact of information exchange.

The structure of the paper is as follows: the next section describes the Multiagent Prediction Market (MPM) and discusses the assumptions to use such mechanism for group judgment; section 3 describes the argumentation processes among agents that models the information exchange among agents; then section 4 presents an empirical evaluation of MPM in a prediction domain and we assess (1) the effect of using a prediction market instead of a voting scheme, and (2) the effect upon prediction markets of information exchange. Finally, section 5 presents related work and section 6 discusses the contributions of the paper and the foreseeable future work.

2 Multiagent Prediction Market

Essentially, a Multiagent Prediction Market (MPM) is composed of (a) a *prediction task domain*, (b) a market broker agent A_D, (c) a collection of participating agents $\mathcal{A}$, and two parameters: M (maximum bet) and X (a percentage bonus).

In this paper we will address only single-issue predictions and we will assume that the prediction task domain is characterized by an enumerated collection of *alternatives* or *solutions* $\mathcal{S} = \{S_1, ..., S_K\}$ and the prediction task is to select the correct one for the current *situation* or *problem* P. The participating agents is a multiagent system composed of n agents $\mathcal{A} = \{A_1, ..., A_n\}$. For a specific market, given a *problem* P every agent receives P, generates its individual prediction, and then it can bet up to a quantity M_P on one single alternative.

Let $B_{A_i} = \langle S, b \rangle$ be the bet made by a particular agent A_i, where S is the predicted solution, and b is the amount bet. Let $\mathcal{B}_{\mathcal{P}} = \{B_{A_i}, ..., B_{A_n}\}$ be the set of all bets made by all the agents in the market MPM_P. We will use the dot notation to refer to elements inside a tuple, e.g. we will write $B.b$ to refer to the amount bet in B. We define $B_P = \sum_{B \in \mathcal{B}} B.b$ as the total amount of money bet by all the agents, and $B_{S_k} = \sum_{B \in \mathcal{B} | B.S = S_k} B.b$ the total amount of money bet for a particular solution S_k.

The broker agent A_D receives those bets (amounting to a total quantity B_P) and determines the joint prediction as the alternative (say S_r) invested with the highest accumulated bet, as follows: $S_r = \arg\max_{S_k \in \mathcal{S}} B_{S_k}$. When the correct solution S_c of P becomes known, the broker agent A_D checks whether the joint prediction was accurate ($S_r = S_c$). If it was, then those agents that bet for S_c receive a reward. Specifically, an agent A_i who bet for the correct solution receive the reward $r_{A_i} = \frac{1}{B_S}(B_P \times B_{A_i}.b \times c)$, where $c = \frac{100+X}{100}$ is a factor that ensures that the agents receive more money than they bet if they win. Intuitively, the winner agents receive all the money bet by all the agents (i.e. B_P), but multiplied by the factor c, to provide an incentive. In our experiments we have set the percentage bonus $X = 10\%$, thus, $c = 1.1$.

The rationale of this design is to provide a twofold incentive: a) for the agents to reveal their true prediction, and b) also to benefit from the the joint accuracy.

Concerning the participating agents, we make the assumptions that (1) the individual agents possess a way to determine the confidence in an individual prediction and (2) the agents possess an argumentative capability that supports the information exchange with other agents regarding the prediction task domain. The first assumption requires that the agent is not only capable of making a prediction, but also establishing the likelihood of that specific prediction to be correct, i.e. a degree of confidence for each specific prediction. Rationality dictates that the more confident an agent with respect to a prediction, the higher the quantity to bet on that prediction. The second assumption allows the agents to perform an information exchange phase (that we model as an argumentation process), and thus generate more informed predictions.

3 Information Exchange in Social Networks

Social networks views social structures as composed of nodes and links, where nodes are individuals or organizations and links are their relationships. For the purpose of this paper, we will focus on individual agents as nodes and acquaintances as their links.

In our framework, a social network is a collection of *acquaintance* directional relations $N = \{(A_{i_1}, A_{j_1}), ..., (A_{i_m}, A_{j_m})\}$, where an agent A_i has another agent A_j as an acquaintance only if $(A_i, A_j) \in N$. Figure 1 shows three examples of social networks: In the leftmost one, each agent has one acquaintance, in the middle one, each agent has two acquaintances, and in the rightmost one each agent has three acquaintances.

Before declaring a prediction on the market, an agent A_i will first try to exchange information with its acquaintances. Thus, A_i will engage in argumentation processes about the correct solution of the problem at hand with each of its acquaintances before making a prediction — following the argumentation formalism we introduced in [13].

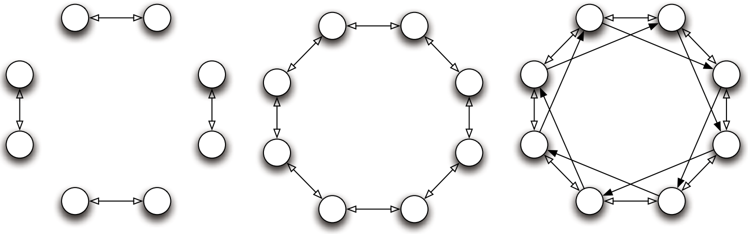

Fig. 1. Three of social networks among 8 agents where each agent has 1, 2 or 3 acquaintances

3.1 Problem-Centered Information Exchange as Argumentation

An agent A_i can obtain new information concerning the solution of a problem P by engaging in an argumentation process with another agent A_j, that might have information unknown to A_i. During an argumentation process, two agents exchange information concerning the solution of a specific problem P. Specifically, an agent may generate an argument in favor of a particular solution and send it to the other agent. Agents can also analyze a received argument, and agree or disagree with it. When an agent disagrees with an argument, it might generate a counterargument or a counterexample. By exchanging arguments and counterarguments two agents may reach a consensus about which is the most plausible solution for a given problem taking into account the information that both of them have. Therefore, the individual solution reached after an argumentation process is in principle more informed, and thus more likely to be correct.

3.2 MPM with CBR Agents

In our framework, each agent uses Case-Based Reasoning (CBR) [1] in order to generate predictions. Thus, each agent A_i owns a case base C_i, composed of a collection of cases, $C_i = \{c_1, ..., c_m\}$. A case is a tuple $c = \langle P, S \rangle$ containing a case description P and a solution $S \in \mathcal{S}$. We will use the terms *problem* and *case description* indistinctly.

CBR agents can solve problems by themselves, using CBR problem solving methods. Moreover, agents can also try to obtain information from other agents in order to increase their prediction accuracy. In a prediction market, given that each individual agent is interested in maximizing its prediction accuracy (in order to obtain a higher reward), it is rational for an agent to try to obtain the maximum information possible from other agents before making its prediction.

Argumentation provides a formal and well founded way to problem-centered information exchange. We will next summarize the case-based approach to multiagent argumentation introduced in [13]: the kind of arguments and counterarguments supported, how CBR agents generate arguments, and how agents compare arguments. Finally, we will present a specific argumentation protocol for information exchange in prediction markets, that agents can use to increase the accuracy of their predictions.

3.3 Arguments and Counterarguments

For our purposes an *argument* α generated by an agent A is composed of a statement S and some information D endorsing the fact that S is correct. In the context of CBR

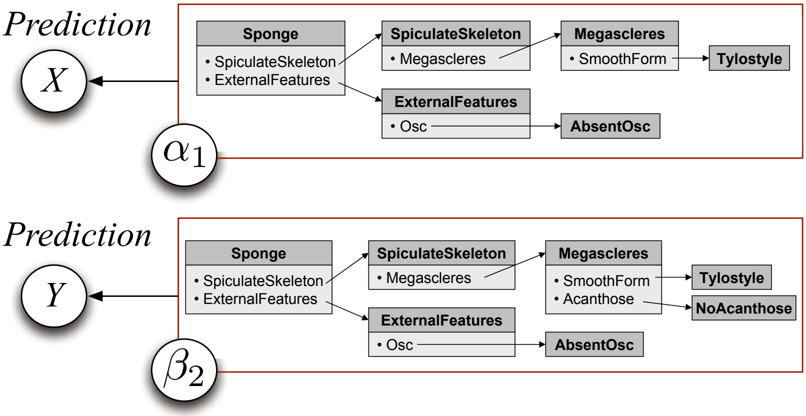

Fig. 2. Relationship between two arguments:β_2 is a counterargument of α_1 because β_2 is a refinement of α_1 and predicts Y that is different from α_1's prediction X

agents, agents argue about predictions for new problems and can provide two kinds of information: a) specific cases $\langle P, S\rangle$, and b) justified predictions: $\langle A, P, S, D\rangle$. Using this information, we can define three types of arguments: justified predictions, counterarguments, and counterexamples.

A *justified prediction* α is generated by an agent A_i to argue that A_i believes that $\alpha.S$ is the correct solution for problem P because of justification $\alpha.D$.

A *counterargument* β is an argument offered in opposition to another argument α. In our framework, a counterargument consists of a justified prediction $\langle A_j, P, S', D'\rangle$ generated by an agent A_j with the intention to rebut an argument α generated by another agent A_i, that endorses a different solution S' with a justification D'.

Figure 2 shows two arguments from our experimental setting in section 6. First notice that each argument is predicting a different solution: α_1 predicts X while β_2 predicts Y. Moreover, α_1 subsumes β_2 (in other words, β_2 is a specialization of α_1), meaning that all problems that satisfy β_2 also satisfy α_1. If the predictions are contradictory ($X \neq Y$) then β_2 is a counterargument of α_1.

A *counterexample* c is a case that contradicts an argument α. Thus, a counterexample is also a counterargument, stating that an argument α is not always true, and the evidence provided is the case c. Specifically, a case c is a counterexample of an argument α if the following conditions hold: $\alpha.D \sqsubseteq c$ and $\alpha.S \neq c.S$, i.e. the case satisfies the justification $\alpha.D$ while determining a solution different to than the predicted by α.

3.4 Argument Generation

In our framework, arguments are generated by the agents from cases, using learning methods. Any learning method able to provide a justified prediction can be used to generate arguments. For instance, decision trees and LID [4] are suitable learning methods. Specifically, in the experiments reported in this paper agents use LID. Thus, when an agent wants to generate an argument endorsing that a specific solution is the correct solution for a problem P, it generates a justified prediction using LID.

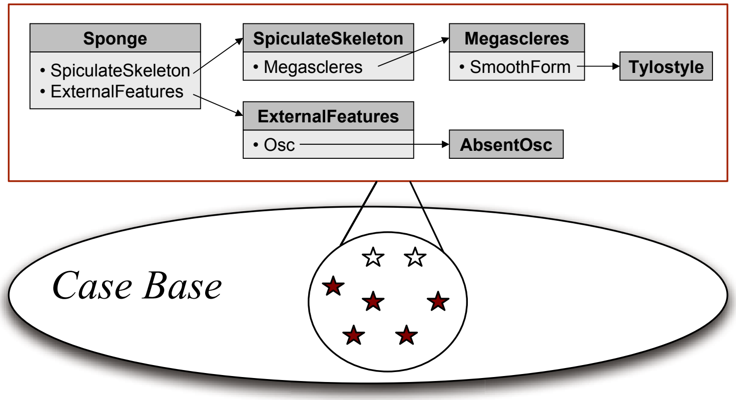

Fig. 3. Relationship between an argument and a case base. Dark stars are cases endorsing the argument while white stars are cases contradicting it.

Agents may try to rebut arguments by generating a counterargument or by finding counterexamples. An agent A_i wants to generate a counterargument β to rebut an argument α when α is in contradiction with the local case base of A_i. Moreover, while generating such a counterargument β, A_i expects that β is preferred over α. For that purpose, agents use a specific policy to generate counterarguments based on the *specificity* criterion [14]. The generation of counterarguments using the specificity criterion puts some requirements on the learning method but techniques LID or ID3 can be easily adapted for this task (as shown in [13]).

For instance, in Figure 2, given an argument α_1 that predicts X asserted by agent A_1 generating a counterargument means that agent A_2 finds a description β_2 such that it is subsumed by α_1 but (according to A_2's experience) predicts a solution $Y \neq X$.

Specifically, in our experiments, when an agent A_i wants to rebut an argument α, uses the following policy: (1) Agent A_i tries to generate a counterargument β more specific than α; if found, β is sent to the other agent as a counterargument of α. If not found, then (2) A_i searches for a counterexample $c \in C_i$ of α. If a case c is found, then c is sent to the other agent as a counterexample of α. If an agent A_i is unable to generate a counterargument or find a counterexample then A_i has no grounds to disagree with argument α and can not rebut that argument.

3.5 Prediction Confidence

We will use a *case-based confidence* measure [13] to determine the degree of confidence of an individual agent in its own argument (justified prediction) and also on the counterarguments received from other agents. The confidence is assessed by the agents via an process of *examination of arguments*. During this examination, an agent will count how many of the cases in its individual case base *endorse* an argument α, and how many cases are counterexamples of α. The more endorsing cases, the higher the confidence; and the more the counterexamples, the lower the confidence.

While examining an argument α, an agent determines the set of cases in its individual case base that are subsumed by $\alpha.D$ (the cases shown as stars in the circle of

Figure 3): the more of these cases that have $\alpha.S$ as solution, the higher the confidence. After examining an argument α, an agent A_i obtains the *aye* and *nay* values: The *aye* value $Y_\alpha^{A_i} = |\{c \in C_i|\ \alpha.D \sqsubseteq c.P \wedge \alpha.S = c.S\}|$ is the number of cases in the agent's case base *subsumed* by the description $\alpha.D$ that has solution $\alpha.S$ proposed by α, while the *nay* value $N_\alpha^{A_i} = |\{c \in C_i|\ \alpha.D \sqsubseteq c.P \wedge \alpha.S \neq c.S\}|$ is the number of cases in the agent's case base *subsumed* by description $\alpha.D$ that *do not* have that solution.

Figure 3 shows an the examination process where, given an argument α, an agent first retrieves all the cases that are subsumed by $\alpha.D$ from the case base, and then counts how many are counterexamples (white stars) or endorsing cases (black stars).

The confidence on an argument α is assessed by an agent A_i as follows:

$$\mathsf{C}_{A_i}(\alpha) = \frac{Y_\alpha^{A_i} + 1}{Y_\alpha^{A_i} + N_\alpha^{A_i} + 2}$$

where the reason for adding 1 to the numerator and 2 to the denominator is akin to the Laplace correction to estimate probabilities.

3.6 Information Exchange Protocol

In this section we will define an information exchange protocol that allows agents in an information market to exchange information with its acquaintances in the social network. Intuitively, an agent will engage into one-to-one argumentation processes with each one of his acquaintances sequentially, trying to improve its prediction at each step. The intuition is that after each discussion, the solution is more likely to be the correct one, since more information has been taken into account to come up with it.

Let us assume that a particular agent A_i wants to generate a prediction for a problem P. Let $F \subseteq \mathcal{A}$ be the set of m acquaintances of A_i. The information exchange protocol initiates a series of argumentation processes between A_i and each of the agents in F in a series of rounds. In the first round $r = 0$, A_i simply generates its individual prediction in the form of an argument α^0. Then, in the next round $r = 1$, A_i will argue with the first agent $A_j \in F$ and refine its prediction into a better one α^1. At the end of round $r = m$, A_i will have a prediction α^m that will be the final one made for the market.

Each one of these argumentation processes in itself consists of a series of cycles. In the initial cycle, each agent states which is its individual prediction for P. Then, at each cycle an agent can try to rebut the prediction made by the other agent. The agents alternate turns in the protocol, and an agent is allowed to send one counterargument or counterexample at its turn. When an agent receives a counterargument or counterexample, it informs the other agent if it accepts the counterargument (and changes its prediction) or not. Moreover, agents have also the opportunity to answer to counterarguments in their turn, by trying to generate a counterargument to the counterargument. At any time the protocol terminates when all the agents agree or when no agent has generated any counterargument during the last two cycles.

During the argumentation protocol, agents can use the following performatives:

- $assert(\alpha)$: the justified prediction held during the next cycle will be α. If multiple asserts are send, only the last one is considered as the currently held prediction.
- $rebut(\beta, \alpha)$: the agent has found a counterargument β to the prediction α.

We will define α_i^t as the prediction that an agent A_i is holding at iteration t of the argumentation protocol, and H_t as the set containing the predictions that each of the two agents hold at a cycle t. The argumentation protocol between an agent A_i, that is currently holding a prediction α_r at a round r of the information exchange protocol, and an acquaintance A_j works as follows:

1. At cycle $t = 0$, the initial argument of A_i will be the one coming from the previous round α_r, thus $\alpha_i^0 = \alpha_r$. The initial argument of A_j will be the result of trying to solve P individually, building a justified prediction using its own CBR method. Then, each agent sends the performatives $assert(\alpha_i^0)$ and $assert(\alpha_j^0)$ respectively to the other agent. Thus, the agents know $H_0 = \langle \alpha_i^0, \alpha_j^0 \rangle$. The turn is given to the first agent A_i.
2. At each cycle t (other than 0), the agents check whether their arguments in H_t agree. If they do, the protocol moves to step 5. If during the last 2 cycles no agent has sent any counterexample or counterargument, the protocol also moves to step 5. Otherwise, the agent A_i who has the turn tries to generate β_i^t (a counterargument or a counterexample) against the argument of the other agent:
   - If β_i^t is a counterargument, then, A_i locally compares α_i^t with β_i^t by assessing their confidence against its individual case base C_i (notice that A_i is comparing its previous argument with the counterargument that A_i itself has just generated and that is about to send to A_j). If $C_{A_i}(\beta_i^t) > C_{A_i}(\alpha_i^t)$, then A_i considers that β_i^t is stronger than its previous argument, changes its argument to β_i^t by sending $assert(\beta_i^t)$ to the rest of the agents (i.e. A_i checks if the new counterargument is a better argument than the one it was previously holding) and $rebut(\beta_i^t, \alpha_j^t)$ to A_j. Otherwise (i.e. $C_{A_i}(\beta_i^t) \leq C_{A_i}(\alpha_i^t)$), A_i will send only $rebut(\beta_i^t, \alpha_j^t)$ to A_j. In any of the two situations the protocol moves to step 3.
   - If β_i^t is a counterexample c, then A_i sends $rebut(c, \alpha_j^t)$ to A_j. The protocol moves to step 4.
   - If A_i cannot generate any counterargument or counterexample, the turn is given to the next agent, a new cycle $t + 1$ starts, and the protocol moves to state 2.
3. The agent A_j that has received the counterargument β_i^t, locally compares it against its own argument, α_j^t, by locally assessing their confidence. If $C_{A_j}(\beta_i^t) > C_{A_j}(\alpha_j^t)$, then A_j will accept the counterargument as stronger than its own argument, and it will send $assert(\beta_i^t)$ to the other agent. Otherwise (i.e. $C_{A_j}(\beta_i^t) \leq C_{A_j}(\alpha_j^t)$), A_j will not accept the counterargument, and will inform the other agent accordingly. Any of the two situations start a new cycle $t+1$, A_i gives the turn to the next agent, and the protocol moves to state 2.
4. The agent A_j that has received the counterexample c retains it into its case base and generates a new argument α_j^{t+1} that takes into account c, and informs the rest of the agents by sending $assert(\alpha_j^{t+1})$ to all of them. Then, A_i gives the turn to the other agent, a new cycle $t + 1$ starts, and the protocol moves to step 2.
5. The argument that A_i is holding is the one that will be carried on to the next round of the information exchange protocol, i.e. when A_i engages in an argumentation with the next agent out of his acquaintances.

Moreover, in order to avoid infinite iterations, if an agent sends twice the same argument or counterargument to the same agent, the message is not considered.

3.7 Bet Generation

At the end of the information exchange protocol, an agent A_i will have a prediction α for a particular solution class. Moreover, in order to participate in a prediction market, the agent has to bet a particular amount of money on its prediction. The more money the agent bets, the bigger the potential reward is, but the bigger the risk. Thus, it is natural for an agent to bet more money when it is more confident that its prediction is correct. For that reason, in our framework, agents bet money proportionally to the confidence (computed as explained in Section 3.5) on their predictions. Since an MPM defines a maximum amount of money M that each agent can bet, each agent will bet $M \times C(\alpha)$, i.e. a proportional amount to its individual confidence. Thus, the bet made by an agent A_i that has a prediction α after the information exchange process will be:

$$B_{A_i} = \langle \alpha.S, M \times C(\alpha) \rangle$$

4 Experimental Evaluation

In this section we will empirically evaluate the performance of prediction markets, comparing it to the performance of normal voting. Moreover we will also study the effect of having different social networks among the agents in the market and how much the quality of data affects the market.

We have made experiments in the sponge data set, a marine sponge identification tasks that contains 280 marine sponges represented in a relational way and pertaining to three different orders of the Demospongiae class. In an experimental run, training cases are distributed among the agents. In the testing stage problems arrive to the market, and each agent will place a bet for the solution they predict is the correct one.

We have performed three sets of experiments. In the first set, we are interested in comparing prediction markets with majority voting, in the second one we want to explore the effect of argumentative information exchange in prediction markets, and finally, the third one explores the effect of varying the quality of the data sample that each agent owns. Each experiment consists of 5 runs of a 5-fold cross validation test. Notice that in step 4 of the argumentation protocol in section 3.6, agents learn from counterexamples coming from other agents. In the experiments we performed, each problem in the test set has to be independent from one another, in order to compute the averages for cross validation. Thus, the learning performed during argumentation is not carried up to the next problem in the test set. We have researched the issue of *learning from communication* in other multiagent scenarios in [12].

4.1 Prediction Markets Versus Majority Voting

For these experiments we evaluated the prediction accuracy of a committee using majority voting consisting of 8 agents with a prediction market consisting of the same 8 agents. The training set is split into 8 parts and each part is sent to an agent. Thus, each agent has an initial case base of about 28 cases.

Agents solving problems using a prediction market didn't do any information exchange for this experiment. The maximum bet was set to $M = 100$, and the incentive

Table 1. Prediction markets accuracy with information exchange along several social networks and with different biases in the individual case bases

social network	*market accuracy*	*individual accuracy*	*average reward*	*majority voting*
0 acquaintances	*89.71%*	*74.21%*	*10.35*	
1 acquaintances	*90.57%*	*83.99%*	*11.42*	
2 acquaintances	*91.29%*	*86.63%*	*12.14*	89.71
3 acquaintances	*91.14%*	*87.64%*	*11.94*	0.20
4 acquaintances	*91.07%*	*88.16%*	*11.85*	0.21

factor was set to $X = 10\%$, thus $c = 1.1$. The results showed that the majority voting achieved a prediction accuracy of 88.93%, while the prediction market achieved an accuracy of 89.71%, a significant improvement. Moreover, agents won an average of 10.35 monetary units per problem solved. In a voting committee, agents are only asked to reveal part of its individual information, namely the preferred alternative for which an individual casts a vote. In a prediction market, however, the amount bet by an individual acts as a "signal" indicating the degree of individual confidence in predicting the preferred alternative as being the correct one. Since the reward is proportional to the bet amount, the agents have an incentive to disclose this additional information.

Since the reward is proportional to the individual prediction confidence, the agents have an incentive to try to improve their individual prediction accuracy and confidence.

4.2 The Effect of Information Exchange

We performed several experiments with different social networks in a prediction market composed of 8 agents. Figure 1 shows some social networks where each agent has 0, 1, 2 or 3 acquaintances; we have performed experiments with 0 to 4 acquaintances and logged the prediction accuracy of the market, the prediction accuracy of each individual agent, and also the average money reward received by each agent per problem.

Table 1 shows that information exchange is positive both for the individual agents and for the market as a whole. We can see that the more acquaintances an agent has, the higher its individual prediction. For instance, agents with 0 acquaintances have an accuracy of 74.21% while agents with 1 acquaintance have an accuracy of 83.99%, and when they have 4 acquaintances, their accuracy is increased to 88.16%. Moreover, the predictive accuracy of the market increases from 89.71% when agents do not perform information exchange, to above 91% when agents have more than 1 acquaintance.

These results also show that the argumentation process of section 3.6 is successful in in acquiring individually valuable information. The increase in individual accuracy and confidence in prediction can only be explained by agents changing their original prediction and confidence value after arguing with other agents.

Another effect we can observe is that the reward that the agents obtain increases when they perform information exchange, starting in 10.35 monetary units per problem when they do not perform information exchange, and going up to close to 12 when agents have 2 acquaintances ore more. It is interesting to notice that the performance of the prediction market doesn't increase linearly with the performance of the individual agents. In fact, the more accurate the individual agents get, the more correlated their individual predictions are, and thus there is less difference between their individual predictions and the prediction of the market as a whole. This is a well known effect in machine learning (known as the *ensemble effect* [9]), or in economics (related to the Condorcet Jury Theorem). Therefore, if the reward signal that the agents get was only related to its individual accuracy, agents might be interested in their classification accuracy to a point were the correlation is too high, and then the market would not achieve it's optimal accuracy. The reward signal presented in Section 2 takes this into account, and rewards the agents when the market as a whole has high accuracy.

Moreover, Table 1 shows that the reward signal is higher when the market accuracy is higher (in our experiments, when agents have 2 acquaintances), instead of when their individual accuracy is higher. Therefore, the agents have an incentive to be highly accurate, but up to a limit, so that the market as a whole has a high accuracy. In our experiments, the agents receive maximum reward when they collaborate with two acquaintances, and thus it is rational for the agents to do so. As a side effect, the accuracy of the market as a whole is also maximum under those conditions, thus the agents have an incentive to do what is better for the market.

Summarizing, the experiments show that prediction markets can provide incentives for agents to disclose more information, and that information improves the accuracy of joint predictions or group judgments. The MPM is based on disclosing further information interpreted as a bet amount that represents the individual confidence on a prediction. The results also show that the case-based confidence function defined in Section 3.5 provides a good estimation, since the prediction market improves the accuracy.

Concerning information exchange, the experiments show that individual and market accuracy improve. This means that the agents make a more informed prediction, and thus that the argumentation protocol of Section 3.6 is effective in providing agents with enough information to correct previously inaccurate predictions.

4.3 Quality of the Data Sample

The results in the previous section assume that each agent have a good sample of data, i.e. that each agent is competent. We performed a set of experiments where we changed the quality of the data sample that each agent has and evaluated how this affects the performance of the market, as well as the individual agents.

Specifically, we performed experiments where agent have *biased* case bases. A biased case base is one that is not a good sample of the complete data set. The bias of a case base C_i with respect to a data set C is defined by:

$$\mathbb{B}(C_i) = \sqrt{\sum_{k=1 \ldots K} \left(\frac{\#(\{c \in C_i | c.S = S_k\})}{\#(C_i)} - \frac{\#(\{c \in C | c.S = S_k\})}{\#(C)} \right)^2}$$

Notice that Case Base Bias is zero when the ratio of cases for each solution class is the same in the case base C_i than in the data set C. The higher the bias, the worst the sample.

Table 2. Prediction markets accuracy with information exchange along several social networks and with different biases in the individual case bases

bias	social network	market accuracy	individual accuracy	average reward	majority voting
0.2	0 acquaintances	90.14%	73.01%	10.44	89.00
	1 acquaintances	89.86%	82.79%	10.29	
	2 acquaintances	90.07%	85.80%	10.53	
	3 acquaintances	91.21%	87.16%	11.54	
	4 acquaintances	91.36%	87.49%	11.79	
0.4	0 acquaintances	88.71%	66.43%	8.86	84.86
	1 acquaintances	87.79%	75.95%	7.43	
	2 acquaintances	89.43%	79.56%	8.63	
	3 acquaintances	90.57%	81.7%	9.59	
	4 acquaintances	90.79%	82.16%	9.84	
0.6	0 acquaintances	86.29%	58.05%	6.7	83.29
	1 acquaintances	88.00%	70.00%	6.75	
	2 acquaintances	90.14%	74.94%	8.05	
	3 acquaintances	89.36%	74.48%	7.33	
	4 acquaintances	89.21%	75.51%	7.21	
0.6	0 acquaintances	49.71%	38.63%	-32.36	55.57
	1 acquaintances	48.00%	41.26%	-23.64	
	2 acquaintances	61.43%	47.52%	-16.55	
	3 acquaintances	67.43%	55.67%	-0.02	
	4 acquaintances	66.93%	55.44%	-0.53	

We performed experiments giving agents case bases with bias 0.2, 0.4, 0.6, and 0.8 (results in Table 1 correspond to bias equal to 0.0). Table 2 shows the accuracy of the market, of the individual agents, and also of majority voting. If we look at the accuracy of majority voting, we can see that it degrades when the bias increases, since the individual agents' predictions also degrades. For instance, the classification accuracy of majority voting degrades from 88.93% when the bias is 0.0 to 83.29% when the bias is 0.6 or to 55.57% when the bias is 0.8 (0.8 is a really large bias in our sponge data set, and each agent almost only knows cases of a single class).

Concerning market's accuracy, increasing the bias also diminishes its accuracy, but remarkably much less than for majority voting; in fact market's accuracy is strongly affected only when the bias is very high. For instance, a prediction market where agents have 4 acquaintances with bias 0.6 has still an accuracy of 89.21% (compared to 83.29% of majority voting). The accuracy of individual agents is much more affected by bias, being reduced from 74.21% with no bias to 58.05% with bias 0.6 when they have no acquaintances. When the bias is even larger (0.8), their accuracy goes down drastically to 38.63%. As the number of acquaintances increases, the individual accuracy largely increases, showing that the argumentation framework allows agents to efficiently exchange information and benefit from information in the case bases of acquaintances. For instance, the individual accuracy with bias 0.6 recovers from 58.05% with no acquaintances to 75.51% with 4 acquaintances. This effect of argumentation is thus responsible for the increase of the market's accuracy with bias 0.6 recovers from 86.29% with no acquaintances to 89.21% with 4 acquaintances. Even in the extreme case with 0.8 bias, argumentation is able to increase the accuracy of individual agents. Therefore, we can conclude that an argumentation-based process of information exchange is very useful in conditions where individuals do not have a perfect (or very good) sample of data in which to base its decisions. Argumentation allows each individual agent to contrast its empirically-based judgements with those of its peers and acquire new information that, albeit partially, help it recover a better sample of data.

An interesting effect is that when the bias is extreme, the confidence the agents computes is not accurate: this is the only scenario where the market has lower accuracy than simple majority voting and where the agents obtain a negative reward. However, by means of information exchange (having acquaintances) even in this extreme scenario the quality of the individual solutions increases, and thus the market accuracy also increases, from 48.0% accuracy with no information exchange to 67.43% and 66.93% when agents have 3 and 4 acquaintances respectively.

Thus, in summary, we see that by means of information exchange using an argumentation process, agents become much more robust to biased data than standard voting mechanisms, by leveraging available information in the acquaintances' case bases. Agents with biased case bases benefit from exchanging information (by means of argumentation processes) with other agents, and the result of argumentation among agents with different biases is a less biased prediction that can largely overcome the effect of bias (except in the extreme case of 0.8 bias, where we see an improvement, but not up to the levels achieved when there is no bias). However, notice that in our experiments, the cases learnt from other agents during argumentation are not stored (for experimentation purposes), it is part of our future work to explore how retention of cases can help agents to further overcome the effect of bias (continuing the work started in previous work [12]).

5 Related Work

Research on *prediction markets* has been focused on exploiting human knowledge [17], and to our knowledge they have not been used in multiagent systems. Research in MAS is generally focused on negotiation processes and much less on social choice, in the sense of modeling and implementing processes where a group of agents achieve a joint judgment. As argued in [8], computational approaches to social choice can benefit both social choice studies and AI. Impossibility theorems proved in theoretical approaches to social choice do not prevent the design of reasonably fair and robust mechanisms [3].

Other approaches in social choice (different from prediction markets) have been applied to MAS. What we have been calling statistical means approaches (that includes voting) have been applied to MAS, from simple voting to complex schemes such as voting for combinatorial domains [10]. Deliberative approaches to group judgment have also been studied, for instance in [13] a committee of agents argue the pros and cons of a group judgment. Market mechanisms have been applied to resource allocation [15] or other types of market goods. Our focus here is rather different: developing an agent-based information or prediction market for group judgment.

Concerning on argumentation in MAS, previous work focuses on several issues like a) logics, protocols and languages that support argumentation, b) argument selection and c) argument interpretation. Approaches for logic and languages that support argumentation include defeasible logic [7] and BDI models [16]. An overview of logical models of reasoning can be found at [6]. Moreover, the most related area of research is case-based argumentation. Combining cases and generalizations for argumentation has been already used in the HYPO system [5], where an argument can contain both specific cases or generalizations. Moreover, generalization in HYPO was limited to selecting a set of predefined dimensions in the system while our framework presents a more flexible way of providing generalizations. Furthermore, HYPO was designed to provide arguments to human users, while we focus on agent to agent argumentation. Case-based argumentation has also been implemented in the CATO system[2], that models ways in which experts compare and contrast cases to generate multi-case arguments to be presented to law students.

6 Conclusions

Mechanisms for group judgment (voting, deliberation, etc) are ubiquitous in human societies. However, in addition to the formal structure of the group judgment mechanism, the informal structure play an important role [17]. We have considered here the effect of an informal structure (social networks used to exchange information mediated by argumentation) in a formal group judgment mechanism (MPM). We have shown that these social networks maybe individually useful for artificial agents, since agents may use argumentation to improve their information about the world. Therefore, artificial multiagent systems will also have to deal with the interplay of informal structures together with formal group judgment mechanism.

We have taken a typical task of prediction form a Machine Learning data set and we had goal of developing a simple market called MPM. The basic idea of MPM is that learning agents can use data concerning a *prediction task domain* to predict new unknown problems and, moreover, use the learnt data to implement a confidence estimate of their own predictions. Then, the prediction market design has to be set up to encourage the expression of the agents confidence as a "price signal". Clearly, this is a quite general approach, and different variations can be explored in future work: improving the confidence estimation functions, modifying the market reward scheme or using other machine learning techniques.

We also introduced a process of deliberation based on an argumentation protocol inside the framework of prediction markets. The reason is twofold: first, we wanted to model the idea that people often consult trusted people before making a decision (i.e. they not only learn from experience, but also from communication). Second, current state of the art in multiagent learning suggests that the individual accuracy and confidence increases after a deliberative process [13]. The experiments shown that this is the case: information exchange supported by an argumentation process increases individual accuracy and confidence. As expected, the information exchange also increases the error correlation among agents [12], decreasing the so-called "ensemble effect" that increases joint accuracy over the individual accuracy. The conclusion thus is that information exchange is beneficial to a certain extent, i.e. among a small number of individuals compared to the total number of participating individuals, in such a way that individual performance is rather increased but error correlation is not much increased. We have also shown that information exchange helps agents with biased views of the problem to overcome their bias and produce more accurate predictions.

Although we presented the results for one data set, any other classification machine learning data set could be used. Current state of the art in multiagent learning suggests that the only difference would be on the degree in which the prediction market surpasses voting [11,13].

As part of our future work, we plan to explore how our techniques will extend to fully open multi-agent systems, where there are several different problems that agents must solve, and not all agents ar competent in all of them, agents use heterogeneous learning mechanisms, and not all agents are trustable. So, agents will have to learn which agents are trustable and which ones are not, and the argumentation process has to be generalized to support heterogeneous learning methods. Our final goal is to define a framework for learning agents with problem solving, learning, collaboration and argumentation capabilities ready to be deployed and be autonomous in an open multi-agent system for real-life application.

Acknowledgements. Support for this work came from projects MID-CBR TIN2006-15140-C03-01, and Agreement Technologies CSD2007-0022.

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. Artificial Intelligence Communications 7(1), 39–59 (1994)
2. Aleven, V.: Teaching Case-Based Argumentation Through a Model and Examples. PhD thesis, University of Pittsburgh (1997)
3. List, C., Pettit, P.: Aggregating sets of judgments: Two impossibility results compared. Synthese 140, 207–235 (2004)
4. Armengol, E., Plaza, E.: Lazy induction of descriptions for relational case-based learning. In: Flach, P.A., De Raedt, L. (eds.) ECML 2001. LNCS, vol. 2167, pp. 13–24. Springer, Heidelberg (2001)
5. Ashley, K.: Reasoning with cases and hypotheticals in hypo. International Journal of Man-Machine Studies 34, 753–796 (1991)
6. Chesñevar, C.I., Mguitman, A., Loui, R.: Logical models or argument. Computing Surveys 32(4), 336–383 (2000)
7. Chesñevar, C.I., Simari, G.R.: Formalizing Defeasible Argumentation using Labelled Deductive Systems. Journal of Computer Science & Technology 1(4), 18–33 (2000)
8. Chevaleyre, Y., Endriss, U., Lang, J., Maudet, N.: A short introduction to computational social choice. In: van Leeuwen, J., Italiano, G.F., van der Hoek, W., Meinel, C., Sack, H., Plášil, F. (eds.) SOFSEM 2007. LNCS, vol. 4362, pp. 51–69. Springer, Heidelberg (2007)
9. Dietterich, T.G.: Ensemble methods in machine learning. In: Kittler, J., Roli, F. (eds.) MCS 2000. LNCS, vol. 1857, pp. 1–15. Springer, Heidelberg (2000)
10. Lang, J.: Logical preference representation and combinatorial vote. Annals of Mathematics and Artificial Intelligence 42, 37–71 (2004)
11. Ontañón, S., Plaza, E.: Justification-based multiagent learning. In: ICML 2003, pp. 576–583. Morgan Kaufmann, San Francisco (2003)
12. Ontañón, S., Plaza, E.: Case-based learning from proactive communication. In: Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI 2007), pp. 999–1004. IJCAI Press (2007)

13. Ontañón, S., Plaza, E.: Learning and joint deliberation through argumentation in multi-agent systems. In: Proc. AAMAS 2007, pp. 971–978. ACM, New York (2007)
14. Poole, D.: On the comparison of theories: Preferring the most specific explanation. In: IJCAI 1985, pp. 144–147 (1985)
15. Rodriguez-Aguilar, J.A., Sousa, P.: Issues in multiagent resource allocation. Informatica 30, 3–31 (2006)
16. Jennings, N.R., Parsons, S., Sierra, C.: Agents that reason and negotiate by arguing. Journal of Logic and Computation 8, 261–292 (1998)
17. Sunstein, C.R.: Group judgments: Deliberation, statistical means, and information markets. New York University Law Review 80, 962–1049 (2005)