

REGRET: A reputation model for gregarious societies

Jordi Sabater

IIIA - Artificial Intelligence Research Institute
CSIC - Spanish Scientific Research Council
Bellaterra, Catalonia, Spain
jsabater@iiia.csic.es

Carles Sierra

IIIA - Artificial Intelligence Research Institute
CSIC - Spanish Scientific Research Council
Bellaterra, Catalonia, Spain
sierra@iiia.csic.es

ABSTRACT

Reputation has been studied in many fields but few formal models have been defined up to now. In this paper we propose a new model for reputation that takes into account the social dimension of agents and a hierarchical ontology structure. This model is a natural extension of other models currently being used in the area of electronic commerce. We show how the model relates to these other systems and provide initial experimental results about the benefits of using a social view on the modeling of reputation.

1. INTRODUCTION

Given the great importance that has in social and commercial relations, the study and modeling of reputation has attracted the interest of scientists from different fields: sociology, economics [7, 10], psychology [5, 9]. Computer science, particularly in the areas of multi-agent systems and online communities [6, 4, 14], is not an exception.

We can define *reputation* as the “opinion or view of one about something”. This opinion is formed and updated along time through direct interactions or through information provided by other members of the society about experiences they had with that entity in the past. As a result of these interactions, individuals record *impressions* that reflect how they value the experience. These individual *impressions* are the bricks that properly combined may be used to build the reputation of other individuals.

We take the stance that reputation is compositional, the overall opinion on an entity is obtained as a result of the combination of different pieces of information. Also, we consider that the reputation on an individual is not a single and abstract concept but rather it is a multi-facet concept. For example, the reputation of being a good flying company summarizes the reputation of having good planes, the reputation of never losing luggage and the reputation of serving good food. In turn, the reputation of having good planes is a summary of the reputation of having a good maintenance

service and the reputation of frequently renewing the fleet. The different types of reputation and how they are combined to obtain new types is what we call the *ontological dimension* of reputation. Note that each individual usually has a different *ontological structure* to combine reputations and a different way to weigh the importance of reputations when they are combined.

In societies where individuals belong to groups that condition their behaviour, besides the direct interaction with an individual (what we call the *individual dimension* of reputation), the interaction with other members of his group is also very influential on the reputation associated to that individual. In human societies, previous experiences of the members of the group to which individual who is assessing a reputation belongs to are also taken into account [9]. This is what we call the *social dimension* of reputation.

Because each member in a society has a particular point of view, each member may record a different impression from the same interaction. Also, each individual has its own *ontological structure* (e.g. some travelers may not consider quality of the food as influencing the reputation of a carrier). This means that each member has a different perception of the reputation of a given entity and, therefore, that reputation is linked to subjectivity. There are some situations, however, where the points of view of all the members of a community related to some specific aspect are supposed to be unified. In these situations, the set of impressions related to that aspect can be shared without taking into account who is the information source. Then, reputation is considered as a global property of an entity and common to all the members of the society. This is what usually happens, for instance, in electronic marketplaces.

In electronic marketplaces, the reputation that a user has is the result of aggregating all the impressions of the other users that interacted with him/her in the past. Amazon Auctions [1], eBay [2] and OnSale Exchange [3], for instance, are online auction houses where users buy and sell goods. Each time a new transaction is finished, the buyer rates the seller. These ratings are used to build the reputation of a seller. Sporas [14] is an evolved version of this kind of reputation models. Sporas introduces the notion of reliability of the reputation and is more robust to changes in the behaviour of a user than reputation systems like Amazon Auctions, based on the average of all the ratings given to the user. In all these systems each user has a global reputation

shared by all the observers instead of having a reputation biased by each observer. Histos [14], also oriented to electronic commerce, is a more personalized reputation system where reputation depends on who makes the query, and how that person rated other users in the online community.

In this paper we present REGRET, a model of reputation that allows to take into account the three dimensions of reputation we have described: the *individual dimension*, the *social dimension* and the *ontological dimension*. We will describe two experiments that show that the model is flexible enough to be applied successfully in two scenarios where the structure of the society is very different.

2. REPUTATION MEASURES

2.1 Outcomes

We define the *outcome* of a dialogue between two agents as both an initial contract to take a particular course of action or to establish the terms and conditions of a transaction, and the actual result of the actions taken or the actual values of the terms of the transaction. An outcome is represented as a conjunction of equalities between variables and constants. The names of these variables and constants are part of the domain ontology. Variables are used to represent the outcome features and constants to represent the values of these features.¹ We differentiate between two types of variables: the *common* variables that reflect those aspects of the contract that are agreed by the two parts and the *expected* variables, that reflect those aspects that are implicitly supposed to happen by one of the two parts (and usually unknown for the other). The value for the *expected* variables, either in the contract part or in the result part of the outcome, are linked to the subjectivity of agents. This means that, for a given dialogue between two agents, there will be two different outcomes, one per agent. We use a suffix to note to which agent a given outcome belongs to.

For instance, the outcome of a dialogue on a commercial transaction between agents a and b from the point of view of agent b could be: $o_b = (\text{Delivery_date} =_c 10/02 \wedge \text{Prize} =_c 2000 \wedge \text{Quality} =_c A \wedge \text{Delivery_date} = 15/02 \wedge \text{Prize} = 2000 \wedge \text{Quality} = C)$, where *Delivery_date* and *Prize* are *common* variables and *Quality* is an *expected* variable. This is, the agreement between a and b says that the product should arrive on 10/02 and its value should be 2000, on the other hand, agent b also expects that the quality of the product be 'A'. The product actually arrives on 15/02, with a value of 2000 and the quality from the point of view of b is 'C'. The same dialogue, from the point of view of a , could generate this other outcome: $o_a = (\text{Delivery_date} =_c 10/02 \wedge \text{Prize} =_c 2000 \wedge \text{Delivery_date} = 15/02 \wedge \text{Prize} = 2000)$ where the quality is not considered.

We note the set of all possible outcomes as \mathbf{O} .

2.2 Impressions

From now on, we will note groups of agents using upper-case letters, $(\mathcal{A}, \mathcal{B}, \dots)$, and agents, using indexed lower-case letters, (a_2, b_3, \dots) . An agent noted b_i is assumed to belong

¹We'll note $=_c$ to represent the relation between a variable and its value on a contract c , and by $=$ the value of the variable in the actual result.

to group \mathcal{B} . In a given society, we note \mathbf{G} as the set of group identifiers and \mathbf{A} as the set of agent identifiers.

We define an *impression* as the subjective evaluation made by an agent on a certain aspect of an outcome. An *impression* ι is represented then by a tuple of the form:

$$\iota = (a, b, o, \varphi, t, W)$$

where $a, b \in \mathbf{A}$ are the agents who dialogue (being a who is judging), $o \in \mathbf{O}$ is the outcome, φ the variable of the outcome that is judged, t is the time when the impression is recorded and $W \in [-1, 1]$ is the rating associated to the specific aspect being evaluated from the agent a point of view. The intuitive meaning of the scale is -1 absolutely negative, 1 absolutely positive and 0 neutral. Using the same example, a possible impression for the outcome o_b could be: $(b, a, o_b, \text{Delivery_date}, 16 : 05, -0.5)$. Note that W represents the subjective opinion of the agent who is judging, with respect to φ . For instance, the same delay for the arrival date (in this case 5 days) could be a disaster for a given agent and, hence, the value of W would be close to -1 , while in other situations it may not be a problem and the value of W can be near 0 .

I is defined as the set of all possible impressions and an agent's *impressions database* $IDB^a \subseteq I$ as a set of impressions judged by agent $a \in \mathbf{A}$. We define $IDB_p^a \subseteq IDB^a$ as the set of impressions in IDB^a that satisfy the pattern p where the general form for a pattern is

$$(a, b, o, \varphi, t, W) : expr$$

with *expr* as a logical formula in FOL over the components of the impression. The ' $_$ ' symbol is used to represent an 'anything' value. For example, IDB_p^b with $p = (-, a, -, \text{Delivery_date}, -, -) : true$ is the set of all b impressions over a that are related with a delivery date. Another example, IDB_p^a with $p = (-, b, -, -, -, W) : W \geq 0$ is the set of all a impressions over b with a positive rating value.

2.3 Subjective reputation

We use the term *subjective reputation* to talk about the reputation calculated directly from an agent's impressions database.

A subjective reputation at time t from agent a point of view and satisfying pattern p is noted as $R^t(IDB_p^a)$. To calculate a subjective reputation we use a weighted mean of the impressions' rating factors, giving more relevance to recent impressions.²

The general formula to calculate a subjective reputation is:

$$R^t(IDB_p^a) = \sum_{\iota_i \in IDB_p^a} \rho(t, t_i) \cdot W_i$$

where $\rho(t, t_i) = \frac{f(t_i, t)}{\sum_{\iota_j \in IDB_p^a} f(t_j, t)}$ and $f(t_i, t)$ is a time dependent function that gives higher values to values closer to t . A simple example of this type of function is $f(t_i, t) = \frac{t_i}{t}$. Finally, given that $W_i \in [-1, 1]$ and that $\rho(t, t_i)$ is a normalized value, it's easy to see that $R^t(IDB_p^a) \in [-1, 1]$.

²There are many psychological studies that support recency as a determinant factor [9]

From now on, we will use the notation $R_{a \rightarrow b}(\text{subject})$ to represent $R^t(IDB_p^a)$ where $p = (a, b, -, \text{subject}, -, -) : \text{true}$ and t is supposed to be the actual time.

Besides the reputation value, it is important to know how reliable is that value. Although there are a lot of elements that can be taken into account to calculate how reliable a subjective reputation is, we will focus on two of them: the number of impressions used to calculate the reputation value and the variability of its rating values (the impressions' rating deviation). This approach is similar to that used in the Sporas system [14].

The intuition behind the number of impressions factor is that in a real society, an isolated experience (or a few of them) is not enough to make a correct judgement of somebody. You need certain amount of experiences before you can say how is that person. As the number of impressions grows, the reliability degree increases until it reaches the maximum value at what we call the *intimate* level of interactions (*itm* from now on). From a social point of view, this stage is what we know as a close relation. More experiences will not increase the reliability of our opinion from then on.

Next function is an example of a simple function that can be used to model this:

$$Ni(IDB_p^a) = \begin{cases} \sin\left(\frac{\pi}{2 \cdot itm} |IDB_p^a|\right) & |IDB_p^a| \in [0, itm] \\ 1 & \text{otherwise} \end{cases}$$

Where $|IDB_p^a|$ is the cardinality of IDB_p^a , this is, the number of impressions used to calculate the reputation.

Note that the value of the *itm* parameter is application dependent. It depends on the interaction frequency of the individuals in that society as well as the quality of the impressions.

The subjective reputation deviation is the other factor that our model takes into account to determine the reliability of a subjective reputation. The greater the variability in the rating values the more volatile will be the other agent in the fulfillment of its agreements.

We calculate the subjective reputation deviation as

$$Dt(IDB_p^a) = 1 - \sum_{t_i \in IDB_p^a} \rho(t, t_i) \cdot |W_i - R^t(IDB_p^a)|$$

This value goes from 0 to 1. A deviation value near 0 indicates a high variability in the rating values (this is, a low credibility of the reputation value from the subjective reputation deviation point of view) while a value close to 1 indicates a low variability (this is, a high credibility of the reputation value).

Finally, we define the reliability of a subjective reputation as a convex combination of the function Ni and the impres-

sions' rating deviation Dt .

$$RL(IDB_p^a) = (1 - \mu) \cdot Ni(IDB_p^a) + \mu \cdot Dt(IDB_p^a)$$

As before, from now on we will use the notation $RL_{a \rightarrow b}(\text{subject})$ to represent $RL^t(IDB_p^a)$ where $p = (a, b, -, \text{subject}, -, -) : \text{true}$ and t is the actual time.

3. THE REGRET SYSTEM

3.1 Individual dimension

The *individual dimension* models the direct interaction between two agents.

We define the reputation measure that takes into account the *individual dimension* as:

$$R_{a \rightarrow b}(\text{subject})$$

3.2 Social dimension

In the *individual dimension* it is only considered the direct interaction between the agent who is evaluating and the agent being evaluated. With the social dimension, we add the possibility to reflect in our model a characteristic of complex societies: the group relation.

In many societies, an individual inherits by default the reputation of the group it belongs to. When direct information due to personal interactions with the entity itself is lacking, this group reputation gives initial expectations about the behaviour of an agent because belonging to a certain group implies, a priori, that its members share a common way of thinking. In a symmetric way, an individual will use the experiences of the other members of his own group to complement his own experiences. That is, the experiences that the members of a group had with a given entity and also with the group that entity belongs to, influences and completes the point of view of each one of its members [9].

Hence, when we take into account the group relation, we are adding three new sources of information to calculate a reputation value. Besides the direct interaction with the agent itself (the *individual dimension*), now we have to consider the interaction with the members of its group, the information that our group has related to that agent and, finally, the information that our group has related to its group.

3.2.1 The personal experience

As the title suggest, this part contributes to the *social dimension* with all the information related with the experience accumulated by the agent who is calculating the reputation. We use a to note the agent who calculates the reputation and b to note the agent being evaluated.

On the one hand we have the direct interaction between a and b , this is, the *individual dimension*:

$$R_{a \rightarrow b}(\text{subject})$$

On the other, the interaction with the other members of the group to which agent b belongs to, represented by:

$$R_{a \rightarrow \mathcal{B}}(\text{subject}) = \sum_{b_i \in \mathcal{B}} \omega^{ab_i} \cdot R_{a \rightarrow b_i}(\text{subject})$$

where $\sum_{b_i \in \mathcal{B}} \omega^{ab_i} = 1$. As in the *subjective reputation* case we need a mean to express how reliable this reputation is:

$$RL_{a \rightarrow \mathcal{B}}(subject) = \sum_{b_i \in \mathcal{B}} \omega^{ab_i} \cdot RL_{a \rightarrow b_i}(subject)$$

Note that if agent a knows how close (from the point of view of the behaviour of the agent) is the agent b to the other members of its group, a can adjust ω^{ab_i} giving more weight to those agents closer to agent b . When uncertainty is maximal $\omega^{ab_i} = 1/|\mathcal{B}|$.

3.2.2 The group experience

Once we have taken into account the personal experience, it is time to consider what the other members of the group think about the agent being evaluated and his group.

To represent what the members of the group think about the agent being evaluated we use the formulae:

$$R_{\mathcal{A} \rightarrow b}(subject) = \sum_{a_i \in \mathcal{A}} \omega^{a_i b} \cdot R_{a_i \rightarrow b}(subject)$$

$$RL_{\mathcal{A} \rightarrow b}(subject) = \sum_{a_i \in \mathcal{A}} \omega^{a_i b} \cdot RL_{a_i \rightarrow b}(subject)$$

where $\sum_{a_i \in \mathcal{A}} \omega^{a_i b} = 1$.

To represent what the members of the group think about the other group we use the formulae:

$$R_{\mathcal{A} \rightarrow \mathcal{B}}(subject) = \sum_{a_i \in \mathcal{A}} \omega^{a_i \mathcal{B}} \cdot R_{a_i \rightarrow \mathcal{B}}(subject)$$

$$RL_{\mathcal{A} \rightarrow \mathcal{B}}(subject) = \sum_{a_i \in \mathcal{A}} \omega^{a_i \mathcal{B}} \cdot RL_{a_i \rightarrow \mathcal{B}}(subject)$$

again, where $\sum_{a_i \in \mathcal{A}} \omega^{a_i \mathcal{B}} = 1$.

The $\omega^{a_i b}$ and $\omega^{a_i \mathcal{B}}$ can reflect the credibility or hierarchical relations inside the group, giving more importance to the information coming from those agents with more credibility or with a higher status in the hierarchy.

3.2.3 Putting all together

We define the reputation measure that takes into account the *social dimension* as:

$$\begin{aligned} SR_{a \rightarrow b}(subject) &= \xi_{ab} \cdot R_{a \rightarrow b}(subject) + \\ &\xi_{a\mathcal{B}} \cdot R_{a \rightarrow \mathcal{B}}(subject) + \\ &\xi_{\mathcal{A}b} \cdot R_{\mathcal{A} \rightarrow b}(subject) + \\ &\xi_{\mathcal{A}\mathcal{B}} \cdot R_{\mathcal{A} \rightarrow \mathcal{B}}(subject) \end{aligned}$$

$$\begin{aligned} SRL_{a \rightarrow b}(subject) &= \xi_{ab} \cdot RL_{a \rightarrow b}(subject) + \\ &\xi_{a\mathcal{B}} \cdot RL_{a \rightarrow \mathcal{B}}(subject) + \\ &\xi_{\mathcal{A}b} \cdot RL_{\mathcal{A} \rightarrow b}(subject) + \\ &\xi_{\mathcal{A}\mathcal{B}} \cdot RL_{\mathcal{A} \rightarrow \mathcal{B}}(subject) \end{aligned}$$

where $\xi_{ab} + \xi_{a\mathcal{B}} + \xi_{\mathcal{A}b} + \xi_{\mathcal{A}\mathcal{B}} = 1$.

Adjusting these parameters, the agent can give more or less importance to each source of opinion. These values are not necessarily static and can change along time according to the state of the agent.

3.3 Ontological dimension

Along the individual and social dimensions, reputation is always linked to a single aspect. As we argued in the introduction, with the ontological dimension we add the possibility to combine reputations on different aspects to calculate a more complex one. To represent the *ontological* dimension we use graph structures like the one showed in Figure 1. In this example, and from the point of view of a particular agent, the reputation of being a good seller is related with the reputation of providing good quality products, offering good prizes and delivering the products quickly.

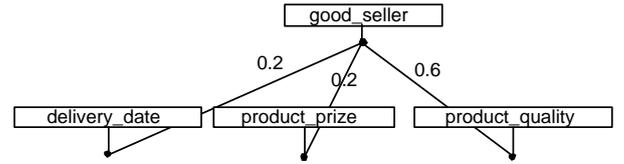


Figure 1: An ontological structure.

Hence, to calculate a given reputation taking into account the *ontological dimension*, an agent has to calculate the reputation of each one of the related aspects that, in turn, can be the node of another subgraph with other aspects associated. The reputation of those nodes that are related with an atomic aspect of the behaviour, are calculated using the individual and social dimensions. The reputation of node i in an ontological graph is computed as follows:

$$OR_{a \rightarrow b}(i) = \sum_{j \in \text{children}(i)} w_{ij} \cdot OR_{a \rightarrow b}(j)$$

$$ORL_{a \rightarrow b}(i) = \sum_{j \in \text{children}(i)} w_{ij} \cdot ORL_{a \rightarrow b}(j)$$

Where $OR_{a \rightarrow b}(j) = SR_{a \rightarrow b}(j)$ when j is an atomic aspect.

For instance, using the ontological structure in Figure 1 we can calculate the reputation of b as a good seller from a 's perspective using the formula:

$$\begin{aligned}
OR_{a \rightarrow b}(\text{good_seller}) &= 0.2 \cdot SR_{a \rightarrow b}(\text{delivery_date}) + \\
&0.2 \cdot SR_{a \rightarrow b}(\text{product_prize}) + \\
&0.6 \cdot SR_{a \rightarrow b}(\text{product_quality})
\end{aligned}$$

and similarly for the reliability.

Note, again, that the importance of each aspect is not necessarily static and could change over time, according to the necessities of the agent.

4. EXPERIMENTS

4.1 Online marketplaces

Actual online marketplaces present one of the most simple kinds of community from the point of view of reputation. Agents are not grouped, so it is not possible to take into account the affiliation of an agent. Only the direct interaction with the agent can be used to evaluate its reputation. Moreover, there is only one type of reputation and this reputation is not personalized, this is, the reputation of an agent is common and known to all the community. Each time there is an interaction between two agents, each agent rates the other and their respective reputations are publicly updated according to that rating.

Given that only a single aspect of the behaviour of an agent is considered (being a reliable trader), the *ontological structure* consist of just a single node. From the *social dimension* point of view, all agents are considered to belong to the same group. The only information that we can take into account is the information coming from this group, this is, we compute $SR_{a \rightarrow b}(\text{subject})$ with $\xi_{ab}, \xi_{aB}, \xi_{AB} = 0$ and $\xi_{Ab} = 1$. In these kind of systems it is assumed that all agents have the same relevance. This is equivalent, in our model, to adjust the weights $w^{a_i b}$ for each subjective reputation in $R_{A \rightarrow b}(\text{subject})$ as $1/|\mathcal{A}|$, where $|\mathcal{A}|$ is the cardinality of group \mathcal{A} .

To see how it works, we have compared our model with the Sporas ³ system [14] and the method used in Amazon Auctions [1].

The experiment is a replica of the experiment described in [14] where a user who joins a marketplace behaves reliably until s/he reaches a high reputation value and then starts committing fraud. In order to be able to compare our results with the results of Sporas and Amazon we have adapted the range of the reputation value from $[-1,1]$ to $[0,3000]$.

In this example, during the first 66 interactions, the user reaches a reputation level of 2400 but during the next 134 interactions, the user starts behaving as having a reputation level of 720. The user receives ratings from 100 users with uniformly distributed reputations. These ratings are normally distributed around his/her actual performance with a standard deviation of 300.

In Figure 2 we plot in the same graph the reputation of the user from the point of view of each model.

³The parameters used in the Sporas systems were $\sigma = 200$, $\Theta = 8$ and $\lambda = 0.8$

As in the Sporas system, we use a forget factor that allows REGRET to consider only the most recent *impressions*. This is why the behaviour of REGRET is similar to Sporas when the user starts to commit fraud in the example (compare with the Amazon's behaviour that doesn't have this capability). In this experiment, time of an impression is made equal to the arrival order of that impression to the agents data base. Then we can define the part of the impressions database used to calculate a reputation as IDB_p^a with $p = (a, -, -, -, t, -) : t \geq (\text{actual_time} - \theta)$, where θ is the window of *impressions* considered to calculate the reputation. For this example we have used $\theta = 20$.

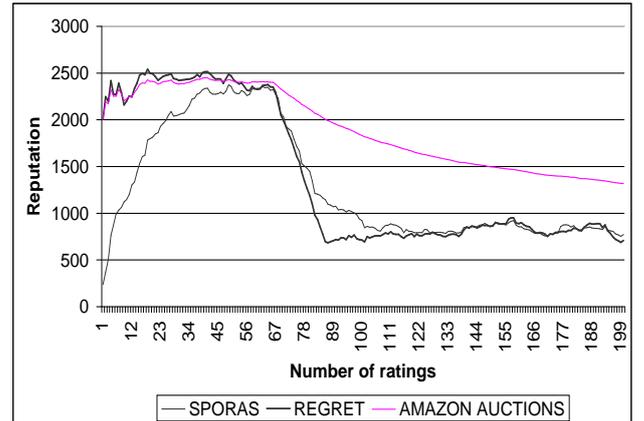


Figure 2: Abuse of prior performance

As we can see from the graph, our model is able to reach equilibrium and react to changes in the behaviour of a user very quickly.

4.2 The tourism scenario

The interaction between travelers and travel agencies and the relation between travel agencies and tour operators define the scenario for this experiment. This scenario has been designed to demonstrate two things: how the model works in a complex environment and the influence of the *social dimension*.

With this experiment we don't intend to be exhaustive. As you will notice, we fix a lot of parameters that influence the behaviour of the system. More experiments are needed in order to prove the generality of our conclusions. Our main aim is only to show how to use REGRET in a complex environment.

4.2.1 Describing the experiment

The society for this experiment is composed by travelers, travel agencies and tour operators. The relation that exists between them is as follows:

- Travelers: They have to travel regularly either for business or for holiday and have to buy these trips to travel

agencies. The kind of services they need are different if it is a business travel or a holiday travel.

- Travel agencies: Travel agencies provide travelers with packages that include all the facilities needed for a given type of travel. Each travel agency is affiliated with a tour operator that is who provides, in the end, the facilities. Hence, agencies affiliated to the same tour operator will offer similar products.
- Tour operators: They provide to their affiliated agencies, hotels, tours, leisure activities and guides. Each tour operator has its own standards of quality that are not necessarily compliant with the standards of quality that the travelers have.

A travel is typified by a set of variables related with different aspects of that travel. These variables are:

- Hotel location: Its value is a natural number that represents the distance from the hotel to the target place. The utility function used by all the travelers in this experiment is showed in Figure 3. It is a *common* variable.
- Hotel quality: The category of the hotel (in number of stars). Figure 3 shows its utility function. It is a *common* variable.
- Leisure Activities: In business travels, the travel agency organizes extra activities for the traveler to relax after work. This variable reflects the quality of these extra activities. It is an *expected* variable.
- Tours itineraries: This variable measures the quality of the itineraries that the agency has organized to visit places during a holiday travel. It is an *expected* variable.
- Tours organization: How the tours are organized in general. To give value to this variable, aspects like the time table, the food or the bus used to do the tour could be considered. It is an *expected* variable.
- Guides: The quality of the guides in a holiday travel. Is is an *expected* variable.

The utility function for all the *expected* variables used in this experiment is showed in Figure 4.

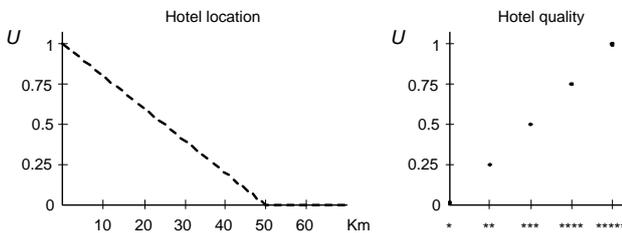


Figure 3: Utility functions: hotel location and hotel quality variables.

Given that, in figure 5 we can see how these variables relate in the *ontological structure* considered by travelers.

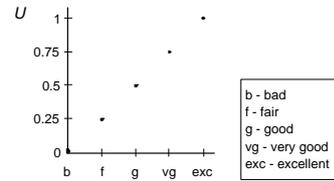


Figure 4: Utility function: *expected* variables

Because in this experiment travelers don't interact between them, in the *social dimension* we only take into account the direct interaction with the travel agency and with the travel agencies that are affiliated to the same tour operator. This is, we consider $SR_{a \rightarrow b}(subject)$ where $\xi_{ab} = 0.5$, $\xi_{aB} = 0.5$, $\xi_{Ab} = 0$ and $\xi_{AB} = 0$.

Figure 6 shows the exact composition of the society in this experiment. There are 4 tour operators, each one with different quality standards that are inherited by their affiliated travel agencies, 36 travel agencies and 4 travelers.

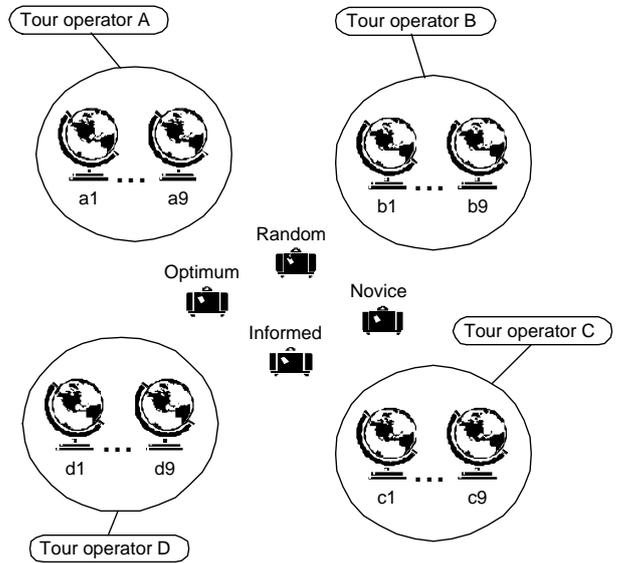


Figure 6: The tourism society

The experiment consists on 45 interactions between travelers and travel agencies. Each interaction follows the next pattern:

- The system chooses a kind of travel and three travel agencies, from the set of all travel agencies available, randomly.
- These three agencies are supposed to be the unique agencies that can provide the kind of travel the travelers need at that moment. Each traveler has to chose one of them trying to chose the travel agency that s/he thinks will provide the best travel.
- The system “simulates” the travel and generates an outcome for each traveler. The travelers use that outcome to generate impressions and feed their reputation model.

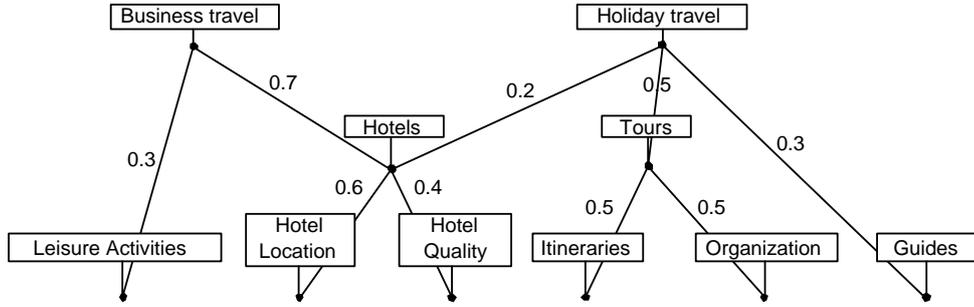


Figure 5: Ontological dimension for the tourism scenario

Next we describe each one of the 4 travelers. Two of these travelers are there for control reasons, one is the *optimum* traveler and the other is the *random* traveler. The other two use the reputation model presented in this paper.

- Optimum traveler: The particularity of this traveler is that he can choose the travel agency after the actual results of each travel are known. Therefore, he always chooses the best agency. We use this traveler to know, in a given run, the best performance that can be achieved.
- Random traveler: This traveler chooses the travel agency randomly.
- Informed traveler: Uses the reputation model presented in this paper to evaluate the reputation of travel agencies. These reputations are used to decide which is the best travel agency. In this case, to calculate the reputation of a given travel agency takes into account all the aspects of the *social dimension* commented before, this is, the direct interaction with the travel agency and the interaction with other travel agencies affiliated with the same tour operator.
- Novice traveler: Like the *informed* traveler but doesn't consider the *social dimension*, just the *individual dimension*.

An outcome⁴ in the tourism scenario is represented by the variables that typify a travel. In a business travel, these variables are the hotel location, the quality of the hotel and the leisure activities. In a holiday travel the variables are the hotel location, the quality of the hotel, the tours itineraries, the tours organization and the quality of the guides. (see again Figure 5)

In our experiment, a traveler hopes always a value of 'good' in the *expected* variables. On the other hand, the agreement with the travel agency related with the *common* variables is always: $\text{Hotel_Quality} =_c \text{'***'} \wedge \text{Hotel_Location} =_c 25$. This values are invariant for all interactions in the experiment. This means that the contract part of an outcome for a business travel will always be $(\text{Hotel_Quality} =_c \text{'***'} \wedge \text{Hotel_Location} =_c 25 \wedge \text{Leisure_Activities} =_c \text{'good'})$ and the

⁴Because we are only interested in travelers, from now on, when we talk about the outcome of a given interaction we will be referring to the outcome of the traveler.

contract part for a holiday travel will be $(\text{Hotel_Quality} =_c \text{'***'} \wedge \text{Hotel_Location} =_c 25 \wedge \text{Tours_Itineraries} =_c \text{'good'} \wedge \text{Tours_Organization} =_c \text{'good'} \wedge \text{Guides} =_c \text{'good'})$.

The result part is calculated using a different function depending on which agency organizes the travel. We define two types of travel agencies:

- Good agencies: The values of the *expected* variables of the result part of an outcome are always 'very_good' or 'excellent', the Hotel_Quality '****' or '*****' and the Hotel_location 10 or less.
- Bad agencies: The values of the *expected* variables of the result part of an outcome are always 'fair' or 'bad', the Hotel_Quality '**' or less and the Hotel_location more than 40.

Tour operators \mathcal{A} and \mathcal{C} have Good agencies affiliated while tour operators \mathcal{B} and \mathcal{D} have Bad agencies.

Finally, we calculate the rating for an impression $\iota = (a, b, o, \text{subject}, t, W)$ where $o = (\dots \text{subject} =_c v_c \wedge \text{subject} = v \dots)$ using the formula $W = U(v) - U(v_c)$ where U is the utility function associated to the variable 'subject'.

4.2.2 The results

In order to compare the results properly, we need a measure that summarizes, in a single value, the degree of satisfaction of a traveler after each interaction. To calculate this value we use the rating of the impressions resulting from that interaction and the ontological structure. Instead of using the individual and social dimension in the leaves of the ontological structure, we use the ratings of the impressions and follow up the tree combining these values to obtain a single value related to the kind of travel. We call this measure the *degree of satisfaction* of a traveler for a given interaction.

Figure 7 shows the results for this experiment where the 'y' axis represents the average of the last 10 *degrees of satisfaction* for a traveler.

Note that the performance of the *informed* traveler increases more quickly and is closer to the optimum than the performance of the *novice* traveler. This is due to the extra information that suppose taking into account the affiliation of the travel agency, in other words, the *social dimension*.

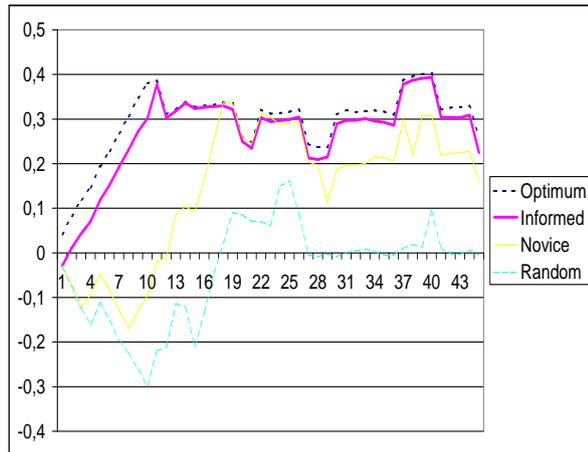


Figure 7: The tourism scenario results

5. RELATED WORK

The idea of using the opinion of other agents to build a reputation is not new. The work of Michael Schillo, Petra Funk and Michael Rovatsos [12] and the work of Bin Yu and Munindar P. Singh [13] are good examples of this. In both cases they use a trust-net for weighting the other agents' opinions. We take into account our group "trust-net" to determine how the information coming from other agents is propagated in order to be merged with the personal experiences to calculate a final reputation measure.

The model described in [13] merges information that comes from agents that have good reputation. In our case, the criteria used to decide whether a piece of information is considered is whether it comes from an agent belonging to our group or not. We both implicitly assume that the information merged comes from trustful agents that do not deliberately manipulate the information (in one case because they have good reputation and in our case because they belongs to our group). On the contrary, in [12] the same agents that can provide you with information are also competing with you. Although agents are assumed to never lie, they can hide information or bias it to favour their goals. Our approach radically differs from these in the sense that we take the stance that the agents' social structure must be an essential factor in weighing the other agents' opinions. We take into account our own group members' opinion about our opponent as well as our group members' opinion about our opponent's group.

Unlike REGRET, these models [13, 12] consider reputation as a single concept instead of a multi-facet concept. We claim that reputation has different dimensions. Distinguishing these dimensions can improve the behaviour of a negotiating agent in those scenarios where the contracts are described as multi-dimensional objects and specially when there are intersections between the issues of these contracts (see the tourism scenario example).

6. CONCLUSIONS AND FUTURE WORK

We have presented a reputation model that takes into account the social dimension of agents and a hierarchical ontology structure that allows to consider several types of reputation at the same time. This model is flexible enough to be successfully used in societies with a different structure as has been shown in the experiments. These experiments have to be extended to validate the model in different situations specially where individuals use the experiences of the other members of the group to complement its own experiences (the group experience). Moreover, the model has to be extended to allow agents to belong to more than one group at a time.

Finally, although these kind of experiments are useful to test the properties of the model, a reputation model only makes sense as part of a negotiation mechanism. The use of our model in negotiation processes as those described in [11, 8] is one of our main goals for the next future.

7. ACKNOWLEDGMENTS

This work has been supported by the European project SLIE, IST-1999-10948.

8. REFERENCES

- [1] *Amazon Auctions*. <http://auctions.amazon.com>.
- [2] *eBay*. <http://www.eBay.com>.
- [3] *OnSale*. <http://www.onsale.com>.
- [4] A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Proceedings of the Hawaii Int. Conference on System Sciences, Maui, Hawaii*, 2000.
- [5] D. B. Bromley. *Reputation, Image and Impression Management*. John Wiley & Sons, 1993.
- [6] C. Castelfranchi and R. Falcone. Principles of trust for mas: Cognitive anatomy, social importance, and quantification. In *Proceedings of the 3th International Conference on Multi-Agent Systems*, pages 72–79, 1998.
- [7] M. Celentani, D. Fudenberg, D.K. Levine, and W. Psendorfer. Maintaining a reputation against a long-lived opponent. *Econometrica*, 64(3):691–704, 1966.
- [8] P. Faratin, C. Sierra, and N.R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, (24):159–182, 1997.
- [9] Marvin Karlins and Herber I. Abelson. *Persuasion, how opinion and attitudes are changed*. Crosby Lockwood & Son, 1970.
- [10] R. Marimon, J.P. Nicolini, and P. Teles. Competition and reputation. In *Proceedings of the World Conference Econometric Society, Seattle*, 2000.
- [11] N. Matos and C. Sierra. Evolutionary computing and negotiating agents. In P. Noriega and C. Sierra, editors, *Agent Mediated Electronic Commerce*, number 1571 in Lecture Notes in Artificial Intelligence, pages 126–150. Springer-Verlag, 1998.

- [12] Michael Schillo, Petra Funk, and Michael Rovatsos. Using trust for detecting deceitful agents in artificial societies. In *In Applied Artificial Intelligence, Special Issue on Trust, Deception and Fraud in Agent Societies*, 2000.
- [13] Bin Yu and Munindar P. Singh. A social mechanism of reputation management in electronic communities. In *Cooperative Information Agents, CIA-2000, Boston, MA, USA*, pages 154—165, 2000.
- [14] Giorgios Zacharia. Collaborative reputation mechanisms for online communities. Master's thesis, Massachusetts Institute of Technology, September 1999.