Learning Agents' Reliability Through Bayesian Conditioning: a Simulation Experiment

Aldo Franco Dragoni, Paolo Giorgini

Istituto di Informatica, Università di Ancona,
via Brecce Bianche, 60131, Ancona (Italy)
{dragon,giorgini}@inform.unian.it

Abstract. This paper reports the first results of a simulation experiment. There are two databases, one containing true propositions and the other containing their respective negations. Five agents in turn access one of them. Each agent has a "capacity" that will be used as the frequency with which the agent accesses (unconsciously) the database with the correct knowledge. Agents randomly exchange information with the others. Since they have limited degrees of capacity, their "cognitive state" quickly becomes inconsistent. Each agent is equipped with the same belief revision mechanism to detect and solve these contradictions. This adopts the Dempster's Rule of Combination to evaluate the credibility of the various pieces of information and Bayesian Conditioning to estimate the relative degrees of reliability of the agents (itself included). The purpose of the experiments was that of evaluating on a statistical basis, the emergent cognitive behavior of the group.

1 Introduction

A biologist, asked whether there will ever be on earth organisms more complex than human brain, answered that such an astonishing super-brain already exists and it is the community of the minds interacting over the surface of the globe. Extending this idea, one might argue that these interacting minds are not necessarily embodied in currently living brains, since amanuenses and printing made it possible the interaction (monodirectional) among intelligence incarnated in brains distant many generations. Present technology greatly improved our ability to interact over time and space; by multiplying and enforcing the links among individuals, it contributed to the development of the global mind mentioned by the biologist. In our days, technology overwhelmed itself by creating artificial (non-biological) bodies in which it is possible to materialize intelligence (intended as the ability to conceive and solve a problem) and communicative abilities. So, the global mind has to be conceptually extended to comprise also such artificial brains.

We concentrate on a very specific problem of collective intelligence: the *distributed elicitation* of the knowledge, i.e., how is it possible that from an enormous variety of inferential schemes and judging capabilities, from different opinions and dogmas, from distinct perspectives and opposite point of views, after a continual interaction, a more uniform (if not unique) vision of the world emerges? Under which singular circumstances the emergent representation of the world results a correct one? The idea is to capture some successful mechanisms in order to replicate them in a world of intelligent highly-engineered (programmed or trained) interacting agents.

In this regard, a central question is how should each agent ascribe a relative degree of reliability to any agent (itself included). Indeed, one of the main differences between classical and multi-agent learning is that, while it makes no sense saying that the external world one should learn about is *incorrect*, the agents from which one get information can do lye, deliberately or not. Even the sensing equipment from which one experiences the external world may be faulty, thus, even for a collaborative group, it is important to define methods to assess each member's relative degree of reliability. According to us, these ascription should be performed under "liberal policies", i.e., each one should be permitted to stand on its own opinion regarding him/herself and the others. Furthermore, these ascription should be performed on the basis of the reciprocal experience and acquaintance, thus we could talk of a form of learning *about* the others in addition to learning *from* the others.

With the distributed elicitation of the knowledge we distinguish two desiderata:

- *convergence*: which local decision strategies favor the convergence of the opinions rather than the chaotic divergence or the static indifference?
- *correctness*: among the local decision strategies of the previous point, which favor the convergence to the most correct pieces of knowledge globally available to the overall agency?

By "local strategies" we mean:

- policies of communication
- methods to form a personal opinion

Each agent's policy of communication should define:

C1. *when* communicate

C2. *what* communicate

C3. to *whom* communicate

These criteria are almost conscious to humans. On the contrary, the ways we form our opinion from:

- information directly perceived from our experience of the world
- information received from the others

seem quite unconscious. Perhaps, one makes at least three kinds of check after the incoming of a previously unbelieved information:

A1. "although I was not aware of it, is this information in accordance with my direct experience of the world?"

A2. "how many people I know believe in it?"

A3. "how much reliable are those people?"

People who prefer the check A1 are confident in themselves; those who prefer A2 are rather conformist, while check A3 is preferred by the suspicious people.

A question studied in this paper is: if all the individuals adopt the same local criteria to evaluate an incoming information, how the global process of knowledge elicitation is affected? For instance, if all the individuals adopt only the control A1, then we would expect that no gain results from the interaction and each remains in his personal degree of correctness. On the other hand, if all the individuals adopt exclusively A2, with no regard toward his own and the others' reliability, then we would expect a global flattening to the medium degree of correctness of the agency. Perhaps we should adopt a reasonable mix of these criteria.

Unfortunately, these criteria are very rough to be studied on a statistical simulation basis. We need more precise techniques. Vaguely, opinion formation deals with:

O1. the way we integrate evidences coming from different sources (ourselves included)

O2. the way we estimate the reliability of the various sources (ourselves included)

O3. the way we revise our opinions after the incoming of a new information that contradicts some of our previous convictions

In section 2 we introduce the techniques our agents adopt to perform O1÷O3. Section 3 introduces the simulation experiment that we are currently carrying on, and reports some of the first results. We will conclude presenting our future work in section 4.

2 Revising/Integrating Knowledge in a Multi-Source Environment

As computer scientists, we are more interested in computer networks than in philosophy, sociology or history of science, so, from now on, we leave the generic distributed scenario for that of a net of artificial (programmed), intelligent and interacting agents.

As the importance of computer networking increases, some tasks traditionally performed by a single workstation are being conceived as performable in a distributed manner by groups of interacting knowledge-based systems. Sometimes, the contribution coming from the other systems could be harmful rather than helpful, for instance because:

B1. local knowledge available at the other systems may not be updated or truthful

B2. there can be some hardware faults at the other workstations

B3. there can be some mistakes in the architecture or software implementation of the other systems

B4. some nodes could join the network with non-cooperative intentions and could supply incorrect information in order to induce wrong decisions and inferences

B5. even if the systems adopt standard protocol and syntax to exchange information, there can still be some semantic differences or ambiguities

B6. there can be noises on the communication channels.

We define as *incompetent* a system that is unreliable because of B1, B2 and B3, and we define as *insincere* a system that is unreliable because of B4. Items B5 and B6 regard communication. We expect that standardization and technology will reduce the effects of B5 and B6; they yield incompleteness rather than incorrectness, i.e., they simply make it impossible to transfer some pieces of information but very seldom cause the reception of a valid piece of information different from the one sent. We concentrate on the kinds of mistakes B1÷B4, which can be ascribed to the responsibility of the sending software/hardware agent.

Many pieces of incorrect knowledge run silently through the network, spreading from node to node. However, occasionally, a symptom that some degree of incompetence or insincerity is insinuated into the network is the appearance of contradictions in the knowledge base of some nodes. Thus, we need a method to integrate information coming from various sources with limited degrees of correctness. From our experience of multi-agent [1] and investigative domains [2], we choose the Dempster's Rule of Combination to perform O1 (previous section), the Bayesian Conditioning to perform O2 and the "Assumption-Based approach" to perform O3. Starting from the beginning, let us recapitulate here our ideas about how to revise/integrate beliefs in a multi-agent environment.

Initially defined as a symbolic model-theoretical problem [3-6], *belief revision* has also been approached both as a qualitative syntactic process [7,8] and as a numerical mathematical issue [9]. Beliefs can be represented either as weighted sentences of a decidable language *L* or as sets of weighted possible worlds (the models of the sentences). Weights can be either *reals* (normally between 0 and 1), representing explicitly the credibility of the sentences/models, or *ordinals*, representing implicitly the believability of the sentences/models w.r.t. the other ones. Essentially, belief revision consists in the redefinition of these weights in the light of the incoming information (see [16] for an overview). Most of the symbolic and numerical models for belief revision developed so far obey the following three rationality principles:

P1. *Consistency*: revision must yield a consistent cognitive state

P2. *Minimal Change*: revision should alter as less as possible the previous opinion

P3. *Priority of Incoming Information*: the revised cognitive state must embody the information that caused revision.

While the last principle is acceptable when updating the representation of an *evolving* world, it is not generally justified when revising the representation of a *static* situation. In this case, the chronological sequence of the informative acts has nothing to do with their credibility or importance. Another point is that changes should not be irrevocable. To make practical and useful belief revision in a multi-agent environment, we substitute the priority of the incoming information with the following principle [10].

P3. *Recoverability*: any previously believed information item must belong to the current cognitive state if it is consistent with it.

We will achieve recoverability by imposing the *maximal consistency* of the revised cognitive state.

Beliefs introduced in the agent's cognitive state directly by the various sources (itself included) are called *assumptions*. Those deductively derived from the assumptions (internally to the agent's cognitive state) are called *consequences.* We call *Knowledge Base* (*KB*) the set of the assumptions, and *Knowledge Space* (*KS*) the set of all the beliefs (assumptions + consequences). *KB* and *KS* grow monotonically since none of their element is ever erased from memory. Normally both contain *contradictions*. Essentially, our method of belief revision consists of five steps:

S1. Detection of all the minimally inconsistent subsets of KB (*nogoods*)

S2. Detection of all the maximally consistent subset of KB (*goods*)

S3. Defining a credibility ordering $\leq_{KB}$ over the assumption in KB

S4. Extending $\leq_{KB}$ into a credibility ordering $\leq_G$ over the goods detected in S2

S5. Selecting the preferred good with its corresponding set of consequences

2.1 S1 and S2

S1 and S2 deal with consistency and work with the symbolic part of the beliefs. Their task is that of providing the maximal consistency of the revised cognitive state; they adopt an ATMS [11] as basic mechanism. Goods and nogoods are detected through the Reiter's set-covering algorithm for model-based diagnoses [18].

2.2 S3

S3 deals with uncertainty and works with the numerical part of the beliefs. Its task is that of integrating evidences coming from different sources; the basic mechanism is the Dempster's Rule of Combination. We refer to [12-14] for a complete presentation of the Dempster-Shafer Theory of Evidence. Here we recapitulate the main concepts and definition as they have been exploited in our agents.

To begin with, we introduce two data structures: the *reliability set* and the *information set*. Let $S=\{S_1,\ldots,S_n\}$ be the set of the sources and $I=\{I_1,\ldots,I_m\}$ be the set of the information items given by these sources. Then:

- *reliability set* = $\{<S_1, R_1>,\ldots,<S_n, R_n>\}$, where R_i (a real in [0,1]) is the reliability of S_i, interpreted as the "a priori" probability that S_i is reliable.
- *information set* = $\{<I_1, Bel_1>,\ldots,<I_m, Bel_m>\}$, where Bel_i (a real in [0,1]) is the credibility of I_i.

The reliability set is one of the two inputs of the belief-function formalism (see figure 1). The other one is the set $\{<S_1, s_1>,\ldots,<S_n, s_n>\}$, where s_i is the subset of I made of all the information items given by S_i. The information set is the main output of the belief-function formalism. Let us see now how the mechanism works. Let Ξ denote the set of the atomic propositions of L. The power set of Ξ, $\Omega=2^{\Xi}$, is called *frame of discernment*. Each element ω of Ω is a "possible world" or an "interpretation" for L (the one in which all the propositional letters in ω are true and the others are false). Given a set of sentences $s\subseteq I$ (i.e., a conjunction of sentences), $[s]$ denotes the interpretations which are a model for all the sentences in s.

The key assumption with this multi-source version of the belief function framework is that *a reliable source cannot give false information, while an unreliable source can give correct information*; the hypothesis that "S_i is reliable" is compatible only with $[s_i]$, while the hypothesis that "S_i is *un*reliable" is compatible with the entire set Ω. Each S_i gives an *evidence* for Ω and generates the following *basic probability assignment* (*bpa*) m_i over the elements X of 2^{Ω}:

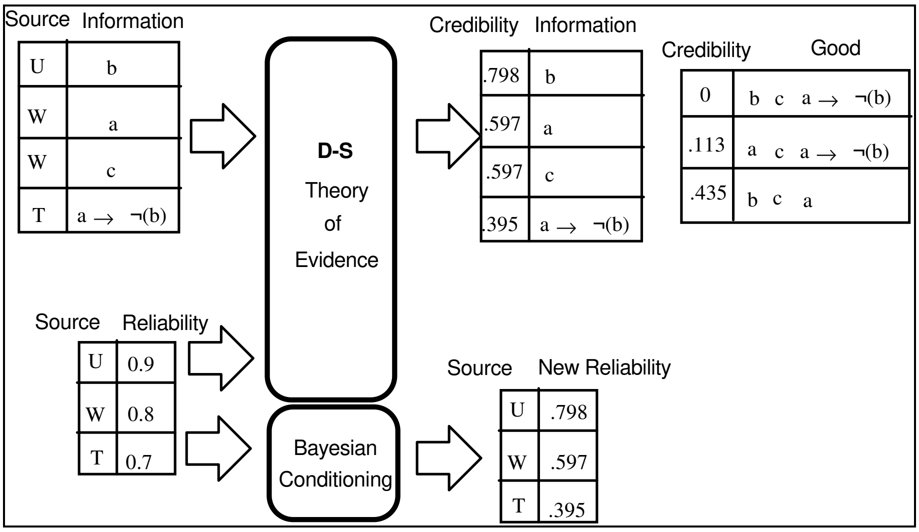

Fig. 1. Basic I/O of the belief-function mechanism

$$m_i(X) = \begin{cases} R_i & \text{if } X = [s_i] \\ 1 - R_i & \text{if } X = \Omega \\ 0 & \text{otherwise} \end{cases}$$

All these *bpa*s will be then combined through the Dempster Rule of Combination:

$$m(X)=m_1(X)\otimes\ldots\otimes m_n(X)=\frac{\sum_{X_1\cap\ldots\cap X_n=X} m_1(X_1)\cdot\ \ldots\ \cdot m_n(X_n)}{\sum_{X_1\cap\ldots\cap X_n\neq\emptyset} m_1(X_1)\cdot\ \ldots\ \cdot m_n(X_n)}$$

From the combined *bpa m*, the credibility of a set of sentences *s* is given by:

$$Bel(s) = \sum_{X\subseteq[s]} m(X)$$

This is the technique adopted to perform O1 in section 1.

Learning each others' reliability

In figure 1, we see another output of the mechanism, obtained through Bayesian Conditioning: the set {<S_1, NR_1>,…,<S_n, NR_n>}, where NR_i is the new reliability of S_i.

Following Shafer and Srivastava, we defined the "a priori" reliability of a source as the *probability* that the source is reliable. These degrees of probability are "translated" into belief-function values on the given pieces of information. However, we may also want to estimate the sources' "a posteriori" degree of reliability from the cross-examination of their evidences. To be congruent with the "a priori" reliability, also the "a posteriori" reliability must be a *probability* value, not a belief-function one. This is the reason why we adopt the Bayesian Conditioning instead of the Theory of Evidence to calculate it. Let us see in detail how it works here.

Let us consider the hypothesis that only the sources belonging to $\Phi \subseteq S$ are reliable. If the sources are independent, then the probability of this hypothesis is $R(\Phi)= \prod_{S_i \in \Phi} R_i \cdot \prod_{S_i \notin \Phi} (1 - R_i)$. We could calculate this "combined reliability" for any subset of S. It holds that $\sum_{\Phi \in 2^S} R(\Phi)=1$. Possibly, the sources belonging to a certain Φ cannot all be considered reliable because they gave contradictory information, i.e., a set of information items s such that $[s]=\varnothing$. In this case, the combined reliabilities of the remaining subsets of S are subjected to the Bayesian Conditioning so that they sum up again to "1"; i.e., we divide each of them by "1- $R(\Phi)$". In the case where there are more subsets of S, say $\Phi_1, \ldots, \Phi_l$, containing sources which cannot all be considered reliable, then $R(\Phi)=R(\Phi_1)+ \ldots +R(\Phi_l)$. We define the revised reliability NR_i of a source S_i as the sum of the conditioned combined reliabilities of the "surviving" subsets of S containing S_i. An important feature of this way to recalculate the sources' reliability is that if S_i is involved in contradictions, then $NR_i \leq R_i$, otherwise $NR_i=R_i$.

This is the technique adopted to perform O2 in section 1.

2.3 S4

S4 also deals with uncertainty, but at the goods' level, i.e. it extends the ordering $\leq_{KB}$, defined at S3 on the assumptions, into an ordering $\leq_G$ onto the goods.

The belief-function formalism is able to attach directly a degree of credibility to any good g, bypassing S4 in our framework (see again Fig. 1). The problem is that if a good contains only part of the information items supplied by a source, then its credibility is null (see [1] for an explanation). Unfortunately, the event is all but infrequent, so that often the credibility of all the goods is null. Thus, to perform S4 we tried the "best-out" method [8] and the "average" method.

- *best-out method*. Let G' and G" be two goods and let g' and g" be the most credible assumptions (according to $\leq_{KB}$), respectively, in KB\G' and KB\G". Then G''$\leq_G$G' iff $g' \leq_{KB} g''$.
- *average method.* The goods are ordered according to the average credibility of their elements

A main difference between the two is that using the *average credibility* method the preferred good(s) may no longer necessarily contain the most credible belief(s).

2.4 S5

The good with the highest priority in $\leq_G$. will be selected in S5 as the preferred good. In case of ties, the revised cognitive state may be based on either one of the goods with the same highest priority (randomly selected) or on their intersection (see [8] and [15]). This latter case means rejecting all the conflicting but equally credible information items. The result is not a good (it is not maximally consistent) and thus implies rejecting more assumptions than necessary to restore consistency.

3 The Simulation Experiment

The task given to the group is that extracting as much correct knowledge as possible from a corrupted knowledge repository. The knowledge repository is trivially implemented as a couple of databases containing the same quantity of information items: a database holds correct information, the other one contains the negations of the information items in the correct database. Agents cannot distinguish the two databases. Each agent is characterized by a degree of "capacity" (between 0 and 1) that will be adopted as the frequency with which it accesses (unconsciously) the correct database.

While learning from each others, agents exchange information, thus, they should be equipped with a communication module (Comm) specifying the three basic communication policies C1÷C3 (section 1). These policies will affect deeply the global cognitive behavior of the group. We adopted the following rules:

H1. the communication is peer-to-peer and the recipient of the communication is selected randomly (C3)

H2. the piece of information to communicate is chosen randomly among those in the preferred good (C2)

H3. agents communicate (and access the databases) once per simulation cycle (C1)

H4. agents do not communicate to the others the sources from which they received the pieces of information, but they present themselves as completely responsible for the knowledge they are passing onto the others; the receiving agents consider the sending agents as the sources of the information they are receiving

H5. agents exchange opinions regarding the credibility of the information they are providing to the others, but they do not exchange opinions regarding the reliability of the other agents with whom they got in touch

With H4 we extend the scope of responsibility: an agent is responsible not only for the information that it provides to the network as the original source, but also for the information that it receives from other agents and, *retaining it credible*, passes it on to the others. With H5 we limit the range of useful information: an agent's opinion regarding the others' (and its own) reliability is drawn out from pure information regarding the knowledge domain under consideration, not from indirect opinions.

For each agent we evaluated three parameters: its average *reliability*, the *quality* and the *quantity* of the beliefs in its preferred good.

Quality and quantity are evaluated as differences w.r.t. the case without interaction.

- *Quality* = $Q - Q_{\text{without communication}}$

  where $Q = \dfrac{|\text{true propositions in G}| + |\text{false propositions outside G}|}{|\text{propositions in KB}|}$

- *Quantity* = $|\text{true propositions in G}| - |\text{true propositions in G}|_{\text{without communication}}$

- *Reliability* of the i^{th} agent $R_i = \frac{\sum_{k=1}^{|agents|} r_{ki}}{|agents|}$ where r_{ki} is the reliability of the i^{th} agent, estimated by the k^{th} one.

Due to severe complexity and duration problems, we fixed at five the number of agents. Each of the two databases contains 100 propositions. From some tentative with 40 accesses for each agent, we realized that 25 accesses were sufficient to reach stable values. To reduce the effects of casuality, each of the sixteen simulations has been repeated 20 times. We calculated the average of each parameter, for each cycle, over the twenty repetitions.

3.1 Experiment 0

In order to evaluate the performances of Bayesian Conditioning to estimate the reliabilities, we made a series of simulations focusing our attention on the perspective of a single agent. This agent receives information from four agents with different degree of capacity. If some of them have a degree of capacity clearly less than the others, Bayesian Conditioning is able to detect the less capable agents (see fig 2; case with only one agent with a very low degree of capacity).

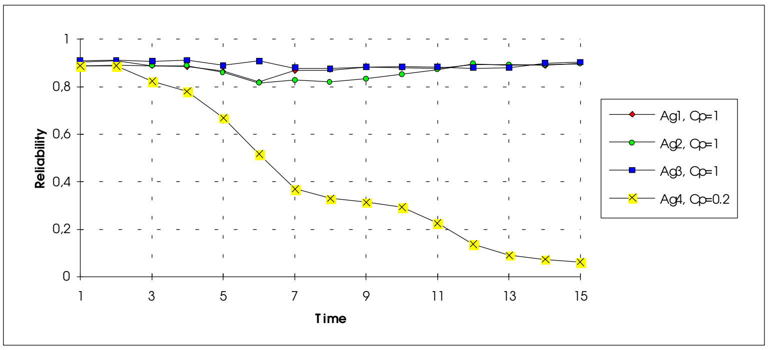

Fig. 2

More generally, when the agents' capacity are sufficiently differentiated (with differences greater than 0.2) there is an acceptable correspondence between estimated reliability and real capacity (see the case in fig 3), but this correspondence is lost when the degrees of the agents' capacity are closed to each others (see the case in fig 4).

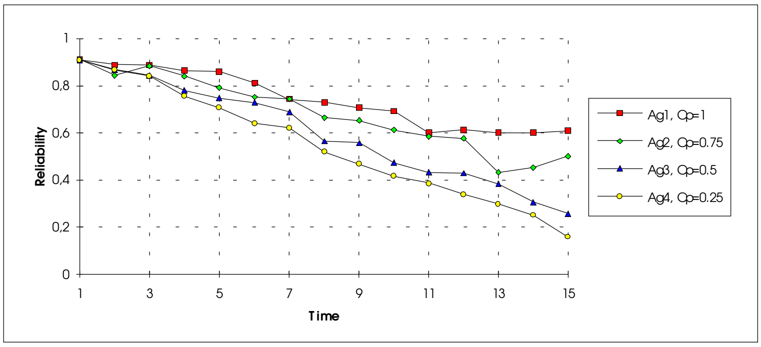

Fig. 3

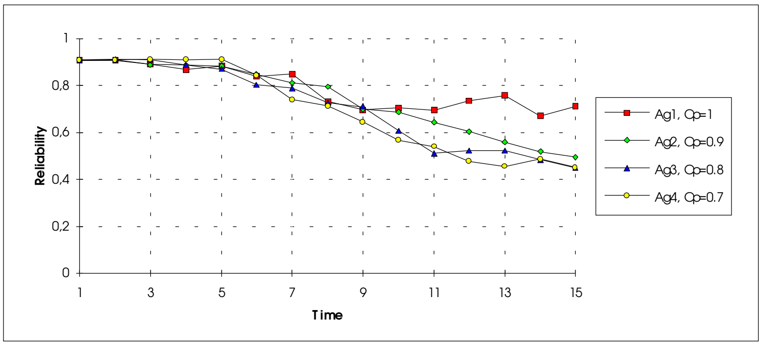

Fig. 4

Experiment 0 showed that, for our purposes, Bayesian Conditioning was an acceptable way to estimate the degrees of reliability of information sources affected by limited degrees of correctness (capacity, competence, sincerity ...).

3.2 Experiment A

We began the "distributed" experiments by making two series of simulations. In the first one we gave to the five agents the same capacity. We made six simulations, from capacity 0 to capacity 1 (fig. 5). In the second series (ten simulations) we gave to the agents the capacities reported in fig. 6.

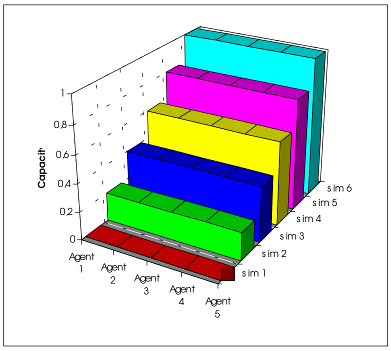

Fig. 5

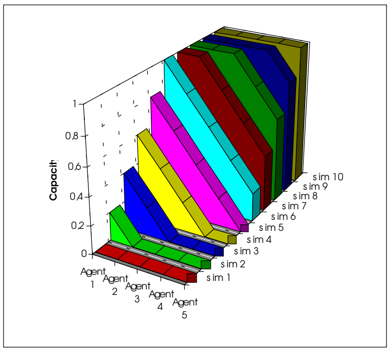

Fig. 6

The results obtained can be summarized as follows.

Quality

Interaction increases the quality of an incapable agent's cognitive state and decreases the quality of a capable one. A typical trend is reported in fig. 7, which refers to the sim 5 in fig. 6 (decreasing capacity from 1 to 5). The average quality (Av) stands at zero. This means that if there are few incapable agents, then they gain much in correctness while the others lose very little.

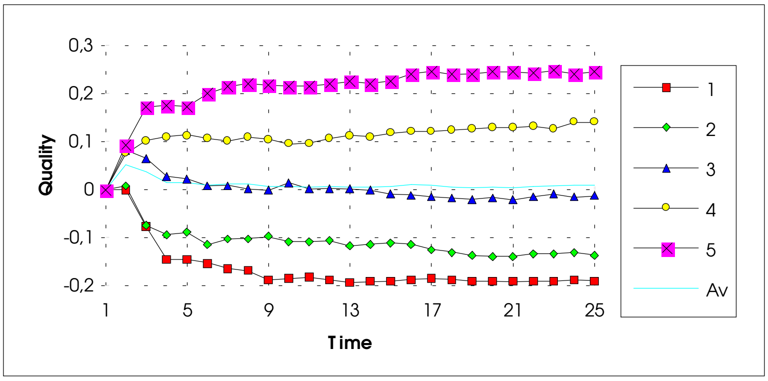

Fig. 7

Quantity

Interaction always increases quantity. However, incapable agents gain more than the capable ones. A typical trend is reported in fig. 8 (sim 5 fig. 6; decreasing capacity from 1 to 5).

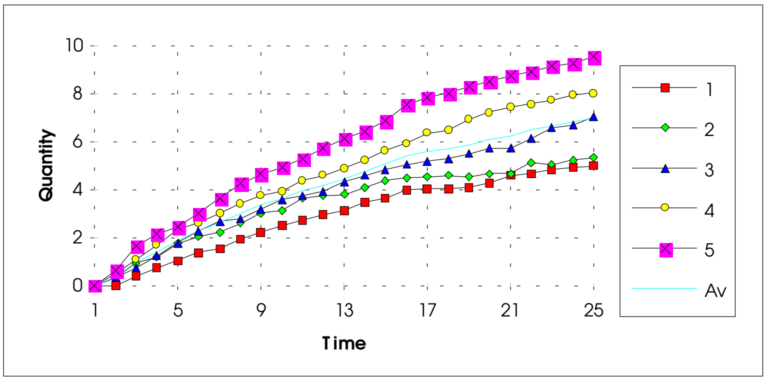

Fig. 8

Reliability

Reliabilities follow typical trends as the one reported in fig. 9 (sim 7 in fig. 6; decreasing capacity from 1 to 5). We see that all the agents lose reliability, but there is what we called "majority effect". If the average capacity of the group is greater than 0.5, then the capable agents lose less than the incapable ones.

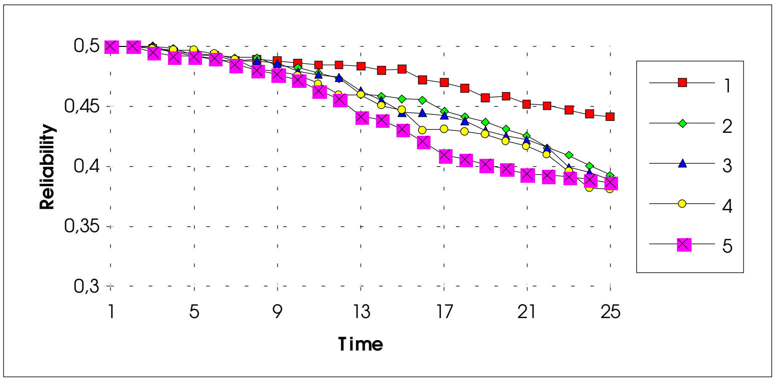

Fig. 9

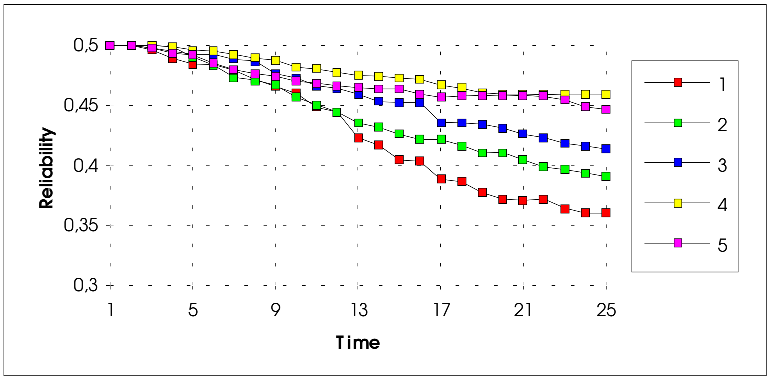

Fig. 10

On the opposite, if the average capacity is less than 0.5, than the capable agents lose more than the incapable ones, as in fig. 10 (sim 3 fig. 6). In this case, since the capable agents are in minority, they are regarded as unreliable by the others. The result is that the estimated reliability is quite inversely proportional to the real capacity.

These results were expected, since this group looks like a class without teacher, or a scientific community without experimental evidence. They are positive in the sense that the members of the group acquire knowledge without loosing (averagely) in correctness.

3.3 Experiment B

Subsequently, we introduced a "teacher" agent with the following features:

- capacity = 1 (it accesses only the correct database)
- for each k, r_k= 1 (all the agents know that the teacher is absolutely reliable)
- it transmits but doesn't receive information (the teacher's cognitive state will not be contaminated by the others)

We repeated the simulations with the teacher and five agents as in the experiment A. Results may be summarized as follows.

Quality

The main difference w.r.t. the case without teacher is that now the average quality is positive and lightly increasing. This is due to the presence of the teacher which, occasionally (communications took place randomly) gave to the others the possibility to choose in the right way. A typical trend is showed in fig. 11 (capacity decreases from agent 1 to agent 5).

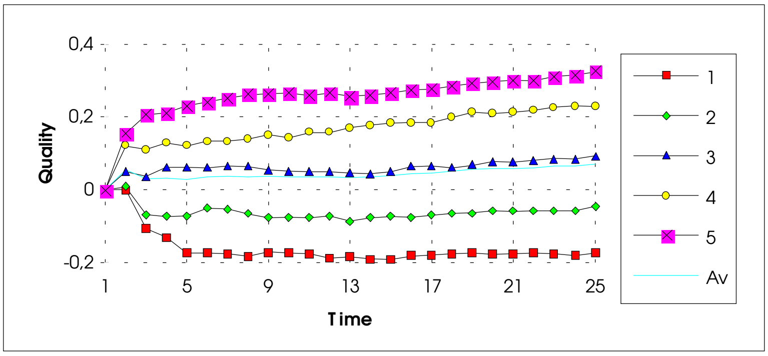

Fig. 11

Quantity

Even in this case, interaction increases the amount of correct data in the cognitive state of every agent, and the gain is inversely proportional to the capacity of the agent. The effect of the teacher is that the gain is higher than before for any agent. A typical trend is showed in fig. 12 (decreasing capacity from 1 to 5).

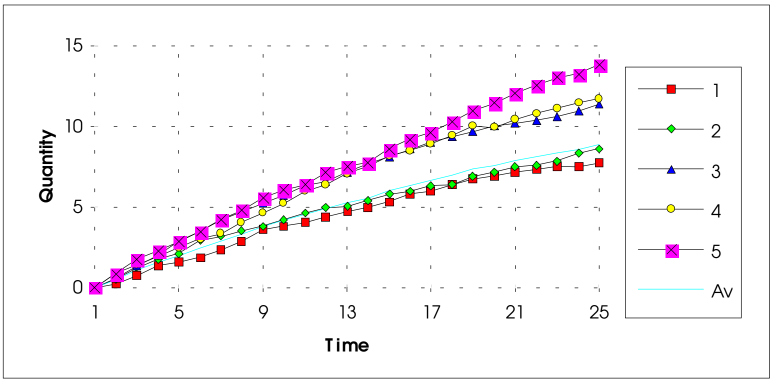

Fig. 12

Reliability

The correspondence between estimated reliability and real capacity holds when the average capacity is greater than 0.5. A typical trend is showed in fig. 13 (decreasing capacity from 1 to 5). Under that value the situation becomes confused. What we called “majority effect” in the experiment A is no longer appreciable when there is a teacher in the group.

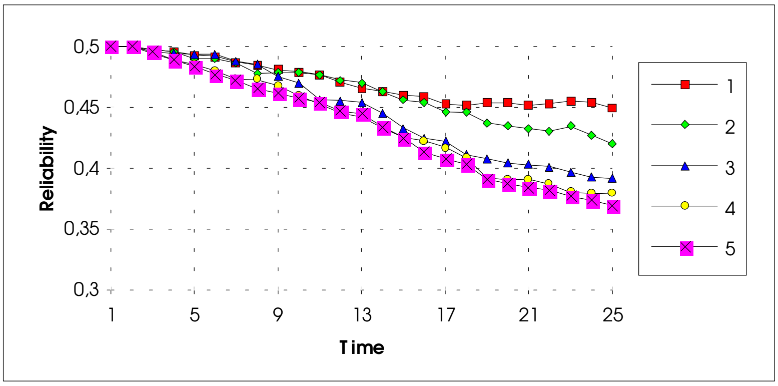

Fig. 13

3.4 Experiment C

In the previous experiments, the characteristics of the agents were static; their capacity did not change during the simulation. We were curious to see if the group would have been able to realize that some of its member changed their degree of capacity, and how long would have it taken to the group to be aware of the change. After several simulations with a agent decreasing or increasing its capacity, we realized that only high quality groups were able to perceive the change. For groups with an average capacity less than 0.6, the situation at the time of the change was sufficiently chaotic to hide it. This implies that only decreases of capacity will be perceived by the group. The graph in fig. 14 reports the "quality" trends of five agents with capacity 0.98, where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.8.

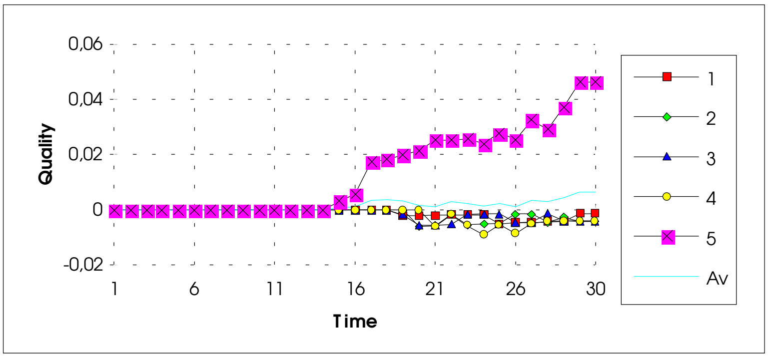

Fig. 14

The graph in fig. 15 reports the "quantity" trends of five agents with capacity 0.8 where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.6.

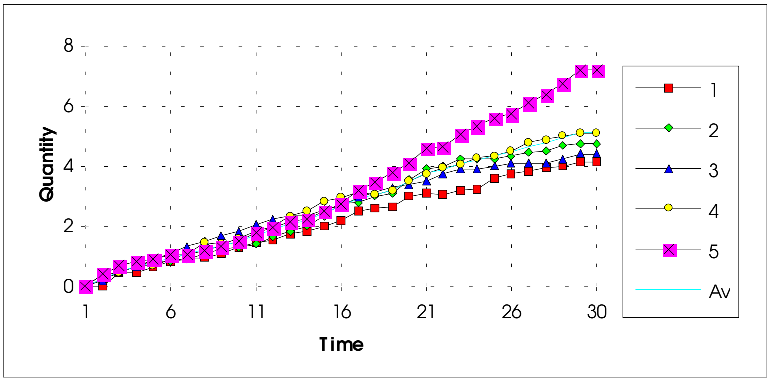

Fig. 15

The graph in fig. 16 reports the “reliability” trends of five agents with capacity 0.98 where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.8.

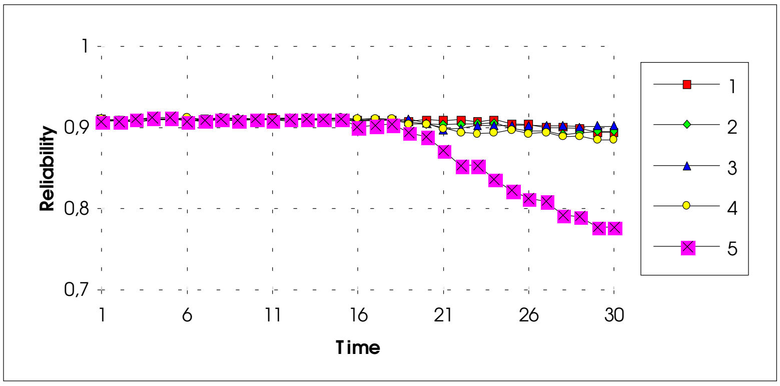

Fig. 16

As an example of a low quality group, fig. 17 reports the “reliability” trends of five agents with capacity 0.6 where the fifth one, at the 15th simulation cycle, decreased its capacity to 0.4. It is evident that the group was not able to perceive the change.

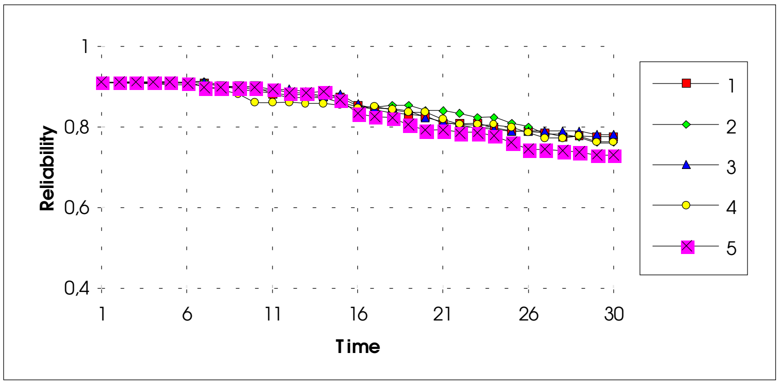

Fig. 17

We also tried to study the effects of the deepness of the reciprocal acquaintance: the longer the period spent together, the more reluctant should be the agents in changing their opinions regarding their companion who decreased its capacity. As a matter of fact, fig. 18 shows the "reliability" trend of the agent who changed its capacity from 0.98 to 0.8 at different time points: t=0, 5, 10, 15, 20 and 25. We can see that the later the change, the lighter the curve's slope: the group which appreciated the agent's reliability for a longer period has more inertia to change its opinion regarding that agent.

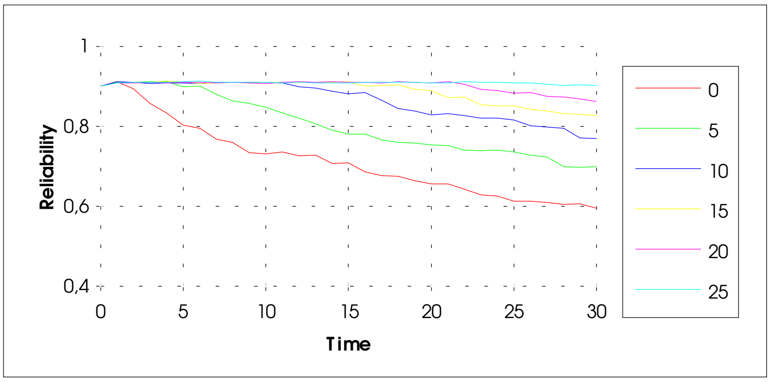

Fig. 18

4 Conclusions

Learning other agents' (and its own) relative degrees of reliability is crucial when interacting within a group of partially competent (correct, sincere, etc.) information sources. The technique proposed in this paper is a simple application of Bayesian Conditioning. Results have been extracted from the first outcomes of a broader minded simulation experiment: studying the global cognitive behavior of a group whose members adopt definite local mechanisms to revise their opinions (regarding agents and knowledge), local policies of interaction and global strategies of integration. The local method to revise beliefs is a form of assumption-based reasoning which adopts the Dempster's Rule of Combination to integrate, information coming from different sources, and embodies Bayesian Conditioning to evaluate the reliability of those sources. Our feeling is that this latter technique is very intuitive and simple. However, while it shows interesting properties from a distributed perspective (the global cognitive behavior of the group seems sociologically intuitive and foreseeable), we still need to deepen into its performances as a centralized mechanism to evaluate the reliability of partially correct information sources (e.g., sensors). We will soon report in another paper the results of a specific simulation experiment that we are still carrying on.

References

[1] A.F. Dragoni, P. Giorgini, "Belief Revision through the Belief Function Formalism in a Multi-Agent Environment", Third International Workshop on Agent Theories, Architectures and Languages, LNAI Series, Springer-Verlag, 1997.

[2] Dragoni A.F., Ceresi, C. and Pasquali, V., A System to Support Complex Inquiries, in Proc. of the “V Congreso Iberoamericano de Derecho e Informatica”, La Habana, 6-11 march 1996.

[3] Alchourrón C.E., Gärdenfors P., and Makinson D., On the Logic of Theory Change: Partial meet Contraction and Revision Functions, in The Journal of Simbolic Logic, 50, pp. 510-530, 1985.

[4] P. Gärdenfors, Knowledge in Flux: Modeling the Dynamics of Epistemic States, Cambridge, Mass., MIT Press, 1988.

[5] P. Gärdenfors, Belief Revision, Cambridge University Press, 1992.

[6] Williams M.A., Iterated Theory Base Change: A Computational Model, in Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence, pp. 1541-1547, 1995.

[7] W. Nebel, Base Revision Operations and Schemes: Semantics, Representation, and Complexity, in Cohn A.G. (eds.), Proc. of the 11th European Conference on Artificial Intelligence, John Wiley & Sons, 1994.

[8] Benferhat S., Cayrol C., Dubois D., Lang J. and Prade H., Inconsistency Management and Prioritized Syntax-Based Entailment, in Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence, pp. 640-645, 1993.

[9] Dubois D. and Prade H., A Survey of Belief Revision and Update Rules in Various Uncertainty Models, in International Journal of Intelligent Systems, 9, pp. 61-100, 1994.

[10] Dragoni A.F., Mascaretti F. and Puliti P., A Generalized Approach to Consistency-Based Belief Revision, in Gori, M. and Soda, G. (Eds.), Topics in Artificial Intelligence, LNAI 992, Springer Verlag, 1995.

[11] de Kleer J., An Assumption Based Truth Maintenance System, in Artificial Intelligence, 28, pp. 127-162, 1986.

[12] Shafer G. and Srivastava R., The Bayesian and Belief-Function Formalisms a General Perpsective for Auditing, in G. Shafer and J. Pearl (eds.), Readings in Uncertain Reasoning, Morgan Kaufmann Publishers, 1990.

[13] Shafer G. (1976), A Mathematical Theory of Evidence, Princeton University Press, Princeton, New Jersey.

[14] Shafer G., Belief Functions, in G. Shafer and J. Pearl (eds.), Readings in Uncertain Reasoning, Morgan Kaufmann Publishers, 1990.

[15] Benferhat S., Dubois D. and Prade H., How to infer from inconsistent beliefs without revising?, in Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence, pp. 1449-1455, 1995.

[16] Dragoni A.F., Belief Revision: from theory to practice, to appear on “The Knowledge Engineering Review”, Cambridge University Press, 1997.

[17] A.F. Dragoni, P. Giorgini and P. Puliti, Distributed Belief Revision vs. Distributed Truth Maintenance, in Proc. 6th IEEE Conf. on Tools with A.I., IEEE Computer Press, 1994.

[18] R. Reiter, A Theory of Diagnosis from First Principles, in Artificial Intelligence, 53, 1987.