Supporting Evolution in Learning Information Agents

D. Rosaci
DIMET, Università "Mediterranea" di Reggio Calabria
Loc. Feo di Vito
89060 Reggio Calabria (Italy)
Email: domenico.rosaci@unirc.it
Tel: (++39) 0965875313
Fax: (++39) 0965875238

G.M.L. Sarné
DIMET, Università "Mediterranea" di Reggio Calabria
Loc. Feo di Vito
89060 Reggio Calabria (Italy)
Email: sarne@unirc.it
Tel: (++39) 0965875438
Fax: (++39) 0965875238

***Abstract*—Learning agents can autonomously improve both knowledge and performances by using learning strategies. Recently, a strategy based on a cloning process has been proposed to obtain more effective recommendations, generating advantages for the whole agent community through individual improvements. In particular, users can substitute unsatisfactory agents with others provided with a good reputation and associated with users having similar interests. This approach is able to support an evolutionary behaviour in the community that allows the better agents to predominate over the less productive agents. However, such an approach is user-centric requiring a user's request to clone an agent. Consequently, the approach slowly generates modifications in the agent population. To speed up this evolutive process, a proactive mechanism is proposed in this paper, where the system autonomously identifies for each user those agents that in the community have a good reputation and share the same interests. The user can check the clones of such suggested agents in order to evaluate their performances and to adopt them. The results of preliminary experiments show significant advantages introduced by the proposed approach.**

I. Introduction

A learning information agent autonomously and proactively analyzes distributed and heterogeneous information sources for building and updating its knowledge and providing its user with useful recommendations [1], [21], [23]. In other words, a learning agent should be capable to improve its performances in time. Recently, some authors proposed communities of intelligent information agents able to modify both their behaviours and their internal knowledge through the use of learning methodologies [14], [19], [20]. For example, in [3] learning agents improve their individual performances by means of a reciprocal mutual monitoring in order to obtain suggestions about the best agents which cooperate and/or integrating their knowledge. While in [20], in presence of an unsatisfactory recommender agent its owner can enrich its knowledge with that of other agents having similar interests in the community. Differently, other proposals in multi-agent systems (MASs) [5], [26], [27], [32] adopt reputation models rather than similarity measures both to promote agent cooperation and to select the most promising agents for collaboration.

However, while the learning capabilities of an agent produce an improvement in the agent performances, they do not contribute to advantage also the other agents belonging to the same community. On the contrary, biologic "evolution" implies that profitable changes in a population are permanently inherited and spread over the future generations [10], [12] transcending the lifetime of single individuals [8]. In such a way, evolution happens when the genetic material changes from one generation to the next. Differently, occasional changes in individual entities do not produce evolutive processes.

By considering the peculiarities both of the learning agents systems and of the "biologic" environments, in [22] an *evolutionary* framework, called *EVolutionary Agents* (EVA), based on cloning processes and exploiting a reputation model has been proposed. In EVA individual agent's improvements in generating recommendations can induce improvements in the whole learning agent population.

The evolutive technique adopted in EVA is similar to the biologic asexual reproductive processes generating clones that initially are the exact copies of their parents. On the contrary, in the sexual reproductive processes the parents' DNA are joined to obtain an individual that mixes their characteristics. The nature is mainly oriented on the sexual reproduction because individual changes, in response to environmental changes, are spread on the next generations more quickly than via asexual reproduction. Cloning can be most effective in difficult or hostile environments in presence of strong selective processes. In this way, the cloning with a suitable mechanism of selection can implement a simple, but effective, mechanism able to induce evolutive phenomena in a population. In EVA [22] cloning and selection (based on reputation criteria) techniques are adopted in a MAS for allowing a user to require the substitution of unsatisfactory agents with other agents having both similar interests and good reputation in the community.

However, the EVA approach is basically user-centric since it compulsorily requires a user's request for cloning and substituting his/her agent. The consequence is that the evolutive processes in the agent population occur slowly and, for speeding up them, in this paper it is proposed a proactive mechanism. More in detail, the system autonomously identifies for each user those agents that in the community have a good reputation and share the same interests. Then the user can evaluate the clones of such promising agents in order to compare their performances with those of his/her current agents and, possibly, adopting them (or in substitution of his/her current agents). Preliminary experiments in a leaning agent-based recommender system show that the performances of the agent population quickly improve (i.e., the recommendations are most effective) when the new strategy of promoting the most performing agents among the users is activated.

The paper is organized as follows. Related work about mutual agent monitoring are presented in Section II. Sections III and IV present an overview of the EVA framework and of the new evolutive strategy, respectively. Some experiments are presented in Section V and, finally, in Section VI some conclusions are drown.

II. Related Work

In the context of the autonomous agents, a relevant issue is represented by the monitoring learning agents that are able to learn and keep up with a dynamically changing world, also interacting with one another [29]. In the literature, some recent works deal with such a problem, as in [4], [31], where each agent is provided with an internal representation of both interests and behaviour of its owner, usually called ontology. To implement mutual monitoring for choosing the best agents for knowledge-sharing purposes, the inter-ontology properties have to be detected. For instance, some approaches use as inter-ontology properties the similarity between ontology concepts [2], also by determining their synonymies and homonymies [4], or, in addition to similarity, other properties defined on the whole agent community, as the *reputation* of an agent within its MAS [3]. Similarly to our proposals, these approaches try to introduce a form of cooperation in a MAS, based on a mutual agent monitoring. However, differently from our approach, none of the cited proposals considers the possibility that, based on a learning process, the effectiveness of the agents can evolve in time.

The approach proposed in [20] induces logical rules to represent agent behaviour in the ontology by means of a connectionist ontology representation, deriving from [7], based on neural-symbolic networks. In this scenario, the mutual monitoring is realized by introducing a similarity measure of the agent ontology that considers also logical representation of the agent behaviour. In this manner, the learning activity can improve in time the effectiveness of the agent but, differently from our approach, this improvement does not involve the whole system with a cooperative behaviour among the agents.

For learning agents, the approach presented in [14] describes an evolutionary MAS to study Web sites usability and navigation paths. Based on the past users' Web activities, such a system *i)* builds a users' model for trying to navigate among URLs, *ii)* simulates the browsing process and *iii)* analyzes the Web pages that can belong to possible paths between two URLs. This proposal, similarly to our one, exploits evolutionary techniques to make adaptive the behaviour of a MAS but without to support mutual monitoring among agents. On the contrary, this characteristic is the main feature in EVA to implement an effective agent cooperation.

Furthermore, trust and reputation within an agent context are concepts widely proposed in the literature (the interesting reader can refer to [5], [6], [13], [16], [18], [24], [26], [27], [32] for a most comprehensive overview). In learning agent-based recommender systems, recently a reputation-based approach has been proposed in [9] to lead the evolution of a community of information software agents with the purpose of improving the agent communication. Although this approach is similar to our in that the agent evolution is driven by a reputation mechanism, it does not realize any evolutionary behaviour.

III. The EVA Framework Overview

This section presents an overview of the EVA framework. The basic idea exploited in EVA is that in presence of an unsatisfactory agent a user can require the system to provide him/her with one or more suitable and performing agents. For each agent in the EVA framework, the system computes a score based on both the similarity with the user's interests and its reputation (considered likely to a genetic component) in the community. The agents having the best scores are cloned and sent to the requester user. In the following, let u be a generic user belonging to the users' community U and assisted by a set $A_u = \{a_i \mid i = 1 \cdots n_u\}$ of n_u information software agents a_i supporting his/her Web activities with recommendations.

A. Evaluation of the user's satisfaction

For each Web page visited by u, each agent a_i generates for him/her some suggestions (i.e., Web links). Considering the life of a_i, let R_i and L_u be the sets, partially overlapping, of the Web links suggested by a_i to u and those selected by u, respectively. To evaluate the quality of these recommendation sets, *precision* and *recall* measures [11], [17] have been used. Precision is the fraction of the recommendations considered as relevant by u with respect to the potentially relevant recommendable links stored in L_u. Recall is the fraction of the links actually selected by u and successfully recommended by a_i but alone it is meaningless because returning all possible links as recommendations it is equal to 1. A good recommender agent should have both high precision and recall values. Precision and recall of R_i can be formally defined as:

$$Pre(R_i) = \frac{|R_i \bigcap L_u|}{|R_i|}; \qquad Rec(R_i) = \frac{|R_i \bigcap L_u|}{|L_u|}$$

To consider together recall and precision, their harmonic mean, known as *F-measure* [30] is used. Weighting the precision with respect to the recall, it is obtained the more general F_β-measure, where β is a non-negative real:

$$F_\beta(R_i) = (1+\beta^2) * \frac{Pre(R_i) * Rec(R_i)}{\beta^2 * Pre(R_i) + Rec(R_i)}$$

In EVA precision, recall and F_β measures are adopted to compute the satisfaction of u for the recommendations

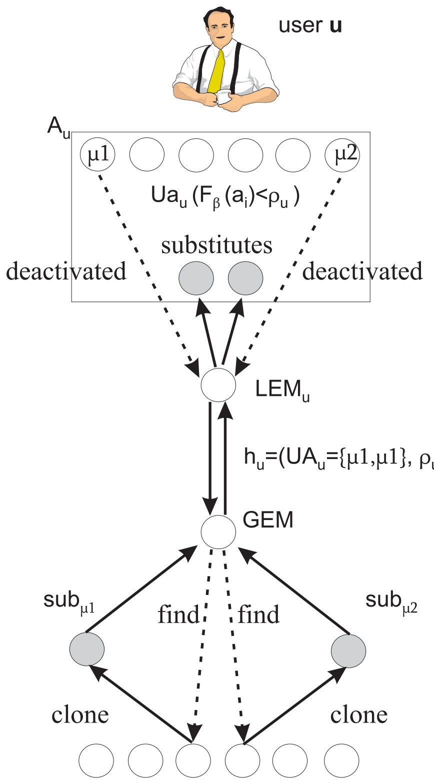

Fig. 1: The evolutionary strategy of the EVA framework

provided both by his/her agent a_i in R_i and by his/her whole agent-set A_u by considering the union of the sets R_i relative to each agent $a_i \in A_u$. Formally:

$$Pre(A_u) = \frac{|\bigcup_{i=1}^{n(u)} R_i \bigcap L_u|}{|\bigcup_{i=1}^{n}(u) R_i|}; \quad Rec(A_u) = \frac{|\bigcup_{i=1}^{n(u)} R_i \bigcap L_u|}{|L_u|}$$

$$F_\beta(A_u) = (1+\beta^2) * \frac{Pre(A_u) * Rec(A_u)}{\beta^2 * Pre(A_u) + Rec(A_u)}$$

Furthermore, the F_β-measure is adopted to synthetically evaluate the user's satisfaction simply by observing the acceptance of the provided recommendations. Other measures as, for instance MAE and ROC, could be used for the same purpose but they require the user to explicitly rate his/her satisfaction. However, the EVA framework confirmed the improvements in the user's satisfaction also with respect to such estimators.

B. Evolutionary strategies to improve user's satisfaction

The EVA framework (depicted in Figure 1), to increase the users' satisfaction about the agents, implements an evolutionary strategy managed by two types of agent, namely: *i)* the *Local Evolution Manager* (LEM_u) agent associated with each user u; *ii)* the *Global Evolution Manager* (GEM) agent associated with the Multi Agent System. The evolutionary strategy is based on the following ideas:

- The satisfaction of a user u for the suggestions provided by his/her agent-set A_u is measured by $F_\beta(A_u)$.
- Each user u can arbitrarily set both the coefficient β, used in computing $F_\beta(A_u)$, and the *satisfaction threshold* ρ_u for $F_\beta(A_u)$ under which u is unsatisfied of the recommendations generated by his/her agent-set.
- For each user u his/her LEM_u agent periodically computes $F_\beta(A_u)$. If $F_\beta(A_u) < \rho_u$ then LEM_u: *i)* identifies the set UA_u of the unsatisfactory agents for which $F_\beta(a_i) < \rho_u$; *ii)* deactivates the agents belonging to UA_u; *iii)* sends a triplet $\langle UA_u, \rho_u, \psi_u \rangle$ (with the set UA_u, the threshold satisfaction ρ_u and the parameter $\psi_u \in [0.0, 1.0]$, that represents how much the user u weights the similarity with respect to the reputation) to the GEM agent; *iv)* requires the substitution of the deactivated agent with other, presumably more satisfactory, to the GEM agent. The GEM agent (see below) will determine a set of substitutes agents based on both their *reputation* in the community and the similarity (represented by ψ_u) with the deactivated agents. For example, if $\psi_u = 0.3$ the user gives a 30% of relevance to the similarity and a 70% of relevance to the reputation.
- The GEM agent maintains a similarity matrix $\Sigma = \{\Sigma_{i,j}\}$, $i, j \in MAS$ where each element belongs to $[0.0, 1.0]$ and represents the similarity between two agents of the MAS computed as in [20]. Moreover, for each agent $a \in MAS$ the GEM agent stores a *reputation coefficient* $r_a \in [0.0, 1.0]$ (see Section III-C) that represents a measure of how much the community considers satisfactory the performances of a. When GEM receives the LEM_u request (i.e., $\langle UA_u, \rho_u, \psi_u \rangle$), it inserts in the set C_μ those agents of the MAS having $F_\beta > \rho_u$ with which to substitute each agent $\mu \in UA_u$. Then, GEM computes for each agent $a \in C_\mu$ the score $s(a, \mu) = \psi_u \cdot \Sigma_{a,\mu} + (1 - \psi_u \cdot r_a)$ and, based on it, chooses as substitute of μ the agent sub_μ with the best score (in the case of equal score, the agent having the best F_β-measure will be chosen).
- The GEM creates, for each agent $\mu \in UA_u$, an agent sub^*_μ cloned by the substitute agent sub_μ and having the same ontology. Similarly that in [20], the ontology of an information agent contains both its categories of interests and the causal implications (i.e., relationships between the considered events) learnt by it during its life. Thus, cloning is the duplication of this information as in the nature is duplicated the genetic material. The clone agent sub^*_μ is then transmitted to the LEM_u agent in substitution of the unsatisfactory agent μ. From now the agent sub^*_μ will be completely independent from its parent μ living in the environment of another user. This way, the agent sub^*_μ monitoring the activity of u probably it will modify its initial personal ontology with new information.

Summarizing, the strategy of EVA consists in permitting to a user u of substituting each his/her unsatisfactory agent μ with another agent $sub^*_\mu \in MAS$ based on a cooperation between the agents LEM_u and GEM. This substitution

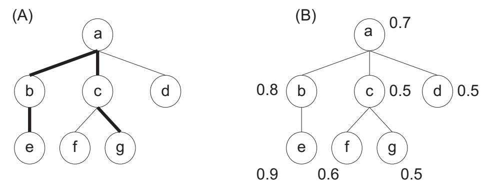

Fig. 2: An example of Descent Tree

should advantage the user u being sub^*_μ the clone of an agent with: *i)* a F_β-measure (computed by its own user) greater than the u's satisfaction threshold ρ_u; *ii)* a top score, computed based on both its reputation in the MAS and its similarity with the substituted agent. The first property assures that the parent agent of sub^*_μ satisfies its own user but not that its clone will produce an F_β-measure satisfactory for u that has a different perception of the satisfaction. The second one guarantees both that the parent agent of sub^*_μ has a good reputation in the community and that its personal ontology is similar to that of the agent μ. Together, these properties provide u with new agent-sets potentially able to improve the $F_\beta(A_u)$ measure.

C. Agent's reputation in the EVA environment

In MASs the reputation (i.e., the opinion of an agent about something [25]) has been studied in a lot of models and surveys (see Section II) and, accordingly with [25], three main issues are recognized: *i)* reputation of an agent is a multi-dimensional concept (For instance, the reputation of a good eBay seller summarizes those of having good products, applying suitable prices, giving appropriate products descriptions, providing fast and secure delivery, etc.); *ii)* each agent has a different *ontological* dimension of the reputation (i.e., it weights each aspect of the reputation differently based on its personal point of view); *iii)* in a MAS there are an *individual* (for each agent) and a *social* (for the MAS) dimension of the reputation.

In particular, in EVA the individual dimension of the reputation is only that to provide effective recommendations to the agent's owner and the social dimension is the cloning activity (remember that an agent can be cloned and its clones supporting other users). As possible ontological dimensions (see Section III-A) both the precision and the recall of the recommendations can be identified. Consequently, as a global measure of the individual reputation of the agent a is adopted the $F_{\beta_{u_a}}(a)$ measure that considers both the two ontological dimensions (u_a denotes the owner of a and β_{u_a} the quantitative representation of the consideration of u_a for the precision with respect to the recall).

The agent reputation has also to consider that the evolutionary strategy implies a cloned agent is moved in a new environment. The relationships introduced by the cloning in the set of agents are described by the same terminology adopted to represent genealogical relationships. For instance, in Figure 2-(A) a "genealogical" tree represents a set of agents, associated with the nodes, involved in cloning processes, associated with edges, and where a *parent* is the agent cloned and a *child* is one of its clones. Furthermore, it is possible to define the following formal definition:

Parent and Sibling Agent - Let a be an agent of the community. We denote by $children_a$ the set of one or more clones of this agent. Two agents b and c, both belonging to $children_a$, are called *sibling agents*. Correspondingly, a is called the *parent agent* of each agent belonging to $children_a$

Ancestor Agent - Let a and p be two agents of the community. We say that p is an *ancestor agent* of a if either: *i)* p is the parent agent of a, or *ii)* recursively there is an agent c in the community such that a is a descendant of p via c.

Relatives, Descent Tree and Kinship Degree - Let a and b be two agents of the community. We say that a and b are *relatives* if they share a common ancestor agent p. We call *family* of a, denoted by $\mathcal{F}_a$ the set of all the relatives of a. We define the *Descent Tree* of a, a tree $DT_a = \langle V, E \rangle$ such that *i)* each agent $x \in \mathcal{F}_a$ is associated with a unique vertex $v_a \in V$ and *ii)* each pair (x, y), $x, y \in \mathcal{F}_a$, such that x is the parent agent of y, is associated with a unique edge $e_{x,y} \in E$. Finally, let a and b be two agents, such that they are relatives. We define the *kinship degree* of a and b, denoted by $k_{a,b}$, the length of the path that links a and b in the Descent Tree DT_a.

As a consequence:

1) At the cloning time, each clone b of an agent a (i.e., $b \in children_a$) is identical to a and inherits its reputation.
2) Since b supports a user, different from that of a, its initial inherited reputation will evolve in time taking into account also the satisfaction degree of its current owner. The inherited reputation and the individual satisfaction are combined in a unique, global, measure of reputation.
3) For the cloning processes, each agent a belongs to a *family* of relatives (i.e., the descent tree DT_a) with which a shares some similarities inherited from the cloning process and that affect its performances. This introduces a social component in the computation of the reputation.

These observations are summarized in the *reputation coefficient* r_a associated with each agent a, with $r_a \in [0.0, 1.0]$ (where 1.0 means a complete reliability of a). This coefficient is weighted using the F_β measures of all the n agents belonging to the descent tree DT_a. Each contribute due to an agent b is weighted in a decreasing manner, based on the kinship degree k between a and b in DT_a, by a coefficient equal to $1/(k_{a,b} + 1)$. This way, the contribution to the satisfaction obtained by each other relative is as smaller as higher is the kinship degree with respect to a. More formally:

$$r_a = \frac{\sum_{b \in \mathcal{F}_a} \frac{F_{\beta_b}(b)}{k_{a,b}}}{\sum_{b \in \mathcal{F}_a} \frac{1}{k_{a,b}}}$$

For example, in Figure 2-(B), the agent e has a F_β-measure (i.e., satisfaction) equal to 0.9 but a reputation of 0.696.

IV. THE NOVEL EVA STRATEGY

To speed up the evolutive process in the agent community a new strategy has been implemented. More in detail, in this new approach, the GEM agent *i)* has to satisfy the user's request to substitute his/her unsatisfactory agents, as in the native EVA strategy, and *ii)* proposes to the user of testing those agents that potentially could enter in his/her agent set in substitution of other agents or in addition to them. In order to perform this proactive mechanism, the native EVA strategy presented in Section III-B is modified as follows:

- The information that the LEM_u agent of each user u sends to the GEM agent are now represented by a tuple $\langle UA_u, \rho_u, \psi_u, T_u, N_u \rangle$ where the first three parameters have the same meaning described in Section III-B, while T_u and N_u are two u's parameters that respectively specify the time (expressed in days) between two consecutive test sessions and the number of agents, ranging in $[0; N_g]$, that u desires to test for each test session (Note that 0 means that u does not want to test any agent, while N_g is the maximum number of agents to test in a single test session and it is a system parameter).
- The GEM agent exploits its similarity matrix Σ and the agents' reputation scores to select for each user u, accordingly to his/her parameters ρ_u, ψ_u, T_u and N_u, a set of agents to clone for a new u's test session.
- After each test session the LEM_u agent evaluates the performances of each clone proposed by the GEM agent. For the agents that really increase the user's satisfaction they can be added to the own agent-set A_u or substitute the less performing agents in A_u.

V. Experimental results

In this section some experiments devoted to test in a MAS the novel strategy implemented in EVA are presented. Experiments have been carried out, similarly to that performed to evaluate the native EVA strategy (see III-B) in [22]), on the top of the CILIOS recommender system [20] for suggesting Web pages to users. In particular, each recommended Web page i is associated with two rates, ranging in $[1, \cdots, 5]$, to represent both the relevances of i for the user esteemed by the system (p_i) and explicitly provided by the user after his/her visit to i (r_i).

The experiments have involved two sets of real users adopting the new and the old EVA strategy, respectively. Furthermore, each set in its turn is constituted of three test-subsets of different cardinality, XML Web sites publicly available at [15] have been exploited and each agent has been provided with a personal ontology, like to that in [20], using the concepts stored in [15]. Moreover, each user u is monitored by a CILIOS agent A_u and its LEM_u agent, while the MAS is managed by a GEM agent. The average satisfaction $F(S)$ of each test-subset S of users is computed as $F(S) = \frac{1}{|S|} \cdot \sum_{u \in S} F_1(A_u)$.

The results of the experiments carried out on the EVA framework for the two user sets confirm an evolutive behaviour in terms of average satisfaction in the MAS population that increases according to the number of users belonging to the test-subset (i.e., the probability to provide suitably clones increases). The values of $F(S_i)$ obtained in the tests are shown in Table I. The first row is referred to the recommendations

T	S_1	S_2	S_3
day 05	0.62/0.62	0.65/0.66	0.69/0.70
day 10	0.64/0.66	0.70/0.72	0.75/0.78
day 15	0.66/0.70	0.73/0.76	0.79/0.83
day 20	0.69/0.72	0.75/0.78	0.83/0.85
day 30	0.71/0.74	0.77/0.80	0.86/0.88
day 45	0.74/0.75	0.81/0.82	0.87/0.89

TABLE I: The temporal evolution of the average F-measure (old strategy/new strategy)

generated only by the CILIOS agents for 5 days. The other rows of the table shown the results obtained by activating the EVA agents for 10, 15, 20, 30 and 45 days.

Analyzing the experimental results reported in Table I it is possible to argue that the difference between the two EVA strategies mainly involve the fact that the new approach is faster then the old one to increase in performances. Besides, after 45 days of using the two EVA approaches, the performances of the system improves for a 21-28 percent in average and the differences between them are not significant.

VI. Conclusions

EVA is an evolutionary agent system based on a cloning process that allows a user to increase the own satisfaction level. EVA, in its first version, admitted only that an owner unsatisfied of his/her agents can require to the system of providing him/her with clones of those agents belonging to the community that are considered similar for interests to the requester user, having a good reputation in the whole agent community and potentially effective for him/her. As a consequence, individual agent improvements in providing recommendations involve the whole agent community supporting an evolutionary behaviour and allowing to the better agents to predominate in time over the less productive agents. The core of the EVA strategy is a reputation model, where a clone agent initially inherits the reputation of its parent agent and then it will autonomously evolves in its own environment, using its learning capabilities to increase this "genetic", initial contribution to its reputation. However, this approach slowly produces changes in the agent population. This characteristic is intrinsic of the exploited user-centric approach that needs of a user's request to clone an agent.

To provide a solution to the problem of speeding up the evolutive process implemented in EVA, in this paper a novel proactive strategy is presented. In particular, the system, autonomously and accordingly to the user's preferences, selects those agents candidates (based on reputation and similarity) to potentially improve the performances of the agent-set supporting the user. Periodically some agent clones are proposed for a test session from the system to the user. After each test session those agents that really increase the performances could be added or substituting the less performing agents in the user's agent-set. To verify if this new strategy effectively promotes evolution in EVA quicker than the native approach, an experimental campaign has been realized and the results have been evaluated by using different well-known metrics. The experiments have confirmed the effectiveness of the novel approach showing that the performances increase more quickly with respect to the previous approach, while differences in terms of F-measure are not significant.

REFERENCES

[1] S. Bergamaschi, P. Edwards and P. Petta, *Intelligent Information Agents*, ser. Lecture Notes in Computer Science. Berlin. Germany: Springer, 2003, vol. 2586.

[2] S. Bowers, K. Lin, B. Ludäscher, “On Integrating Scientific Resources Through Semantic Registration”, in *Proc. of the 3rd International Semantic Web Conference, ISWC2004*. Hiroshima, Japan, Nov. 7-11, 2004, pp. 620–634.

[3] F. Buccafurri, L. Palopoli, D. Rosaci and G.M.L. Sarné, “Modeling Cooperation in Multi-Agent Communities”, *Cognitive Systems Research*, vol. 6, no. 3, pp. 171–190, 2004.

[4] F. Buccafurri, D. Rosaci, G.M.L. Sarné and D. Ursino, “An Agent-Based Hierarchical Clustering Approach for E-Commerce Environments”, in *Proc. of the 3rd Conf. on e-Commerce and Web Technologies, (EC-Web2002)*. Aix-en-Provence, France, Sep. 2-6, ser. Lecture Notes in Computer Science. Berlin. Germany: Springer, 2002, vol. 2455, pp. 115–118.

[5] J. Carbo, J.M. Molina and J. Davila, “Comparing Predictions of SPORAS vs. a Fuzzy Reputation Agent System”, in *Proc. of 3rd International Conference on Fuzzy Sets and Fuzzy Systems*, Interlaken, Switzerland, 2002. pp. 147–153.

[6] R. Conte and M. Paolucci, *Reputation in Artificial Societies: Social Beliefs for Social Order*. Kluwer Academic Publishers, 2002.

[7] A.S. d’Avila Garcez and G. Zaverucha, “The Connectionist Inductive Learning and Logic Programming System”, in *Proc. of the 3rd International Conference on Fuzzy Sets and Fuzzy Systems*, Interlaken, Switzerland, 2002. pp. 59–77.

[8] D.J. Futuyma, *Evolution*. Sunderland: Sinauer Associates, 2005.

[9] S. Garruzzo and D. Rosaci, “Agent Clustering Based on Semantic Negotiation”, *ACM Trans. on Autonomous Adaptive Systems (TAAS)*, vol. 3, no. 2, pp. 1–40, 2008.

[10] S.J. Gould, *The Structure of Eveolutionary Theory The Structure of Eveolutionary Theory*. Belknap Press, 2002.

[11] J.L. Herlocker, J.A. Konstan, L.G. Terveen and J.T. Riedl, “Evaluating Collaborative Filtering Recommender Systems”, *ACM Trans. on Information Systems (TOIS)*, vol. 22, no. 1, pp. 5–53, 2004.

[12] E. Mayr, *What Evolution is*. Basic Books, 2001.

[13] P. Massa, “A Survey of Trust Use and Modeling in Current Real Systems”, in *Trust in E-Services: Technologies, Practices and Challenges*, R. Song, L. Korba, and G. Yee, Eds. Idea Group Publishing, Hershey, Pennsylvania (USA), 2006.

[14] E. Mosqueira-Rey, D. Alonso-Rios, A. Vazquez-Garcia, B. Baldonedo Del Rio, A. Alonso Betanzos and Moret-Bonillo V., “An Evolutionary Multiagent System for Studying the Usability of Websites”, in *Proc. of the International Workshop on Web-based Tools*, Oct. 29-31, Patras, Greece, 2007, pp. 34–41.

[15] (2009) MUADDIB URL. [Online]. Available: http:*//muad.altervista.org*, 2009.

[16] P. Resnick, R. Zeckhauser, E. Friedman and K. Kuwabara, “Reputation Systems”, *Communication of ACM*, vol. 43, no. 12, pp. 45–48, 2000.

[17] V. Raghavan, P. Bollmann and G.S. Jung, “A Critical Investigation of Recall and Precision as Measures of Retrieval System Performance”, *ACM Transactions on Information Systems (TOIS)*, vol. 7, no. 3, pp. 205–229, 1989.

[18] S.D. Ramchurn, D. Huynh and N.R. Jennings, “Trust in Multi-Agent Systems”, *Knowledge Engeenering Review*, vol. 19, no. 1, pp. 1–25, 2004.

[19] J. Reisinger and R. Miikkulainen, “Evolving Neural Networks”, in *Proc. of the 9th Conference on Genetic and Evolutionary Computation (GECCO’07)*, London, England, Jul. 7-11, 2007, pp. 3415–3434.

[20] D. Rosaci, “CILIOS: Connectionist Inductive Learning and Inter-Ontology Similarities for Recommending Information Agents”, *Information Systems*, vol. 32, no. 6, 2007.

[21] D. Rosaci D. and G.M.L. Sarné, “MASHA: A Multi-Agent System Handling User and Device Adaptivity of Web Sites”, *User Modeling and User-Adapted Interaction*, vol. 16, no. 5, pp. 435–462, 2006.

[22] D. Rosaci D. and G.M.L. Sarné, “EVA: An Evolutionary Approach to Mutual Monitoring of Learning Information Agents”, *Applied Artificial Intelligence*, Forthcoming, 2011.

[23] D. Rosaci, G.M.L. Sarnè and S. Garruzzo, “MUADDIB: A Distributed Recommender System Supporting Device Adaptivity”, *ACM Transaction on Information Systems (TOIS)*, vol. 27, no. 4, pp. 1-41, 2009.

[24] J. Sabater-Mir and M. Paoulucci, “On Open Representation and Aggregation of Social Evaluationa in Computational Trust and Reputation Models”, *International Journal of Approximate Reasoning*, vol. 46, no. 3, pp 458–483, 2007

[25] J. Sabater and C. Sierra, “REGRET: A Reputation Model for Gregarious Societies”, in *Proc of the 4th Work. on Deception Fraud and Trust in Agent Societies (AGENTS’01)*, Montreal, Quebec, May 28 - Jun. 1 Canada, 2001, pp. 61–70.

[26] J. Sabater and C. Sierra, “Review on Computational Trust and Reputation Models”, *Artificial Intelligence Review*, vol. 24, no. 1, pp. 33–60, 2004.

[27] M. Schillo, P. Funk and M. Rovatsos, “Using Trust for Detecting Deceitful Agents in Artificial Societies”, *Applied Artificial Intelligence*, vol. 14, no. 8, pp. 825–848, 2000.

[28] A.I. Schein, A. Popescul, L.H. Ungar and D.M. Pennock, “CROC: A New Evaluation Criterion for Recommender Systems”, *Electronic Commerce Research*, vol. 5, no. 1, pp. 51–74, 2005.

[29] P. Stone, “Learning and Multiagent Reasoning for Autonomous Agents”, in *Proc. of the 20th International Joint Conference on Artificial Intelligence*, Hyderabard, India, Jan. 6-12, 2007, pp. 12–30.

[30] C.J. Van Rijsbergen, *Information Retrieval*. Butterworths, 1979.

[31] M. Wooldridge, *Reasoning About Rational Agents*. MIT Press, 2000.

[32] B. Yu and M.P. Singh, in “An Evidential Model of Distributed Reputation Management”, *Proc. of Autonomous Agent and Multi-Agent Systems*, Bologna, Italy, Jul. 15-19, 2002, pp. 294–301.