

Integrating Knowledge Centered MAS through Organizational Links

Gustavo Giménez-Lugo, Jaime Simão Sichman and Jomi Fred Hübner

LTI / EP / USP

Av. Prof. Luciano Gualberto, 158 tv. 3

05508-900 São Paulo, SP, Brazil

Fax: +55 11 3091 5294 - Telephone: +55 11 3091 5397

{gustavo.lugo,jaime.sichman,jomi.hubner}@poli.usp.br

Abstract—*This work presents a model in which concepts in ontologies are extended with organizational information to explicitly express the situation in which they were learned and used. It is discussed how autonomous agents are allowed to reason about concept usage and privacy in terms of organizational constructs, paving the way to reason about social roles in open Web communities. A peer-to-peer application following the model is described. We depart from a specific organization model, MOISE⁺, briefly presented here.*

1. INTRODUCTION

Knowledge sharing presents some challenging issues from representation to reasoning. Context is always elusive, and the emergence and huge usage of the Web has simply emphasized its importance. Even though ontologies are increasingly making part of operational Multi-Agent Systems(MAS), they are seen mostly as a kind of agreed vocabulary of which there are usually no traces of the social interactions that shaped their emergence and change.

This work is situated in a crossroad of several research areas, namely ontology engineering, MAS, and web based communities. It tries to consider how the individual knowledge is affected by the groups the individuals historically belonged to; not forgetting the other way round, i.e., that social agreements upon views can be reached on more abstract levels through social interaction. The approach binds ontologies to a MAS organization model through a mechanism that annotates ontology concepts to hold explicitly the organizational context in which they were learned and used. Such annotations are analogous to the work by Tamma and Bench-Capon [1] that extended the

OKBC (Open Knowledge Base Connectivity) [2] protocol model to allow, among other things, the representation of the behavior of the attributes over time.

The added information allows agents to reason about the organizational context in which portions of an ontology where used. This is a desirable property of an agent since it permits a more detailed search concerning its knowledge representation. Furthermore, dependencies [3] can be used along with temporal associations to specific organizational states to reshape and share such a knowledge representation. Agents can also reason about concept privacy in terms of organizational constructs such as groups and roles.

The content of this work is organized as follows: after a brief revision of ontology definitions, the organizational model called *MOISE*⁺[4] is presented; then, it is outlined how such an organizational model can be used with ontologies through a mechanism for including in the ontologies the organizational informations as they are expressed in *MOISE*⁺; follows a discussion of a reasoning mechanism; finally, a peer-to-peer MAS, applied to information exchange in a community environment over the Web is described, in which agents take advantage of the proposed model.

2. ONTOLOGIES

As suggested by Guarino [5], a conceptualization is an intensional semantic structure which encodes the implicit rules constraining the structure of a piece of reality. An ontology then is a commitment to a particular conceptualization of the world. This refines a definition that states that an ontology is an explicit specification of a conceptualization [6].

The use of ontologies for the explanation of implicit and hidden knowledge is a possible approach to overcome the problem of semantic heterogeneity [7], i.e., within a MAS, agents are characterized by “different views of the world” that can be explicitly defined by ontologies [1]. Interoperability between agents is achieved through reconciliation of these views by a commitment to common ontologies [7]. However, this reconciliation is usually partial, depending on

the organizational environment where the agent is inserted (e.g., we would not expect that a colleague working with us in perfect coordination at the office will accept as her/his own our concepts about beauty). Although there may be a shared vocabulary, there is a need for local ontologies to represent the agents subjective views. Thus, partially global (e.g. group) and local (each agent's) ontologies appear as clearly different, though interrelated. This picture can fit into a flexible organizational metaphor.

This work addresses frame-based representations, more specifically OKBC [2], which is based on *classes*, *slots* and *facets*. Organized into a multiple inheritance (*IS – A*) hierarchy, *classes* (concepts), i.e. collections of objects sharing the same properties, are described in terms of *slots* (attributes) that have *facets*(additional constraints). Follows a simple ad hoc example, that will be used further to illustrate the proposal:

```
c:MachineLearning;
s: NeuronSimulationActivationPotencial
  f: Domain: [0.0,1.0];
  f: Value: [0.4,0.6];
```

In the example, the Machine Learning concept is represented by the class (**c**) *MachineLearning*; one of the attributes or *slots* (**s**) of the concept is the firing threshold for neuronal simulation, expressed as the ordered pair (*minValue*, *maxValue*); the facets (**f**) express the domain of the threshold value, which is considered to range from zero to one and the usual value is between 0.4 and 0.6.

The OKBC model was extended in [1] to represent the behavior of attributes over time by describing the permitted changes in a property and also how properties are inherited by subconcepts. Also, the prototypical (usual) and exceptional properties holding for a concept can be described. Using these facets the cited example may look as follows:

```
c:MachineLearning;
s: NeuronSimulationActivationPotencial
  f: Domain: [0.00,1.00];
  f: Value: [0.40,0.60];
  f: Typeofvalue: prototypical;
  f: Exceptions: [0.00,0.39] ∪ [0.61,1.00];
  f: Ranking: 3;
  f: Changefrequency: Volatile;
  f: Event: (stimulated, [0.20,0.60], after, R);
  f: Event: (damaged, [0.90-1.00], after, I);
```

The extended example expresses that the usual threshold value is between 0.40 and 0.60; below and over that range points to an *exceptional* situation (as shown by *exceptions*). The *type of value* tells that the attribute *value* holds for any prototypical instance of the concept;

exceptions are permitted with the degree expresed by the *ranking*. The *type of value* can also be *inherited* or *distinguishing* (to differentiate among siblings) [1]. A *ranking* of 3 tells the degree of confidence on observing the usual value is for *most* cases (the scale ranges from 1, the highest, to 7). A volatile *change frequency* indicates that the threshold value can change more than once (other values can be *regular*, *once only* and *never*). *Events* are described which can change the value of a neuron's threshold along with the new values, e.g. if a neuron is simulated as being stimulated by some neuro-transmitter substance the threshold value can be as low as 0.20 *after* this event takes place; finally, the change in the threshold value is reversible(*R*), probably after finishing the effect of the substance. If a neuron is damaged, the threshold changes irreversibly (*I*).

3. THE MOISE⁺ORGANIZATION MODEL

The *MOISE⁺*(Model of Organization for multi-agent SystEms)[4] is an attempt to join roles with plans in a coherent organizational specification, seen under three aspects: structural, functional and deontic (not discussed here). While the roles build the structure of an organization, the plans build their functioning.

The *MOISE⁺*structure is built up in three levels: (i) the behaviors that an agent playing a role is responsible for (*individual*), (ii) the structure and interconnection of the roles with each other (*social*), and (iii) the aggregation of roles in large structures (*collective*).

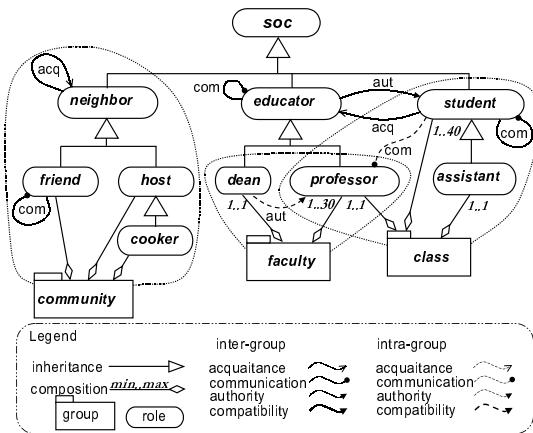


Figure 1 - A structural specification example.

In the individual level, roles are defined and structured by an inheritance relation (see fig. 1). A *role* is a set of constraints that an agent ought to follow when it accepts to enter a group playing that role. Roles are bound in the social level by links (acquaintance, communication, and authority). The links constrain an agent in relation to another agent. E.g., an agent is allowed to communicate with another agent only if there is a communication link among them.

Roles can only be played in the collective level, i.e., in a group already created from a group specification. In *MOISE⁺*, a group specification is defined by: a set of roles; a set of links; a set of role compatibilities; and the number of agents that have to play roles in this group type. E.g., in fig. 1 the group class must have one agent playing the professor role, one agent for the assistant role and, at most, 40 agents playing the student role. As in one organization there can be more than one group, a structural specification is, roughly speaking, a collection of groups and roles.

The functional specification (FS) describes how a MAS usually achieves its global goals, i.e., how these goals are decomposed (by *plans*) and distributed to agents (by *missions*). Such decompositions are called *social schemes* (SCH). A *mission* is a set of global goals to which an agent commits. The SCH of the FS may be set either by the MAS designer who specifies its expertise in a SCH or by the agents that store their past (best) solutions.

A set of agents that instantiates an organizational specification constitutes an organizational entity (OE). The dynamics of an OE (agent entrance or leaving, group creation, role adoption, SCH starting or finishing, mission commitment, etc.) is represented by a sequence of OE states. A state is described by the current organizational specification, the set of agents that created this OE, the created groups with their specification, the roles the agents are playing, the SCHs that are being performed with their specification, and the missions the agents are committed to in those SCHs.

We will assume that agents can represent and reason not only about the groups they presently belong to, but also that they can, at least to a some extent, do the same with groups they made part of in the past.

4. ONTOLOGIES AND ORGANIZATIONS

We will consider that a society is the top level entity. Accordingly, an uppermost ontology is to be attached to it. On the same basis, groups, roles and agents may have associated specific ontologies. The strings that connect them all are rendered through the organizational(e.g., temporal, authority, etc.) links. An agent's knowledge should not be considered as coinciding entirely with some role ontology as role specification can change over time. A role tells what to expect from an agent, what it should know, and not what it actually knows. This depends on the agent's lifetime experience. A role may be enriched with the previous experience of agents that played it and vice-versa [8].

The ultimate target of this work is to use MAS to support currently existent human communities. A human user usually keeps his/her organizational links as indexes to concepts, not necessarily distinguishing different ontologies, but accessing (not necessarily disjoint) portions of his/her knowledge as they are socially needed(e.g depending on

groups), as suggested by Halbwachs [9]. An agent acting on behalf of the user may take advantage of the same approach.

We tried to harness the knowledge context problem applying the MAS metaphor. A valuable information that can be recorded, given an organization model like *MOISE⁺*, is related to the group were the knowledge was used and more specifically, the SCH, role and plan that needed it. The instance specific representation is the subjective view of an agent recalling its participation on a dated instance of a SCH. To recall a particular portion of a knowledge representation an agent can use temporal links developed through her/his memory about the groups (i.e, knowledgeable agents, acquaintance links, etc.) s/he made part of. To do so, a suitable mechanism should be added to existing ontology access models.

The present work is thought as complementar to that of Tamma and Bench-Capon's, notably addressing a distinct aspect of ontology manipulation which is the notion of *situating* social/individual concepts in an organizational context. As stated by Halbwachs [9] memory has an eminently social character. Thus, the three following *facets* are defined to hold the organizational context in which concepts are used:

a)Knowledge Usage Organizational Situation (KUOS): allows to evoke the organizational state where the attribute (*slot*) was used by an agent. There can be two kinds of situations: (i) an agent uses the attribute, and (ii) an agent observes another agent using the attribute and wants to remember how it was used. The latter can be regarded as a special case of learning when an agent records a fact but doesn't adopt a new concept. As there can be multiple entries of this type, it is defined as a set { (ag_k, gi_k, ss_i_k, r_k, t_k, d_k, a_k, i_k) | k = 1, ..., n } (see fig. 2).

b)Knowledge Acquisition Organizational Situation (KAOS): allows to evoke the organizational state where the *value* was learned. As there can be multiple entries of this type (a concept can be modified many times) it is defined as a set { (lv_k, tid_k, gi_k, ss_i_k, r_k, t_k, d_k, ok, i_k) | k = 1, ..., m } (see fig. 2). This field is intended for the case when an agent has acquired a portion of its ontology through word of mouth or some other more normative way. It can keep the identity of the organizational construct that taught it. Eventually, it can be used in a group to determine who was the individual whose knowledge was socially adopted, i.e., this *frame* should record when was something new learned and from whom.

c)Accessibility: allows reasoning about concepts that are to be kept private by an agent, distinguishing them from those that can be shared (say publicly in a group). It will indicate whether and which is the construct (e.g agent, role, group, etc.) to which it is allowed to publicize the attribute. There can be multiple entries of this type, since a concept can have different permission values for roles inside the same group. It is defined as a set { gt_k, gi_k, r_k, ag_k, spv_k, | k = 1, ..., p }

(see fig. 2). This *facet* reflects, in the case of agents, their beliefs about their knowledge as socially situated. The need to differentiate among what is to be considered “public” or “private” is specially critical when considering collaborative agents in open MAS, as they should be able to count on each other, considering the goals of other agents and having attitudes towards those goals [10] [3].

The parameters of the defined facets are not detailed for the sake of space, but fig. 2 should help in clarifying them.

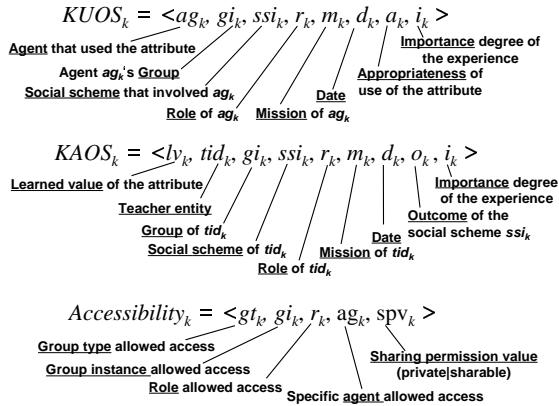


Figure 2 - The KAOS, KUOS and Accessibility facets.

Consider the ontology of an agent *John* (i.e., in *his* mind), and also assume that *now* corresponds to December 1st, 2010. Using just the newly defined *facets* the example showed in section 2 may eventually come to look as follows:

```
c:MachineLearning;
s: NeuronSimulationActivationPotencial
f: Domain:[0,0,1,0];
f: Value: [0,4,0,6];
f: KUOS:[Peter,G4,71,student,m5,3May2010,0,6,1,0];
f: KUOS:[Myself,G4,70,professor,m3,4Jun2010,0,9,0,6];
f: KUOS:[Myself,G5,81,professor,m5,1Dez2010,0,9,0,8];
f: KAOS:[[0,5,1,0],Al,G73,v1,professor,m2,2000,0,8,0,3];
f: KAOS:[[0,4,0,6],Don,G7,w,colleague,m6,2003,0,9,1,0];
f: Access:[faculty,-,-,-,shareable];
f: Access:[-,G5,assistant,Peter,shareable];
```

Now, the example expresses that the *slot* was used 3 times (*KUOS*), that there were two acquisition situations (*KAOS*) and that are two types of organizational constructs that have permission to access it(*Access*).

The first *KUOS* line tells that *John* observed that *Peter* from a group instance *G4* (say an instance of the *class* group type) was executing some *71 scheme instance* as a *student*. *Peter* used the concept with some mission *m5* in *May 3, 2010* and *John* judged that the degree of appropriateness was *0.6* which was highly regarded by him, as the *importance* degree recorded *1.0*. We may suppose that *John* assigned this *importance* value because,

although *Peter* was not really correct (according to *John*), he observed that *Peter* developed an authority (link) over all his colleagues (playing the *student* role) after the use of the concept. Note that by the second *KUOS* line *John* only “taught” the value almost a month after *Peter* used it. Another approach to obtain the *importance* value can take into account the combination of *John*'s indices with those of *Peter*'s as in belief networks.

The first *KAOS* line tells that *[0,5,1,0]* was the range learned for the *value* and that the teacher was *Al*, while in a group instance *G73*, executing a *v1* scheme instance; also, *Al* was playing as a *professor* and executing some *m2* mission. Still, that situation took place “ten years ago” in the year 2000 and that the outcome of the *v1* scheme instance was considered highly satisfactory for *John* (see the value of *0.8*) although *John* didn't find it really important (only an importance of *0.3*). As this was the first value of *KAOS* considering the *date*, this was the first *value* for this attribute that was learned by *John*, and he kept believing on it for three years until *Don* “taught” him well (a highly important fact as the importance stated by *John* for the second *KAOS* importance was *1.0*).

The first *Access* line tells that *any* agent playing *any* role belonging to *any* group instance of the *faculty* group type will have access to the attribute *value*. The same stands, in a more restrictive sense, for *Peter*, given that he is playing as *assistant* in the group instance *G5*(an instance of the *class* group type) as stated by the second *Access* line.

5. REASONING

Though necessary, the presented mechanism is useless without a reasoning scheme which will depend on the agent's mind. An ontology to which this type of information was added is organizationally *situated*. Actually, the *facets* allow processing information in a level of detail that may go beyond the granularity used in a *MOISE⁺* organizational specification. This is done to handle situations that may arise at the agent, i.e. instance, level.

An example of such a scheme is in fig. 3. Two mental constructions of an agent *John* are depicted. The table represents part of the information concerning *his* representations of the organizational stances he participated in (with his impressions about the agents *Joe*, *Peter* and *Mary*). *O_{John}* is the whole ontology of *John*, *his* representation of his own knowledge.

Using the new *facets* that were presented, *John* now has multiple ways of accessing his, in another way rather “monolithic”, ontology *O_{John}*. If somebody asks him about Spanish, and if he thinks his knowledge of Spanish (denoted as *O_{Spanish}* and making part of the bigger *O_{John}*) is not enough for a given situation he may now fetch among his acquaintances those that showed valuable (say accessing the

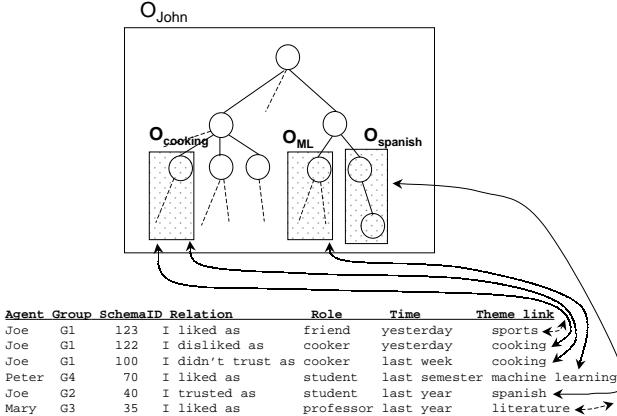


Figure 3 - An agent *John* reasoning about other agents in an organizational context and relating them to portions of his own ontology.

i_k 's, o_k 's or a_k 's) when using (*KUOS*) and/or acquiring (*KAOS*) knowledge of spanish. The “relation” field may be substituted by any of the cited indices (i_k , o_k or a_k).

6. COMMUNET: A PEER-TO-PEER MAS FOR INFORMATION EXCHANGE

To obtain information, people refer to their contact networks[11]. Peer-to-peer systems can be used as computational infrastructures to make available different services in those networks, e.g. expertise location and information exchange. The adoption of the proposed approach, by autonomous agents in an open and distributed MAS, can lead to broader forms of collaboration among members of Internet communities.

As a first step to put the proposed approach into practice, it was applied in COMMUNET, a peer-to-peer MAS dedicated to information exchange in a community environment. It departs from the assumption that each user has a social network that can be used as a starting point to assign social groups and roles to contacts. This is a much weaker assumption regarding usual organization approaches which depend mostly on teams tied by processes.

Each agent is dedicated to one human user. Initially, the agent processes the user's documents and tries to map which are the matching concepts in a general ontology using a selected vocabulary for each concept. The ontology made available by the Open Directory Project (ODP) (www.dmoz.org) was selected as a top level ontology. Such an ontology is initially needed as Steels [12] showed that agents converge to a common ontology only through interaction. For each concept a small set of pages is previously chosen (as in [13])) as examples for it, to extract a representative vocabulary using summarization techniques [14]. Each agent starts with the ODP ontology along with the

extracted vocabulary for each concept. Then, it builds up the user personal network based on its e-mail messages and tries to map which are the concepts that appear on them (and also in his/her documents), spotting the contacts that appear related to it.

Using the presented facets with the personal network data, the agent can have a better view of the relationships between the user's knowledge and its social contacts (as shown in the didactic example in fig.4).

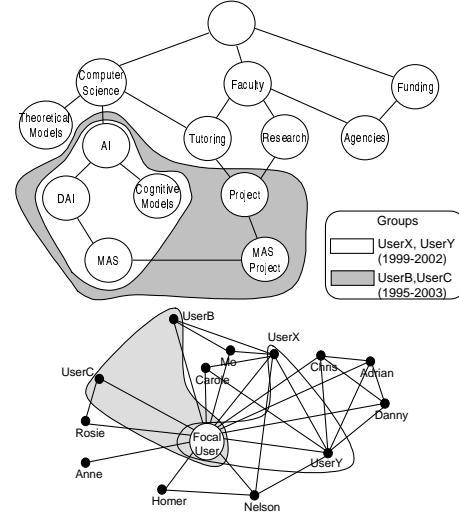


Figure 4 - Interests of two groups, referencing different time periods, as subsets of the *situated ontology* of a focal user and also as contacts in his/her social network.

To finish the setup, the agent interacts with the user to confirm the existent organizational constructs and get the groups/roles permissions. A real example of a role view is shown in fig.5.

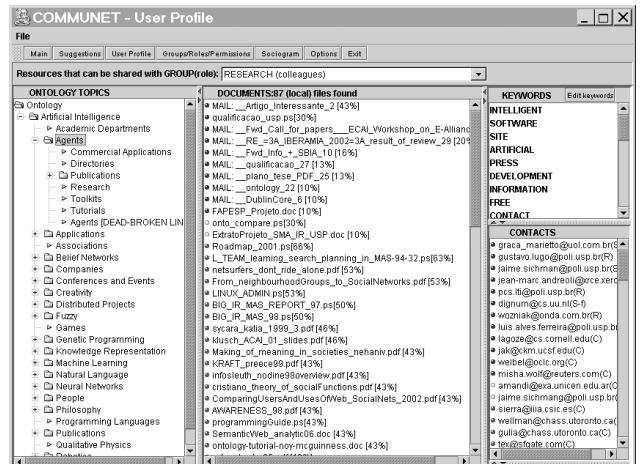


Figure 5 - Documents and contacts of a COMMUNET user, linked to the concept *Agents*, viewed through the filter of the role *colleagues* from the group *research*.

The agent has an interface that allows the user to query

a general purpose search engine. At the same time, the user query is matched against the concepts of the ontology, recommending files and contacts retrieved from other COMMUNET agents in a peer-to-peer basis (in a Gnutella like fashion). As each agent is autonomous, the user knowledge privacy is kept intact and is by no means relaxed besides the limits set by the permissions stored in the accessibility facet.

A further implementation step will allow agents whose users have changed their ontology to publicize and negotiate the changed concepts with other agents that belong in the groups that share those concepts.

7. CONCLUSIONS

This paper presented a model that helps to bridge different world views when agents interact to use/share knowledge in MAS, referring it through organizational information. The knowledge representation model extension is provided as a means of facilitating integration of ontologies based on organizational experiences. Even though there is some related work [15] [16], none of the known approaches tackles the organizational dimension of concepts as part of autonomous agents belonging in an open MAS. Also, agent reputation [17] [18] [19] can be better exploited. The extension overhead pays off bringing a facility to reason about explicit representations of knowledge and organization altogether in a way that it was not previously available.

8. ACKNOWLEDGEMENTS

The first author, supported by CNPq, Brazil, grant 140372/00-8, would like to thank Andrea Wozniak G. for the enlightening discussions on the social meanings of time, organization and knowledge. The second author is partially financed by CNPq, Brazil, grant 301041/95-4; and by CNPq/NSF PROTEM-CC MAPPEL project, grant number 680033/99-8. The third author is financed by FURB and CAPES, Brazil.

REFERENCES

- [1] V. Tamma and T. Bench-Capon, “A conceptual model to facilitate knowledge sharing in multi-agent systems,” in *Ontologies in Agent Systems AS 2001 Workshop Proc.* Canada, May 2001, pp. 69–76.
- [2] V. Chaudhri, A. Farquhar, R. Fikes, P. Karp, and J. Rice, “OKBC: A programmatic foundation for knowledge base interoperability,” in *AAAI-98*, 1998.
- [3] J. Sichman, R. Conte, Y. Demazeau, and C. Castelfranchi, “A social reasoning mechanism based on dependence networks,” in *Proc. of the 11th ECAI*. T. Cohn(ed.), August 1994, pp. 188–192.
- [4] J. F. Hubner, J. S. Sichman, and O. Boissier, “A Model for the Structural, Functional and Deontic Specification of Organizations in Multi-Agent Systems,” in *SBIA’02 Conference. LNCS 2507, Springer*, 2002.
- [5] N. Guarino, “Formal ontological distinctions for information organization, extraction, and integration,” in *Information Extraction: A Multidisciplinary Approach to an Emerging Inf. Technology. LNAI 1299*. M. T. Pazienza(ed). Springer, 1997.
- [6] T. Gruber, “Towards principles for the design of ontologies used for knowledge sharing,” *Intl. Journal of Human and Computer Studies*, vol. 43, no. 5/6, 1995.
- [7] H. Wache, T. Vögele, U. Visser, H. Stuckenschmidt, G. Schuster, H. Neumann, and S. Hübner, “Ontology-based integration of information - a survey of existing approaches,” in *IJCAI Workshop on Ontologies and Information Sharing*, 2001.
- [8] J. Sichman and R. Conte, “On personal and role mental attitudes: A preliminary dependence-based analysis,” in *Proc. of the SBIA’98 Conference*, 1998.
- [9] M.-C. Lavabre, “For a sociology of collective memory,” in *CNRS Cross-disciplinary Encounters: Memories*. <http://www.cnrs.fr/cw/en/pres/compress/memoire/lavabre.htm>, 2000.
- [10] C. Castelfranchi, “Modeling social action for AI agents,” *Artificial Intelligence*, no. 103, pp. 157–182, 1998.
- [11] J. Travers and S. Milgram, “An Experimental Study of the Small World Problem,” *Sociometry*, vol. 32, no. 4, pp. 425–443, 1969.
- [12] L. Steels, “The origin of ontologies and communication conventions in multi-agent systems,” *Autonomous Agents and Multi-Agent Systems*, vol. 1, no. 2, pp. 169–194, October 1998.
- [13] X. Zhu, S. Gauch, N. Kral, and A. Pretschner, “Ontology-Based Web Site Mapping for Information Exploration,” in *CIKM’99 Conference*, 1999.
- [14] D. Lawrie, W. Croft, and A. Rosenberg, “Finding Topic Words for Hierarchical Summarization (draft version),” in *SIGIR’01*. Sept., 2001.
- [15] M. S. Lacher and G. Groh, “Facilitating the exchange of explicit knowledge through ontology mappings,” in *14th FLAIRS Conference*. AAAI Press, May 2001.
- [16] D. G. Schwartz and D. Te’eni, “Tying Knowledge to Action with kMail,” *IEEE Intelligent Systems*, vol. 15, no. 3, pp. 33–39, 2000.
- [17] J. Sabater and C. Sierra, “Social ReGreT, a reputation model based on social relations,” *ACM SIGecom Exchanges*, vol. 3.1, pp. 44–56, 2002.
- [18] B. Yu and M. P. Singh, “An Evidential Model of Distributed Reputation Management,” in *Proceedings of the AAMAS’02 Conference*, 2002.
- [19] J. Pujol, R. Sanguesa, and J. Delgado, “Extracting Reputation in Multi Agent Systems by Means of Social Network Topology,” in *AAMAS 2002*, July 2002.