

Towards Modeling the Social Layer of Emergent Knowledge Using Open Ontologies

Tina Froehner¹ and Matthias Nickles² and Gerhard Weiss²

Abstract. In the context of *open environments* like Open Multiagent and Peer2Peer Systems, the Semantic Web and Communities of Practice, knowledge-based applications with self-interested, *autonomous* knowledge sources like information agents and humans have gained increased interest in recent times. For such environments, neither the consistency nor the reliability or the definiteness of the input knowledge can be ensured. Whereas there already exist promising approaches to the filtering, mapping and homogenization of inconsistent and dynamic knowledge, the modeling of knowledge heterogeneity and dynamics *itself* has been largely neglected so far. In this work, we propose Open Ontologies as a “social” approach to this issue. Open Ontologies are based on a distinct communication-oriented paradigm as they emerge from and evolve with communication processes, and allow for the modeling and processing of semantically heterogeneous, possibly conflicting conceptual knowledge by means of its reification according to its social (i.e. communicational) meaning, impact and relevance, and the probabilistic weighting of inconsistent knowledge facets.

Keywords: *Formal Ontologies, Emergent Semantics, Knowledge Management, Information Integration, Agent Communication, Semantic Web, Peer2Peer Systems*

1. INTRODUCTION

From the viewpoint of knowledge management, open multiagent systems and other open environments like the Semantic Web, open Peer2Peer systems and Communities of Practice can be described as networks of autonomous knowledge sources. The benefits of agent-based knowledge management systems in comparison to traditional systems are their high degree of robustness and flexibility, which enables the system to quickly adapt to changes of its environment. On the other hand, agent autonomy can lead to a high level of complexity and unpredictable system behavior. The specific kind of symbolic, autonomy-

preserving interaction of truly autonomous black- or gray-box agents (i.e., agents with more or less unknown internal states, cognition and goals) is *communication*, which brings along the possibility of contradiction, ambiguity, insincerity and indefiniteness of meaning. Important media to enable communication are *ontologies* in order to facilitate a common understanding by means of a shared set of concept definitions.

Traditional approaches to ontology modeling and acquisition have certain shortcomings in regard to open environments as i) they seldom handle changes of the meaning of concepts, ii) they seldom consider conceptual knowledge as being contextualized with intentions, processes and effects from/on the world “outside” their respective domains (e.g. social goals), iii) they often pursue a “one ontology fits all” policy without an adequate consideration of user needs variety, and iv) they have no concept for the treatment of semantic heterogeneity which is sufficient in our opinion. Whereas approaches like *emergent semantics* [1], *dynamic ontologies* [2] and semantical ontology *alignment* and *clustering* in order to merge heterogeneous knowledge [e.g. 3] have caused very significant improvements regarding some of these issues, semantical inconsistencies are almost always still taken for something which either should be avoided, or should be filtered out (e.g., using criteria like *(dis-)trust* and *reputation* [e.g. 8], or *collaborative rating/majority voting* [e.g. 17,23,11]), or that it would be sufficient in this respect to annotate inconsistent views with adequate local truth *contexts* [4] to make them manageable. Although all of these approaches have their particular benefits within their areas of applicability, it should be recognized, that semantical inconsistencies are not just unfavorable or local states, but that they are in real-world environments often unpreventable due to *stable* social belief or goal conflicts [5] of knowledge sources, that they can provide the knowledge users with valuable meta-information about the intentions, goals and social relations among the knowledge sources, and - provided that they have been made explicit, visible and weighted - that they can be prerequisites for subsequent argumentation and conflict resolution and thus knowledge evolution. Ultimately, mechanisms for knowledge integration and sharing e.g. within the field of knowledge management, can only found preliminary decisions about the reasonable modeling of communicated knowledge artifacts, because within a heterogeneous group of knowledge sources and users, in the end each user can only decide for herself about the relevance and correctness of the given information, which

¹ Research Center Knowledge Management (RCKM), University of Applied Sciences Cologne, D-50678 Köln, Germany, email: tina.froehner@fh-koeln.de

² AI/Cognition Group, Chair for Theoretical Computer Science and Foundations of Artificial Intelligence, Technical University Munich, D-85748 Garching bei München, Germany, email: {nickles,weissg}@tiki.informatik.tu-muenchen.de

provides a strong argument for the leveled conservation of knowledge heterogeneity *while integrating into a shared ontology or knowledge base*. Currently, this social context of knowledge is maintained mostly implicitly and disintegrated in distributed systems of local entities like agents. Only recently, social structures of multiagent systems have been related to knowledge management issues [22,20,11].

Since the Semantic Web can be seen as typical for open environments, most of the issues arising from knowledge source autonomy can be found there, and due to the extend and the dynamics of the Semantic Web they are expected to appear in their most severe shaping. Similar problems arise in open multiagent systems, or when creating and sharing knowledge within e.g. large organizational enterprise knowledge bases, *Communities of Practice*, or between organizations (e.g. *virtual organizations*, outsourcing and offshore), or in certain Peer2Peer systems. Although the Semantic Web effort addresses the problem of missing machine-understandability of web site descriptions, it currently focuses primarily on the specification of languages and tools for the representation of consistent semantics and ontologies, not on the inherently social processes of information gathering and descriptions, and it is just beginning to take into consideration phenomena like the social impact of resource descriptions [6], the computational handling of conflicting opinions, information biasing by competing commercial interests, and inconsistent or intentionally incorrect information. So, the fact that probably the most important aspect of such open environments is the explicit or implicit communication among a possibly very large number of autonomous actors is largely neglected so far, at least from a formal, computation-oriented point of view. Bringing information (e.g. via websites, services and their annotations with meta-data) into a public space like the web is in fact a social act, and the relationship between informational artifacts on the web is communicational (i.e. *specifying, agreeing, contradicting* etc.). This can produce intentional and unavoidable inconsistencies of knowledge concepts (e.g. company interests versus customer interests or various conceptualizations due to differences in human culture). If these are filtered out or homogenized to early, important decision-relevant information for the users or the applications might be lost.

With this paper, we propose a novel approach to the emergence and modeling of ontologies aiming at a combination of recent advances in distributed artificial intelligence regarding knowledge-based agent communication and the research on knowledge management for open environments. We propose *Open Ontologies* that are derived from the communication of multiple knowledge sources and users, and maintain semantical heterogeneity and social (i.e. communication) structures. In demarcation to e.g. *Communication Systems* [18], *discourse/interaction ontologies* [e.g. 24] and *Social Networks* [21], Open Ontologies do not aim for the exclusive modeling of social structures, but for the *integration* of heterogeneous knowledge together with a model of its social meaning, focusing on the *link* between the knowledge level and its social level. The formal representations usable for each of

both levels is to some degree variable, allowing for the transformation of “conventional” ontologies into Open Ontologies, and the other way round the usage of homogenized or highly generalized, definite Open Ontologies as conventional ontologies.

The reminder of this paper is structured as follows: The next section outlines the basic properties of Open Ontologies and section 3 informs about the principles of representation and acquisition. Section 4 applies our concept exemplary to the Semantic Web, and section 5 closes with an outlook on open research issues.

2. OPEN ONTOLOGIES

Formal ontologies are traditionally defined as agreed formal conceptualizations of certain domains that serve as common ground for tasks like knowledge exchange and modeling. As we have seen, this understanding leads to difficulties if the informational input the ontology is build from is likely to be intentionally inconsistent, and there either does not yet exist enough meta-knowledge like trust to identify and filter out “inappropriate” or “wrong” data *a priory*, or there does not even exist a concept of global inappropriateness or correctness at all. On the other hand, consistent and agreed ontologies are doubtless an inevitable prerequisite for efficient knowledge creation, representation and exchange, whereby we consider implicate and emerged ontologies and schemata (e.g. in the context of semi-structured data modeling) to be such ontologies too. Open Ontologies aim at the solution for this dilemma by embedding conceptual knowledge facets gained from a heterogeneous set of self-interested autonomous knowledge sources (e.g. information agents or humans) within contextual information about their communicational (i.e. social) origin, impact, and relationships (e.g., contradiction, approval, revision or specification) to other communicated facets (which can be communicated by means of formal communication languages, but also be provided as, e.g., structured, semi-structured or natural language documents) and their sources. Doing so, in Open Ontologies information as it can be found as first-order concept descriptions in conventional Ontologies, is lifted to the *social level* and thus to a level where the sources and the users of the ontology are likely to achieve an agreement with the social assessments of possibly inconsistent and uncertain facts. The agreement or disagreement with certain assessed facts, based on the meta-information the Open Ontology provides is then a subsequent task based on rich social knowledge instead of binary decisions like to trust or not to trust particular agents. Open Ontologies are thus *formal ontologies which receive their content from the communication of multiple autonomous sources and users, and provide a dynamic representation of socially annotated heterogeneous knowledge*.

In the philosophical tradition of *Pragmatism* [7], the works of Wittgenstein on language meaning [9] and *Social Systems Theory* [13], communication is thereby not so much to be understood as the exchange of symbols with fixed meaning, but the other way round as a means to generate supra-individual meaning from interrelated interactions among

black- or gray-box agents (i.e., agents with more or less unknown internal states, cognition and goals). The practical consequences arising from this are that Open Ontologies are necessarily *dynamic* ontologies that need to be continuously adapted in reaction to new information, and that the processes of communicational creation, understanding and usage of ontological information are *integral aspects* of Open Ontologies i.e., Open Ontologies are not just evolving from and influencing communication processes, but these processes (precisely: suitable representations of certain process structures that proved themselves to be more or less stable over time, cf. section 3) are part of the ontologies themselves, providing contexts of communicated knowledge facets (such contexts must not be confused with contexts in the sense of truth conditions, as they represent the situation of knowledge emergence and usage, not situations in which the respective statements become true). In addition, communication among multiple agents likely requires mechanisms for the *generalization* of emergent meaning, since otherwise the complexity of an Open Ontology would grow too large due to the sheer number of individual

knowledge contributions. Generalization is also a way to make Open Ontologies look like homogeneous ontologies if necessary, because at its highest level, generalization causes semantical homogenization among contradicting knowledge sources. Summing it up, Open Ontologies have the following main characteristics:

- **Openness:** No assumptions are made regarding the benevolence, trustworthiness, relevance, informedness and cooperativeness of its knowledge sources.
- **Dynamical derivation from communication:** Open Ontologies are emergent from and evolving with ongoing communication of its sources (e.g. information agents) that asserts (denies, specifies...) information and expresses and specifies informational needs.
- **Explicitness and social annotation of semantical heterogeneity:** Open Ontologies maintain semantical inconsistencies arising from contradictions and conflicts, and contain annotations of conceptual




Figure 1. Emergence of Open Ontologies

knowledge with higher-level information about its social meaning within the courses of communication (social knowledge). Nevertheless, although Open Ontologies allow logical inconsistencies of *1st-level knowledge* (cf. below), they are consistent on the top level of social knowledge.

For the social annotation, various kinds of social structure formalizations can be used, whereby we recommend an empirical approach as described in section 3.

- **Multiple levels of social generalization:** They allow multiple levels of generalization of social concepts, like agent roles and social groups, weighting the degree of inconsistency and meta-information. For this purpose, stochastic methods might need to be applied (e.g. on the topmost level of generalization in order to calculate consistent group opinions from divergent single opinions).

In the following, we will sketch the framework for the generation and representation of Open Ontologies, for lack of space without providing technical or formal details (cf. [16]). An application of some of the concepts of Open Ontologies to the multiagent-based rating of RDF statements can be found in [11,10]. The emergence of semantics from the empirical observation of agent communication is described in detail in [15]. Typically, Open Ontologies have a data \leftarrow meta-data \leftarrow meta-data \leftarrow ... structure, containing as most simple first-order objects knowledge facets having the form '*1st-level knowledge* \leftarrow *2nd-level knowledge*', with *1st-level knowledge* being probably logically inconsistent with other *1st-level knowledge* within the same Open Ontology. Since Open Ontologies are primarily an abstract meta-concept build upon conventional approaches for the representation of knowledge, we do not restrict the sort of concrete entities that are to be "wrapped" within an Open Ontology or at the content level of agent messages, like first-order statements, classes or frames. For the same reason, we do also not make any assumptions relating to ontology domains or concrete areas of application here. In contrast to *1st-level knowledge*, *2nd-level knowledge* (also called *social knowledge*) depicts the social context of *1st-level knowledge*, the latter taken as generated from a communication act of an autonomous source of knowledge. We call this kind of annotation of knowledge with higher-level knowledge about its social meaning *Social Reification* [16]. The most trivial kind of Social Reification is *quoting* (e.g., "Sue says: "..."), but in general, all kind of information which describes how and to what effect certain data is produced within a process of communication can be understood as *2nd-level knowledge*. Of course, we can apply Social Reification recursively i.e., annotate *2nd-level knowledge* with *3rd-level social knowledge* as in "Sue says: 'Tom says: "...'" and so on.

Very simple forms of such social meta-data are considered to be agent speech act types like *assertion*, *denial* or *query*, inducing relations among single communication like "Sue contradicts Tom's statement saying "...'" and rich *2nd-level knowledge* types such as emergent knowledge source and user profiles and even complex social systems like

organizations [13,19]. In our empirical communication model [15,18] these symbolic acts gain their meaning from their expected effect on the subsequent trajectory of communications, which can be learned empirically from past interactions (although we recommend this so-called *empirical-rational communication semantics* to disregard mentalistic details which are unknown for autonomous agents and allow for the handling of uncertain meanings, the usage of such a semantics is not required to define an Open Ontology). Because meaning is contextualized by the meaning (expected effect) of the situation (history) of the respective act occurrence, in general *2nd-level knowledge* describes *communication processes* (this applies even to simple quotations: In Sue says: "...", "Sue" is in fact just an abbreviation for the pragmatic impact utterances from Sue are expected to have).

As mentioned earlier, Open Ontologies require the generalization of meaning in order to reduce complexity. Generalization in this sense has two dimensions: i) the merging of *2nd-level knowledge* (and, if required, *3rd-level knowledge* and so on), and ii) the subsequent merging of related lower-level knowledge facets. Typically, i) comprises the merging of similar social processes to *interactions patterns*, and the combination of multiple similar behaving agents to social *groups* or *roles*. After applying such generalization rules to *2nd-level knowledge*, the annotated *1st-level knowledge* needs to be merged accordingly. If, for example, multiple agents forming a single social group make inconsistent assertions, within the Open Ontology each of these assertions obtains a probabilistic weight expressing the degree of expected approval this assertions gets from the role or group as a whole (calculated e.g. from the frequency this assertion has been uttered by different agents within this role or group as proposed in [12,11,10]).

We propose the usefulness of a co-presence of multiple levels of generalization, tailored to the desired levels of heterogeneity of the Open Ontology. Open Ontologies are continuously learned by an observer of knowledge source and knowledge user communications. The technical requirements for this learning process are i) *information agents* able to communicate *1st-level knowledge* facets like "c₁ is-a c₂" or "c₁ has-a c₂" at the content level of their communication language (since Open Ontologies do not require agent cooperativeness, performatives used for (e.g.) negotiation are not required, although they would make up a useful extension), ii) a facility for the acquisition of Open Ontologies from the observation of agent communication, e.g., a dedicated middle agent within in the infrastructures of the respective application, called a *semantic observer*, iii) a facility for the low-level storage and querying of persistent ontological data (e.g., a database management system), and optionally iv) a facility for the social reasoning upon the *2nd-order knowledge* within the Open Ontology (to deduce new facts like "Sue is likely to contradict or specify Toms information", but also to derive trust relationships among the participants).

For a two-level structure, the acquisition of Open Ontologies comprises the following main tasks, which are to be

performed in a loop as a continuous, evolutionary learning process for the whole period of agent communication:

1. Observation and filtering of agent communication (e.g., along restrictions to a certain concept domain)
2. Derivation and/or adaptation of 2nd-level knowledge according to the respective semantical model (e.g. empirically)
3. Stochastic generalization of 2nd-level knowledge (multiple levels, if necessary)
4. Extraction and generalization of 1st-level knowledge from communication (e.g., from the content level of speech acts)
5. Social Reification and generalization of 1st-level knowledge (multiple levels, if necessary - depending on application)
6. Constraining of the results with given, obligatory 1st-level knowledge (e.g. some normative top-level ontology) and normative 2nd-level knowledge (normative social structures like rules or *laws* forbidding or demanding the utterance of certain information), if necessary.

3. REPRESENTATION AND AQUISITION

In the following, we will provide a description of the representation and empirical acquisition of Open Ontologies from communication (for details, please see [16]). As already said Open Ontologies emerge from communication processes and consist of socially annotated heterogeneous knowledge facets. In fact, both 2nd-level and 1st-level knowledge like

(homogeneous) ontological concept descriptions can be seen as special cases of *social structures*, since commonly agreed ontologies as well as social knowledge have a structuring effect on communication [13]. Therefore, from the viewpoint of the Semantics Observer, the main task in order to derive Open Ontologies from communication is (i) to compute a rich model of sociality and subsequently (ii) to extract and combine 1st-, 2nd-,... level knowledge.

Action expectations regarding the expected continuation of communication processes can be seen as the most general model of social structures and the most elementary semantics of communications [13]. As we have shown in [15], these structures can be computed and represented as stochastic processes, called *Expectation Networks* [18], which we want to use as an application-independent starting point for Open Ontologies. Nevertheless, in principle other representations of agent sociality like organizational structures [e.g. 19] or rules are likely usable also. Expectation Networks are building upon the most basic social ontology consisting of the terms “expectation”, “action” and “communication process”, as inspired from Social Systems Theory [13]. Since Expectation Networks should be seen as an universal, elementary modeling approach to social structures, they can also be used as a kind of “assembler language” for traditional ontological representations of sociality (e.g., in terms of Expectation Networks, agent “personalities” and roles are certain expected interaction sequence patterns that are bound to variables having agents as their instances).

An Expectation Network is a graphical data structure (not necessarily coherent), which consists of correlated stochastic expectations regarding the future communicational behavior of a set of agents or agent roles (obtained from a temporal generalization of observed communication trajectories). Its basic constituents are event nodes that represent expected utterances and other agent-generated events, and probabilistically weighted edges connecting event nodes, which represent significant correlations of the respective utterances. Expectation Networks always represent dynamic




Figure 2. Simple Expectation Network

belief (i.e., belief that is constantly under revision according to the evolution of social structures as they show up). Figure 2 shows an example of an Expectation Network that represents the structure of a discourse of two information agents (for simplicity, this graph slightly deviates from the formalism presented in [18]). Nodes (squares) are labeled with message templates (in a formal speech act-based agent communication language) and the special symbols “⊥” (denoting the end of a conversation) and “?” (denoting a yet unknown continuation). Nodes are connected by edges (arrows) labeled with numerical *expectabilities*, which denote the probability that the respective message(s) will occur subsequently. These probabilities are computed from observed frequencies of the respective message sequences in the past and assumptions about rational agent attitudes induced from agent behavior. The thickness of an edge represents the *normativity* of the respective expectability (e.g. denoting the conformity of the expected action with a certain law or public commitment), and the numerical value in square brackets denotes its *deviancy* (deviation of empirical expectability from social norms). An edge with high normativeness (thick arrow) also likely represents an expectation that has proved itself as empirically stable in the long term, which is a typical property of expectations obtained from laws and other social norms. In the most simple method for the construction of Expectation Networks from agent observations, the numerical expectabilities correspond to the frequency of observed message trajectories

that unify syntactically with the respective paths. A more elaborated algorithm that considers rational agent attitudes also is presented in [15].

Substitution lists appear in rounded boxes. A substitution list denotes a social role the listed agents can impersonate. For this purpose, the message templates contain role variables ($Role_A$ and $Role_B$) that can be bounded to each of the list entries. Being *generalized*, a single Expectation Network might describe the expectations regarding multiple message sequences due to different instantiations of role variables (the Semantics Observer obtains these roles automatically from stochastic analysis) and variables for external objects.

The expectation structures of a multiagent system contain as subsets the behavior-expectations that can be addressed to agents and agent roles and provide stochastic models of such actors. Therefore, Expectation Networks contain implicitly emergent agent and agent role profiles.

Theoretically, an Expectation Network contains paths for every possible sequence of messages, but in practice, edges with a very low (or unknown) probability are omitted. The message templates attached to the nodes do only serve as examples, formalized in a pseudo speech-act notation (“Sender→Addressee:Performative[Content]”). Generally, any speech-act based ACL, which supports the assertion of logical propositions as with FIPA ACL or KQML/KIF can be used for this purpose.

As one possible, simple approach to the extraction and




Figure 3. Derivation of Open Ontologies using an empirical approach

integration of distributed knowledge from Expectation Networks, one can exploit the propositional attitudes of utterances. The idea is to interpret the messages contents within their pragmatic context (network paths) as domain descriptions and to weight these descriptions according to the amount of consent/dissent (using predefined performatives like Assert and Deny). These weighted propositions are collected as knowledge facets within a knowledge repository by the Semantics Observer (see Figure 3, upper right corner). At each of these steps, generalizations can be performed in order to reduce the complexity of the Expectation Network and/or the knowledge repository by means of, for example, the combination of multiple, (expectably) similar behaving agents to single agent roles. Although it might seem reasonable to solely operate on this generated knowledge repository, the weighted 1st-order knowledge facets within an Open Ontologies should maintain the links to the contextualizing parts of the Expectation Network. This offers the possibility to fully expose the intentional context of knowledge facets and the effect of their utterance in terms of subsequent agents behavior. Now, the 2nd-level knowledge is present in form of the communication structures modeled by the Expectation Network, and the 1st-level knowledge is contained within the knowledge repository retrieved from the Expectation Network also. The annotation (Social Reification) of 1st-level knowledge with 2nd-level knowledge is then achieved by means of linkages between Expectation Network paths and the respective knowledge facets within the knowledge repository.

These facets are then interpreted as the contextualized information the respective communication sequences (paths within the Expectation Network) produce. Furthermore, the annotating 2nd-level knowledge represents the consequences (in terms of subsequent agent actions) these processes are expected to have. Typically, each process includes messages from multiple agents, and thus might produce semantically heterogeneous knowledge facets. The 2nd-level knowledge also determines the weights the 1st-level knowledge have within the knowledge base facts (in terms of expected consent/dissent among, e.g., the agent roles the network models), as these weights correspond with the *expectabilities* of utterances, which themselves are (in the most simple way - see above) calculated from the frequencies of the occurrences of these utterances in the past. In figure 3, homogeneous sets of 1st-level knowledge facets produced by certain communication processes are depicted as rectangular areas that might overlap with other such areas (upper right). The number of overlappings (depicted as grayscales) roughly corresponds with the weight the knowledge facets have within the area. In order to reduce the complexity of the Open Ontology, the Semantics Observer further needs to restrict his focus within the knowledge base depending on (different) user interest concerning the ontology (depicted as concentric circles in figure 3). The larger these foci are the more likely they contain inconsistent knowledge facets. Finally, hierarchically organized ontology clusters can be obtained (the pyramids in figure 3), containing dissented knowledge on the pyramid bottoms (weak generalization) and consented or stochastically leveled knowledge on the tops, organized in clusters corresponding to levels of

generalization. In addition, meta-clusters corresponding to demarcated ontology domains (the “topics” of the agent communication) can be obtained, leading to multiple pyramids as in figure 3.

4. OPEN ONTOLOGIES FOR THE SEMANTIC WEB

Figure 4 shows the semantical levels proposed by Tim Berners-Lee for the structure of the forthcoming Semantic Web [14].

We recommend for some aspects of this concept an extension or specification to provide solutions for the issues mentioned earlier in this work (figure 5).

In particular, it appears to us to provide formalisms and calculi that explicitly consider semantically heterogeneous meta-data like resource descriptions and ontologies created from the contributions of multiple sources that compete for the assertion of their individual “truths” and interests. Of course, the Semantic Web would already be implicitly some sort of Open Ontology, but for a broad acceptance and to provide value to its users, we strongly suppose that communicational relationships among closed “islands” of knowledge like contradiction or agreement need to be made explicit formally and technically. In this regard, the empirical derivation and stochastic modeling of “open” meta-data also seems inevitable if the set of knowledge sources is either very large, or fluctuates, or generates indefinite information. These mechanisms are not meant to be replacements for the usage of first-order predicate logic for Web reasoning, but instead as a completion which could




Figure 4. Proposed Semantic Web layers

be introduced gradually. E.g., the Resource Description Framework RDF(S) already has low-level reification capabilities, which could be used for elementary Social Reifications as described in [11,10], but in general this would require an appropriate specification of this kind of usage and technical and theoretical support for its realization.

As an example for the problems arising from traditional approaches to the annotation of web resources of meta-data,

a well-known problem is constituted through the notorious lack of reliable, impartial descriptions of publicly accessible resources such as web sites (or the content they provide, respectively), shared files or web services. If a resource description (RD) is available, in most cases the description is provided by the original resource provider, which at least for commercial supplies restricts its trustworthiness and might makes it as questionable as any other kind of advertisement. In contrast, recommender systems based on the evaluation of access statistics, voting or resource content analysis try to ascertain the “objective” value of resources. E.g., collaborative filtering recommender systems provide filter criteria for site classification, that classify the rated site in terms of “appropriate / inappropriate” or “interesting / uninteresting”, based on the surfing behavior of a more or less homogenous group of users with a common interest profile (user community) or implicit majority voting processes like Google’s “PageRank” algorithm, which essentially count the number of hyperlinks linking to the rated site. As a supplement or as a competing approach, content-based filtering recommender systems try to analyze the content of web resources (usually by means of keyword counting and bayesian classification) and compare the results with the interest profiles of the web surfers. The main drawback of such filtering systems is their limitation to one-dimensional descriptions (amounting to something like “you like/dislike...” or “...suits your needs”) based on presumed

descriptions which are in addition reliable and semantically rich can currently only be provided by humans, for example journalists and experts, or through discussion forums (e.g., newsgroups and threaded message boards), with well-known shortcomings like the absence of a machine readable encoding, which makes it almost impossible for information agents like web spiders to analyze these descriptions. Although the Semantic Web effort addresses the problem of missing machine-understandability of web site descriptions, it currently focuses primarily on the specification of languages and tools for the representation of consistent semantics and ontologies, not on the process of information gathering and rating itself, and it is just beginning to take into consideration phenomena like social RD impact [6], conflicting opinions, information biasing by commercial interests, and inconsistent or intentionally incorrect information. In contrast to traditional approaches, our goal is to provide social resource descriptions, which are obtained from the contribution of multiple, conflicting opinions represented by interacting information agents (so-called (resource) description agents), and which are rich (i.e., with unrestricted multidimensional description criteria, multiple levels of generalization and unveiled information about the social relationship of their contributors) - somewhat comparable to a computational “resume” of a human discussion. Our approach tackling this issue is presented in [12,11,10].

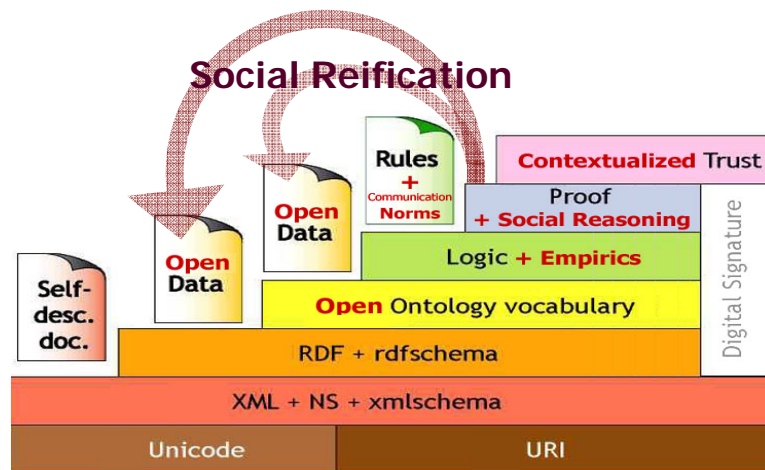


Figure 5. Socially enhanced layers of the Semantic Web

predilections. This approach does not provide much help for the process of interest forming, which should precede any recommendation and filtering. Likewise, systems based upon trust networks [8], which are expected to be especially important for the identification of trustworthy web services, cannot do much if there is no trust yet related to a specific rater, topic or object, or the potential trusting person is anonymous. Another problem is the apparent black-box character of many (commercial) recommender systems, which on the one hand provides some protection against manipulation, but on the other hand seriously restricts their trustability. For these reasons, transparent, balanced

5. CONCLUSION

There is an obvious and rapidly growing need for knowledge-based systems capable of running in open environments with autonomous knowledge sources and users, given the increasing interoperability and interconnectivity among computing platforms. Taking the key properties of such environments like information agent autonomy and the emergence of meaning from interaction seriously concerning the development of ontologies is a great challenge. On the one hand, ontologies should provide a stable ground for

communication and subsequent knowledge modeling, on the other hand, in open environments concept descriptions tend to be semantically inconsistent, they emerge from competing beliefs and goals, and a priori there is no such thing as a commonly agreed “truth” (in general, maybe not even a trend towards such a thing is assumable). To come up to each of these two contradictory aspects must be a core concern of the communication-oriented paradigm of knowledge modeling and management, and is the basic motivation underlying the work described here. To this end, we have proposed Open Ontologies as the basic step towards the modeling and representation of knowledge heterogeneity using a “sociality aware”, communication-based approach. To explore and to work out such a new perspective constitutes a long-term scientific and practical endeavor of considerable complexity. This is why it is not surprising that this introductory work does not answer all relevant issues and leaves room for theoretical and practical specification. First, the quite broad and general concept of Open Ontologies needs to be specialized in order to provide multiple, simplified types of Open Ontologies tailored for concrete fields of applications. E.g., Expectation Networks provide a descriptive power similar to distributed Markov processes, which might be not necessary for simple applications. In such cases, other social structure formalisms might be usable to obtain Social Reification. Second, there are relationships of our approach to the use of discourse ontologies for the annotation of user-created documents [e.g. 24] that could be further investigated. Whereas these approaches focus on more or less informal decision support instead of formal ontologies and automated reasoning, synergies with our proposal could be worth investigating. The same applies for data mining techniques, especially those used to uncover hidden social relationships among users (social data mining and *reality mining*), or for user prediction, recommender systems and natural language documents classification. Very likely, these techniques can be used in combination with Social Reification and Open Ontologies to widen the applicability scope of both areas.

Acknowledgements. We would like to thank the reviewers of the workshop on Agent-Mediated Knowledge Management (AMKM’04) at ECAI’04 for their useful comments on an earlier version of this paper.

REFERENCES

[1] A. Maedche, F. Nack, S. Santini, S. Staab, L. Steels. Emergent Semantics. IEEE Intelligent Systems, Trends & Controversies, 17(2), 2002.
 [2] J. Heflin, J. A. Hendler. Dynamic Ontologies on the Web. Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence, p. 443 - 449, 2000.
 [3] V.A.M. Tamma. An Ontology Model Supporting Multiple Ontologies for Knowledge Sharing. PhD Thesis, The University of Liverpool, 2002.
 [4] A. Farquhar, A. Dappert, R. Fikes, W. Pratt. Integrating Information Sources using Context Logic. Proceedings of the

AAAI Spring Symposium on Information Gathering from Distributed Heterogeneous Environments, 1995.
 [5] C. Castelfranchi. Conflicts Ontology. In R. Dieng, H.J. Mueller (Eds.), Conflicts in Artificial Intelligence. Springer, 2000.
 [6] <http://www.w3.org/2001/sw/meetings/tech-200303/social-meaning/>
 [7] G. H. Mead. Mind, Self, and Society. University of Chicago Press, 1934.
 [8] J. Golbeck, B. Parsia, J. Hendler. Trust Networks on the Semantic Web. Proceedings of Cooperative Intelligent Agents, 2003.
 [9] L. Wittgenstein. Tractatus Logico-Philosophicus. Routledge, Cambridge, Mass., 2001.
 [10] M. Nickles, Towards a Multiagent System for Competitive Website Ratings. Research Report FKI-243-01, Technical University Munich, 2001.
 [11] M. Nickles, G. Weiss. A Framework for the Social Description of Resources in Open Environments. Proceedings of the Seventh International Workshop on Cooperative Information Agents (CIA). Lecture Notes in Computer Science, Volume 2782. Springer, 2003.
 [12] M. Nickles, G. Weiss. Agent-based Social Assessment of Shared Resources. In Proceedings of the 2nd International Workshop on Agent and Peer-to-Peer Computing (AP2PC’03). Lecture Notes in Computer Science, Springer, 2003.
 [13] N. Luhmann. Social Systems. Stanford University Press, Palo Alto, 1995.
 [14] T. Berners-Lee, J. Hendler, O. Lassila. The Semantic Web. Scientific American 184(5), 2001.
 [15] M. Nickles, M. Rovatsos, G. Weiss. Empirical-Rational Semantics of Agent Communication. Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS’04), 2004.
 [16] M. Nickles, T. Froehner. Social Reification for the Semantic Web. Research Report FKI-248-04, AI/Cognition Group, Department of Informatics, Technical University Munich, 2004. To appear.
 [17] J. Delgado. Agent-based Recommender Systems and Information Filtering on the Internet. PhD. Thesis. Nagoya Institute of Technology, 2000.
 [18] M. Nickles, M. Rovatsos, W. Brauer, G. Weiss. Towards a Unified Model of Sociality in Multiagent Systems. International Journal of Computer and Information Science (IJCIS), Vol. 5, No. 1, 2004.
 [19] V. Dignum, J.-J. Meyer, H. Weigand, F. Dignum. An Organizational-oriented Model for Agent Societies. In: Proceedings of International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA’02), 2002.
 [20] M. Bonifacio, P. Bouquet, Paolo, R. Cuel. Knowledge Nodes: The Building Blocks of a Distributed Approach to Knowledge Management. Technical Report DIT-02-066, Informatica e Telecomunicazioni, University of Trento, 2002.
 [21] M.A. Smith. Netscan: Measuring and Mapping the Social Structure of Usenet. 17th Annual International Sunbelt Social Network Conference, 1997.

- [22] L. van Elst, V. Dignum, A. Abecker (Eds). Agent-Mediated Knowledge Management: Selected Papers, Lecture Notes in Artificial Intelligence, Springer, Volume 2926, 2004.
- [23] R. Cobos, X. Alamán. KNOWCAT: A Tool for Knowledge Crystallisation. First Technical Workshop of the Department of Computer Science & Engineering. Universidad Autónoma de Madrid, 2000.
- [24] V. Uren, B. Sereno, S. Buckingham Shum, G. Li. Interfaces for Capturing Interpretations of Research Literature, Distributed and Collaborative Knowledge Capture (DC-KCAP) Workshop, 2003.