Social Acquisition of Ontologies from Communication Processes

Special Issue "Formal Ontologies for Communicating Agents"

Matthias Nickles [a],
[a] *Artificial Intelligence/Cognition Group, Department of Computer Science, Technical University of Munich, Boltzmannstr. 3, D-85748 Garching bei München, Germany.*
E-mail: nickles@model.in.tum.de, Tel.: +49(89)289-17213, Fax: +49(89)289-17207

Abstract. This work introduces a formal framework for the social acquisition of ontologies which are constructed dynamically from overhearing the possibly conflicting symbolic interaction of autonomous information sources, and an approach to the pragmatics of communicated ontological axioms. Technically, the framework is based on distributed variants of description logic for the formal contextualization of statements wrt. their respective provenance, speaker's attitude, addressees, and subjective degree of confidence. Doing so, our approach demarcates from the dominating more or less informal approaches to context and provenance representation on the semantic web, and carefully distinguishes between communication attitudes such as public assertion and intention exhibited on the (semantic) web on the one hand, and mental attitudes such as private belief on the other. Furthermore, our framework provides formal means for the probabilistic fusion of controversial opinions, and presents a semantics of information publishing acts. Our approach provides an incomplex and genuinely social approach to knowledge acquisition and representation, and is thus expected to be widely applicable in fields such as the semantic web and social software, and possibly also in other distributed environments such as P2P systems.

Keywords: Contexts, Ontologies, Uncertainty Reasoning, Social Computing, Agent Communication, Belief Merging, Emergent Semantics

1. Introduction

From the viewpoint of Distribute Artificial Intelligence and research on multiagent systems, an ontology [12, 10, 19, 20, 7, 17, 9] is a conceptual description of a domain for communication and knowledge sharing purposes. The last few years have seen a tremendous rise of interest in such computational ontologies for the use in distributed settings with multiple, interacting participants - not only for multiagent systems, but also for open and distributed environments like the semantic web. Examples are the meaningful description of web services using ontologies [8, 1, 15, 16], *semantic blogging* [4], the *semantical annotation* of resources with meta-data [13]), and semantic approaches to P2P networking (e.g., [5, 16]). Applications in such environments require a shared domain semantics in order to support a mutual understanding among the distributed participants. The use of ontologies in such settings focuses thus mainly on the enabling of knowledge communication, sharing and reuse by means of the generation and provision of a conceptual common ground among the interacting parties.

But while the compilation, integration and sharing of information among heterogeneous information sources and users become more and more important, this development also leads to a sort of dilemma: on the one hand, ontologies should ease the reuse and sharing of knowledge, and thus should be stable and agreed. On the other hand, ontologies (being a special kind of knowledge) are themselves subject to difficulties known from the field of distributed knowledge and knowledge sharing, such as problems arising from misunderstandings, different names for the same concepts, or controversial viewpoints. Whereas some of these issues can be resolved at least in principle using techniques such as *ontology mapping* and *schema-based matching* (e.g. [18, 32, 33]), other issues might be much more difficult to resolve, or the dissolution is impossible on the level of ontology processing at least without disproportional measures that would lead to unwanted authoritative restrictions. For example, semantically inconsistent definitions of the same concept might have their origin in divergent world views of the ontology sources; an alignment of these world views, establishing an agreement, would - if practicable at all - lead to a loss in source autonomy and therefore decrease the openness, flexibility, and robustness of the application. In addition to the possibility that some of the above issues cannot be resolved, and the possible undesired loss of ontology source autonomy when attempting to get rid of them, there are further considerations to be taken into account when it comes to the integration of ontologies in distributed environments such as the web which have been largely neglected in traditional approaches: first, stable semantic conflicts are not just something one should get rid of by any means. Instead, *conflict knowledge* [14] (knowledge *about* conflicts such as controversies about the truth of some information) can provide valuable information about the attitudes, world views and goals of the respective knowledge sources. Furthermore, a set of distributed ontology sources and users forms a social layer. The explication and evaluation of this layer can provide the knowledge users with valuable meta-knowledge, and - if made explicit and visible - can be prerequisites for a subsequent resolution of conflicts regarding controversial information. Such meta-knowledge can be extremely useful - e.g., to gain a picture of the opinion landscape in controversial domains such as politics, product assessment and culture, and in current and forthcoming web applications which support explicitly or implicitly people interaction, like (semantic) blogging, discussion forums, collaborative tagging and folksonomies, and social computing in general. In this regard, it is also important to bear in mind that subjective intentions and goals do not just exist for communicating, intelligent agents, but also indirectly for other kinds of ontology sources like web documents, simply by virtue of their human creators. Second, in the absence of normative meaning governance, and due to the inherently dynamic nature of knowledge in open environments, mechanisms for ontology integration can only provide preliminary decisions about the reasonable modeling of domains, because in the end each ontology user can only decide for himself about the meaning, relevance and correctness of the given information, and these decisions might need to be revised in the course of time. In order to support such decisions, the decision makers (and eventually the web surfers) should be provided with as much information about the context of information found on the web as possible.

All these issues eventually stem from the autonomy of the involved information sources, since being autonomous means that these sources are possibly insincere and malevolent black-box agents with hidden goals. Since it is impossible or at least difficult to impose ontologies and other kinds of knowledge normatively in open environments like the internet, and because communication is the only feasible way of interaction among and with autonomous agents, communication plays a crucial role with respect to knowledge exchange in open environments. This is even the case where the communication takes place indirectly or asynchronously, such as by publishing and reading information on the web or in P2P networks.

These insights give rise to methods for knowledge acquisition and representation which focus on making formally explicit *pragmatic* aspects of information and in doing so its social, heterogeneous and possibly controversial nature. Most traditional approaches which tackle ontologies in the context of agent communication deal with ontologies as a pre-existing stable common ground for making communication possible, or try to resolve misunderstandings or to map equivalent concepts used by different agents, or rely on some kind of consensus finding process such as *ontology negotiation* [3]. In contrast, our work assumes already functioning communication with some basic, predefined and shared ontology. Communication is then used to establish another, emergent ontology in a social but not necessarily cooperative way: the agents exchange possibly contradicting opinions and intentions regarding axioms, and subjective knowledge is stored in the emergent ontology by a kind of "objectifying" technique called *social contextualization*. The "objectified" knowledge can then be used for knowledge sharing, trust assessment etc.

Doing so, this work builds upon [40] and presents a quite lean approach to i) the formal representation of semantical heterogeneity by means of *social contexts*, ii) the *social rating* of possibly contradictious or uncertain statements using probabilistic weights which can be determined from the respective degree of social approval and the trustability of the respective information sources, and iii) the computational acquisition and evolution of socially contextualized information contributions in the course of ongoing agent communication.

We consider our approach to be a contribution to the utilization of semantic web technologies for the development of social software - "social software" both in the sense of [31] (i.e., *social procedures*) and in the sense this term is commonly used by the Web 2.0 community, i.e., for software for the support of human collaboration and social networking. From a pragmatic point of view, the purpose of our framework is not far away from the purpose of the various *dynamic epistemic logics* in the context of agent communication (e.g., [44]), which can also be used for the modeling of knowledge update processes in the course of communication processes. But apart from the different formal background (contexts vs. modalities), our work is essentially neither about knowledge nor about belief in the usual sense, but about *public, ostensible opinions and intentions* as introduced in [39, 37]. What this means precisely will be pointed out later.

The remainder of this article is structured as follows: the following section outlines the key properties of our framework and Section 3 introduces the crucial notion of social context. Section 4 introduces a formal, description logic-based framework for the representation of social contexts, and Section 5 extends the formerly presented framework in order to allow for the fusion and probabilistic weighting of controversial opinions and inconsistent assertions in general. Section 6 presents and illustrates our formal communication model. Section 7 discusses related works, and Section 8 concludes.

2. Approach outline

2.1. Agree to disagree

The idea we start from - descending from the concept of so-called *Open Ontologies* [48] - is simple: in order to represent communicatively expressed information (which includes most of the information on the (semantic) web) within an ontology or a knowledge base, one should be able to make explicit the social contexts of communicated content, including and most important the context of assertion (assertional context) and public belief of a person ("Within the public belief system of person x, y is true"). Doing so, communicated statements can be assessed as being true within the subjective opinion system of its respective asserter (and often *only* there). I.e., in order to store such information within the knowledge base or ontology, it needs to be properly contextualized or reified first, at least in case its trustability can not sufficiently be assessed positively yet. This fact is so far no designated part of the formal semantics of prevailing web ontology representation languages. While there exist several approaches to provenance annotation and statement naming (cf. Section 7), these are either informal or semi-formal, or semantically transparent wrt. the original statement, i.e., they "just" provide meta-information about the provenance and possible other attributes of a dubious statement without stating a new contextually-enhanced true statement. In contrast, our approach aims for formal statements which express the *social meaning* of other statements which have been communicated (e.g., by putting them on the web).

We call the general concept *social contextualization*, being defined as "Creating statements by putting other statements into their social context(s)". An ontology which contains suchlike statements is called an *Open Ontology* [48, 35].

If we refer to a certain technique, we make distinctions between

social contextualization if, as in this article and formerly in [40], formal *contexts* of some kind are used (like in form of context logic, modular or package-based ontologies (e.g., [11]), or microtheories as in Cyc [34])

social reification [35] in case the underlying technique for formalizing statements about statements is reification, such as with RDF (e.g., [42])

social higher-order modeling in case a real higher-order logic is used to make statements about statements

social modalities if modal logic is used (es we did in [37, 36])

social annotation in case informal annotations are used (using, e.g., OWL annotation properties).

All of these variants amount to the "lifting" of statements to the social level, and in doing so the relativization of the statements' truth values. This comprises stating information about the communication-related (i.e., social) aspects of the original statement - minimally, information in order to identify the asserter(s) of the statement and optionally further socially-relevant information such as its addressees, the *alleged* degree of subjective certainty or social acceptance of the respective statement, or being the content of a speech act its communicative function (e.g., assertion, request, denial, proposal etc.).

In the ideal case (which is of course not always feasible in practice), social contextualization, reification etc. create objective knowledge from subjective opinions. E.g., a simple socially contextualized or reified knowledge base would contain a consistent set of statements in a kind of quotations form, like {"*Tom asserts that* sheep are pink", "*Tim denies that* sheep are pink"} instead of the inconsistent set of un-contextualized/un-reified statements {"Sheep are pink", "Sheep are not pink"}.
Social contextualization (we will stick to contextualization in this paper) could also provide means for the weighting of personal opinions against others in order to rank information, and for the aggregation of multiple informational contributions using techniques from the field of belief merging or judgement aggregation.
Social contextualization is of course not a universal solution for the issue of merging heterogeneous ontologies, but should be considered as a strong alternative meta-modeling approach suitable for cases where no sufficient basis for decisions about how to map, align or filter inconsistent information facets exists.
Social contextualization should not be confused with the provision of provenance information such as the use of *named graphs* [50] in the context of RDF. Instead, provenance annotation could be formalized as a special case of social contextualization. Traditional forms of provenance annotation enable the informal "reasoning" about the provenance of some statement or document only, for example in order to find out if the origin of the respective statement can be trusted. In contrast, social contextualization allows for the reasoning about explicitly subjective or group statements in a formal way, i.e., the social dimension is here an integral part of the logic formalism and not just a tag-like annotation. In addition, knowing a provenance is not necessarily identical with knowing that someone holds an opinion (the crucial propositional attitude of social contextualization), since the technical origin of some information does not necessarily manifest any significant social or propositional attitude. So, the annotation of information with meta-information about provenance is a related approach, but has a different semantics and different purposes - namely, our approach does not just attach some information like names to statements, but actually creates new statements on the formal level.
Social contextualization is also not directly related to social networks, as social contextualization is focussed on information (knowledge, opinions, beliefs...), communication and relationships of information, whereas social networks are about persons and relationships of persons. We believe that social contextualization/reification/higher-order modeling is the more general and adequate concept, since it is information- and not person-centric. Nevertheless, we believe that social networks and social contextualization could contribute complementary solutions to similar issues in semantic web research.

In order to illustrate and motivate social contextualization informally, we introduce the notion of a so-called *social semantic network* (SSN) . Social semantic networks can be seen as a simple graphical variant of the formal languages presented later. They are forms of *semantic networks* [49], which are one of the oldest graphical knowledge representation formats (in fact, a family of similar formats) known in artificial intelligence, psychology, philosophy, and linguistics. SSNs provide an intuitive extension of

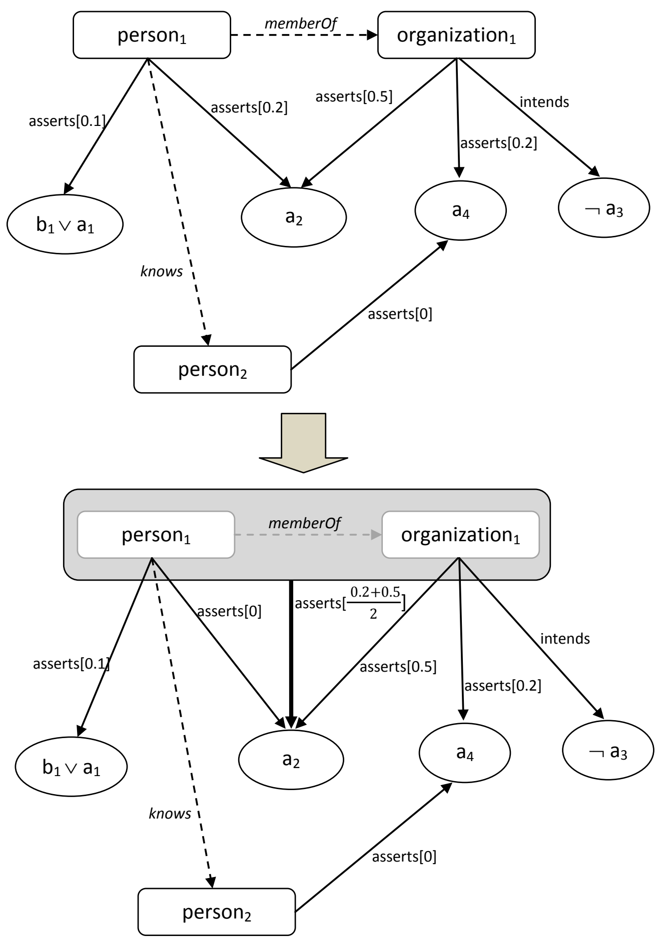

Fig. 1. Two Social Semantic Networks

traditional forms of semantic networks - effectively, a merger with the notion of *social networks* - in order to allow for the modeling of social entities like persons and organizations, and for the representation of their propositional attitudes towards logical statements, and optionally the respective degrees of belief (or interest, or desire, or any other attitude) regarding these statements[1]. Figure 1 provides two simple examples for SSNs.

In a SSN, the oval nodes stand for statements (in any language, e.g., logical formulas, RDF graphs, OWL axioms, or even semantic networks or SSNs). The rectangular nodes (boxes) (in our example labeled with "Person$_n$" or "Organization$_n$") represent social entities, and the edges denote either existing relationships among such entities (as in social networks, e.g., friendship or trust), or attitudes or "sentiments" of social entities regarding statements (edges between rectangular and oval nodes). In Fig. 1, the latter kind of edges denote either *assertions*, i.e., they state that the respective statement is from the point of view of the respective social entity true with a certain degree of certainty given in square brackets, or that an entity *intends* that a statement becomes true. In our example graph and throughout this article, the weights

[1] Social semantic networks are not to be confused with *semantic social networks*, which are semantic models of social networks and not social models of semantic networks.

for assertions are subjectively assessed probabilities, but SSNs can likewise be used to depict, e.g., the assignment of fuzzy values by agents.

SSNs also provide a simple yet flexible means for the graphical representation of all kinds of staged aggregations of individual and group opinions. For this purpose, rectangular nodes are aggregated in order to retrieve fused assessments of statements. In Fig. 1 (lower part), Person_1 and Organization_1 are integrated into a comprising social entity (large gray box), which asserts $b_1 \vee a_2$ with certainty 0.35 by averaging over the input degrees of certainty. But this way of fusion is again only for example - further fusion operators are proposed in Section 5. Note that the use of SSNs does not imply any form of truth functionalism or any specific way of reasoning.

In case the oval nodes in a SSN stand for or could be mapped to RDF graphs, simple SSNs like the ones in Fig. 1 can be translated into a somewhat corresponding form as RDF code, using RDF's reification-related vocabulary according to the following scheme (with the triple Subject/Predicate/ObjectS1 being the reified statement, corresponding to $b_1 \vee a_2$ in Fig. 1). Other Semantic Web-related means such as the use of *Notation3* (*N3*) [6] (a RDF syntax) could of course be used analogously, as well as OWL 1.1 (in OWL 1.1 it will be possible to reify axioms via assigning them URIs).

The example given below, similar to those presented in [42], is just for illustration - we do not encourage this specific way of "formalization" of truth or certainty assignment, since RDF's reification has no proper semantics (not to speak from the issue of uncertainty reasoning).

```
ssn:id1a2 rdf:type rdf:Statement .
ssn:id1a2 rdf:subject somesubject .
ssn:id1a2 rdf:predicate somepredicate .
ssn:id1a2 rdf:object someobject .

ssn:id1a2 ssn:assertedBy ssn:group1 .
ssn:id1a2 ssn:degreeOfCertainty "0.35"^^xsd:decimal .

ssn:group1 rdf:type rdf:Bag .
ssn:group1 rdf:_1 ssn:person1 .
ssn:group1 rdf:_2 ssn:organization1 .
```

Note that with this particular way of representing SSNs using RDF, one would need a dedicated reification quad for each "asserts", "intends" etc. edge within the SSN.

2.2. *Dynamic update of socially contextualized ontologies*

The concept of *dynamic ontologies* [21] provides a simple and general approach to the incremental revision of formal ontologies. After providing formal means for social contextualization in the following sections, we will show how such a dynamic ontology can be updated in dependency from observed communication acts.

The technical requirements for this acquisition process are:

- information agents able to communicate un-contextualized/un-reified (i.e., "1^{st}-level") statements like "c_1 *is a* c_2" or "c_1 *has a* c_2" at the content level of their communication language (note that we do *not* assume any cooperativeness or sincerity of the agents).

- a facility for the acquisition of ontologies from the observation of agent communication, e.g., a dedicated middle agent placed in the infrastructures of the respective application, called the *observer*. The observer is possibly one of the participating agents themselves.

- a facility for the low-level storage and querying of persistent knowledge (e.g., a database management system). This technical issue is not discussed in this work.
- optionally, a facility for the social reasoning upon the contextualized/reified (i.e., "2^{nd}-level") knowledge within the ontology.

The dynamic acquisition of ontologies comprises the following main tasks, which have to be performed in a loop as a continuous, incremental acquisition process during the whole period of agent communication:

1. Overhearing and, if necessary, filtering of agent communication
2. Derivation and/or adaptation of 2^{nd}-level knowledge according to the respective semantical model (e.g., empirically), i.e., the actual social contextualization/reification of 1^{st}-level.
3. Optionally, the stochastic generalization in terms of opinion fusion. This requires the grouping of multiple individual information sources in order to create a social group, and a social choice procedure for the weighting of controversial 1^{st}-level information.
4. Optionally, an alignment with normative 1^{st}-level information (e.g., some normative top-level ontology) or normative 2^{nd}-level knowledge (normative social structures, such as legal norms in order to prohibit the utterance of certain information), if desired. This is optional, and would require a given set of norms, like a normative ontology. We will not consider this case in our framework.

3. Representing communication attitudes using contexts

3.1. Basic social structures

Technically, our approach is based on providing an interrelationship of a social ontology or social knowledge base (KB) for the description of social concepts and individuals (like persons, agents and organizations, and maybe their relationships) on the one hand, and a set of possibly controversial or uncertain statements on the other. Special terms consisting of names from the social ontology/KB then identify so-called *social contexts* for the contextualization and optionally the fusion of semantically heterogeneous statements. This amounts to a technique which makes use of the context-driven partitioning of the knowledge space analogously to the approach presented in, e.g., [23, 25].

There is no canonical social ontology to be used with our approach. Basically any ontology could be used as long as it provides concepts, roles and instances for the modeling of the communication participants, like "Author", "Publisher" or "Reader", or, most basic, "Actor". Our approach also works in case the knowledge sources are given indirectly in form of the web locations or documents they use to articulate themselves. But we believe that information sources shall even in such cases be seen as active, autonomous, and communicating actors, since a description of the (semantic) web as a kind of huge distributed document base or database would be inadequate. We see the web rather as a place where actively pursued opinions and intentions compete with each other interactively, independently from the concrete ways this interaction is performed technically (directly or indirectly, synchronously or asynchronously...).

The following example ontology fragment (an improved version of the very simple social ontology used for social contextualization in [40]) will do for the purpose of this work:

Definition 1: Social ontology SO (example)

$Actor(person_1), Actor(person_2), Actor(person_3)$
...
$Communication(com_1), Communication(com_2), Communication(com_3), Communication(com_4)$
...
$Sender(com_1, person_2), Receiver(com_1, person_3)$
$Sender(com_2, person_1), Receiver(com_2, person_2)$
...
$Content(com_1, \text{"}a\_reified\_statement\text{"})$
$DegreeOfCertainty(com_1, 0.75)$
..
$SocialGroup(group_1), SocialGroup(group_2)$
...
$hasMember(group_1, person_1), hasMember(group_2, person_1)$
$Actor(group_1)$
...
$Actor(organization_1)$
...
$Sender(com_4, group_1), Receiver(com_4, organization_1)$
...
$CA(assertion), CA(information), CA(publicIntention)$
$Attitude(com_1, information), Attitude(com_2, assertion), Attitude(com_3, publicIntention)$
...
$Aggregation(fusedInformation)$

At this, $Actor$ is the category of the participating actors, whereby these can be any kind of information sources or receivers, like persons, organizations, documents, web services, as well as the holder of a so-called public intention (cf. below). $Communication$ is the category of elementary communication acts, described by the properties $Sender$, $Receiver$, $Attitude$ and $Content$ (the uttered statement or intention). A full-fledged approach would add further properties such as a time-stamp, but for many applications it will not be required to make SO explicit at all.

Senders and receivers can be the roles of any kind of actors, not only individual persons. E.g., a social group or an organization such as a company can also act as a sender. Social groups are modeled extensionally as sets, whereas organizations are legal entities. At this, it is very important to see that in our framework opinions and public intentions uttered by a certain group or organization can be modeled fully independently from the opinions and intentions of its members and subgroups. I.e., a social group as a whole could exhibit opinion p, whereas each individual group member exhibits $\neg p$ simultaneously. Of course, in reality the opinions of group members influence the opinion of the group, by way of social choice. But we think that no single way of group opinion settlement should be statically fixed. Instead, we will later introduce a special aggregation operator (informally denoted as $fusedInformation$ in SO) in order to model the quasi-democratic emergence of group opinions from individual opinions. But again, this is only one possibility: likewise, our framework allows to, e.g., model the case that a group always communicates the opinions of some dedicated opinion leader (dictatorship). It is also not necessarily the case that a social group as a whole forms a single actor at all.

In CA ("Communication Attitude"), $assertion$ denotes an *ostensibly* positive (i.e., approving wrt. truth) *communication attitude* (CA) regarding a certain statement, and at the same time the *ostensible* intention to make the addressee(-s) adopt a positive attitude also (e.g., "This product is the best-buy!"). This corresponds more or less to the communication act semantics we have introduced in detail in [39, 37, 38], and to Grice's conceptualization of speech acts as communications of intentions. $information$ is the same as $assertion$, but without the intention to make the addressees approve the respective statement too (but note that it is not possible to communicate an information p without the implicit assertion that p is indeed an information...). Both $information$s and $assertion$s are sometimes called "opinions" in this work. The pragmatic status of $information$, being a kind of "weak assertion", is debatable and mainly introduced for compatibility reasons wrt. [36], and we believe that $assertion$ is sufficient to model most cases of information dissemination using communication on the (Semantic) Web.

$publicIntention$ finally is the communication attitude of ostensibly intending that a statement shall become true. The attitude of *requesting* something from another actor is a subtype of $publicIntention$. As a simplification, we consider the attitude of *denying* something as identical with the positive attitude towards the negation of the denied statement. This would maybe too simple for the modeling of inter-human dialogs, but should do in the context of the less dynamic information exchange on the web. These attitudes should be sufficient to represent most information, publishing and desiring acts on the internet.
Note that $assertion$, $information$ and $publicIntention$ are no propositional attitudes in the usual mentalistic sense, as they do not need to correspond to any sincere (i.e., mental) beliefs or intentions of the actors. Instead, they are possibly insincere *communication* or *social* attitudes. As a consequence, they can not be treated like their mental counterparts. E.g., an actor might hold the opinion ϕ towards addressee one and at the same time $\neg\phi$ informing addressee two (while believing neither ϕ nor $\neg\phi$ privately). As another example, opinions could even be bought, in contrast to sincere beliefs: it is known that opinions uttered in, e.g., web blogs have sometimes been payed for by advertising agencies. In some sense, even *all* information on the web is "just" opinion due to the absence of a commonly accepted truth assessment authority.
$fusedInformation$ will be described later. It is used in place of communication attitudes, but it actually stands for the merging of opinions by some observer.

3.2. Social contexts and social semantics

Contexts (aka *microtheories*) have been widely used in AI since the early nineties, originally intended by McCarthy as a replacement of modal logic. [22, 23] propose a context operator $ist(context, statement)$ which denotes that $statement$ is true ("ist") within $context$. Building upon general approaches to contexts (specifically [23, 25]), and following our initial approach to social contextualization of description logic statements [40] and earlier works on social reification [42, 41, 35, 43], we will use the notation of context to express formally that certain statements are being publicly asserted (informed about, ostensibly intended to become true, denied...) on the web by some information-$Source$(s), optionally facing some specific $Addressee$(s). The latter implies that our use of the term "public" optionally comprises "limited publics" in form of closed social groups also. Thus, such *social contexts* model the *social semantics* of the contextualized information. Here, the term "social semantics" has a twofold meaning itself: first, it refers to the pragmatic effects of the *communicative function* information publication on the web has (essentially, our contexts correspond to kinds of speech acts). Second, the semantics is social in the sense that a fusion context can denote the meaning of a certain statement ascribed by *multiple* actors using some aggregation rule, e.g., the degree of truth assigned via consensus finding or voting, or other kinds of social choice.
Defined as conceptualizations of domains, formal ontologies are usually associated with consensual and relatively stable and abstract knowledge. Contexts in contrast provide a powerful concept underlying approaches which aim at coping with the distributiveness and heterogeneity of environments by means of localizing information. This dichotomy of ontologies on the one hand and contexts on the other has been recognized already, but only since recently, the synergies of both concepts are being systematically explored.
Such *social contexts* are special contexts which are used for the social contextualization of statements, i.e., their purpose is to express the social (= communicative) meaning of statements in a scenario like the web, with multiple synchronously or asynchronously communicating information providers and addressees.

The major task now is thus to define a type of context which allows to model the communicative function of asserted or otherwise published information.

The idea is now to use specific parts of the descriptions of individual elementary communications as defined in SO as identifiers of contexts, and the various contents (i.e., the contextualized statements) of these communications as context elements. That is, we maintain two ontologies: SO and a dynamic ontology with contexts, which uses parts of SO as identifiers. But for some applications, it will be sufficient to actually create and maintain only the latter ontology, whereas SO is given implicitly in form of the context identifiers only.

Definition 2: Social contexts

A *social context* is defined as a pair (id, c), with id being either a term which identifies communications in SO, or a *fusion context identifier* as specified below. c is the set of mutually consistent description logic statements (cf. next section) which corresponds to the set of contents $\{c : Content(com_i, c)\}$ of all communications com_i which share the respective partial description id. id is called the *context identifier*. A "partial description" of a communication means the description of the communication in terms of the properties $Sender$, $Receiver$ and $Attitude$. I.e., it comprises all role assertions for this communication, excluding those for the role $Content$ (which flows into c instead). Thus, social contextualization essentially puts statements into the same context iff the communications which contain these statement as their content share the same properties speaker, hearer, and attitude. In some sense, this "un-reifies" the reified statements within SO in order to obtain contextualized logical statements, and reifies other parts of SO in order to obtain context identifiers.

We use the following syntax for (non-fusion) context identifiers:

$$\stackrel{attitude}{source \longrightarrow addresse}$$

This term is obtained from a SO fragment

$$Sender(com, source), Receiver(com, addressee), Attitude(com, attitude)$$

for a certain com with $Communication(com)$. We also allow for context identifiers with sets of actors in place of the sender and/or the receiver (curly brackets omitted):

$$\stackrel{attitude}{source_1,...,source_n \longrightarrow addresse_1,...,addressee_n}$$

But note that social groups like $source_1, ..., source_n$ can still only occur in the sender role in (non-fusion) context identifiers if they *act as a group* as a sender or a receiver.

In contrast, *fusion context identifiers* will be used later in order to merge possibly inconsistent opinions uttered by multiple sources which do not necessarily form a social group with role $Sender$. The syntax of fusion context identifiers is

$$\stackrel{fusedInformation}{source_1,...,source_n \longrightarrow addressee}.$$

or in case $addressee$ is a social group alternatively:

$$\stackrel{fusedInformation}{source_1,...,source_n \longrightarrow addresse_1,...,addressee_n}.$$

As an abbreviation, we define ${}^{attitude}_{source_1,...,source_n} = {}^{attitude}_{source_1,...,source_n \longrightarrow Actor}$, with $Actor$ being the extension of $Actor$ in SO. I.e., the communication is here addressed to the group of *all* potential addressees like it is the case with information found on an ordinary web page. But note that at the same time a certain source can hold mutually inconsistent attitudes even towards different members or subgroups of $Actor$ (but not towards the same addressee).

4. A description logic with support for social contexts

We settle on the $\mathcal{SHOIN}(D)$ description logic (over data types D), because ontology entailment in the current quasi-standard OWL DL can be reduced to $\mathcal{SHOIN}(D)$ KB satisfiability [46]. Since we don't make use of any special features of this specific description language, our approach could trivially be adapted to any other description language or OWL variant, RDF(S), or first-order logic.

Definition 3: $\mathcal{SHOIN}(D)$-ontologies

The context-free grammar of $\mathcal{SHOIN}(D)$ concepts C is as follows:

$$C \rightarrow A|\neg C|C_1 \sqcap C_2|C_1 \sqcup C_2|\exists R.C|\forall R.C$$
$$| \geq nS| \leq nS|\{a_1, ..., a_n\}| \geq nT| \leq nT|\exists T_1, ..., T_n.D|\forall T_1, ..., T_n.D$$
$$D \rightarrow d|\{c_1, ..., c_n\}.$$

At this, C denote *concepts*, A denote *atomic concepts*, R denote *abstract roles* or *inverse roles* of abstract roles (R^-), S denote *abstract simple roles*, the T_i denote *concrete roles*, d denotes a concrete *domain predicate*, and the a_i / c_i denote abstract / concrete *individuals*.

A $\mathcal{SHOIN}(D) - ontology$ is then a finite, non-empty set of TBox and ABox axioms $C_1 \sqsubseteq C_2$ (inclusion of concepts), $Trans(R)$ (transitivity), $R \sqsubseteq S, T \sqsubseteq U$ (role inclusion), $C(a)$ (concept assertion), $R(a, b)$ (role assertion), $a = b$ (equality of individuals), and $a \neq b$ (inequality of individuals). Concept equality can be expressed via mutual inclusion, i.e., $C_1 \sqsubseteq C_2, C_2 \sqsubseteq C_1$. Spelling out the semantics of $\mathcal{SHOIN}(D)$ is not required within the scope of this work, it can be found in [46].

Definition 4: SOC-OWL

Introducing ontologies and at the same time description logic knowledge bases with social contexts, we define *SOC-OWL* (*Social-Context-OWL* or simply "Social OWL") similarly to C-OWL [25]. Essentially, SOC-OWL adds a kind of "S-Box" ("social box", i.e., social contexts) to formal ontology languages. In the next section, an advanced language P-SOC-OWL will be introduced, which also allows for uncertainty reasoning.

A SOC-OWL ontology/KB parameterized with a social ontology SO is a finite, non-empty set $O = \{(id, s) : id \in Id, s \in AF\} \cup AF^i \cup B$, with AF being the set of all $\mathcal{SHOIN}(D)$ axioms, AF^i being such axioms but with concepts, individuals and roles directly indexed with social contexts (i.e., $AF^i = \{(id_i, C_h) \sqsubseteq (id_j, C_k), (id_i, a_h) = (id_j, a_k), ... : id_i, id_j \in Id\}$), and B being a set of *bridge rules* (see 4.1). A *social context* within O is a pair $(id, \{s : (id, s) \in O\})$.

Id is the set of all social context identifiers according to the social ontology SO (cf. Definition 1). The s within (id, s) (i.e., OWL-DL axioms) are called *inner statements* which are said to "be *true* (or *intended* in case of $publicIntention$) within the respective context".

Examples (with multiple axioms per row and (id, a) written as $id\ a$):

$ControversialPerson(columbus)$ ${}^{assertion}_{tina \longrightarrow tim,tom} Hero(columbus)$
${}^{assertion}_{tim,tom \longrightarrow tina}(\neg Hero)(columbus)$ ${}^{assertion}_{tim,tom \longrightarrow tina} Exploiter(columbus)$

${}^{assertion}_{tom}(\neg Hero)(columbus)$

This SOC-OWL ontology (modeling as a whole somewhat a neutral point of view, like taken by an ideal Wikipedia article) expresses that the (fictive) persons Tim and Tom hold the opinion towards Tina that Christopher Columbus was not a hero but an exploiter (of the natives), while Tina does allegedly believe that the opposite is true. But there is consensus of the whole group that Christopher Columbus is a controversial person. Notice that without explicit further constraints, bridge rules or meta-axioms, different social contexts are logically fully separated. E.g., using only the above ontology it could *not* be inferred that ${}^{information}_{tina \longrightarrow tim}ControversialPerson(columbus)$, because $ControversialPerson(columbus)$ as an abbreviation of ${}^{information}_{tina,tim,tom \longrightarrow tina,tim,tom}ControversialPerson(columbus)$ in the example above is uttered/addressed *exactly* by/to *the single social group of all* participants. This principle allows to model the realistic case that someone conforms with some group opinion, but states some inconsistent opinion towards other groups (even a subgroup of the former group). Of course the co-presence of two or more inconsistent statements which indicate that a certain actor is insincere (as it would be the case with ${}^{assertion}_{tina \longrightarrow tim}(\neg C)(x)$ and ${}^{assertion}_{tina \longrightarrow tom}C(x)$ were contained within the same SOC-OWL ontology, which would be perfectly legal) could usually not be acquired directly from the web, since such actors would likely exhibit inconsistent opinions using different nicknames. Instead, some social reasoning or social data mining techniques would be required to obtain such SOC-OWL knowledge.

Obviously, each SOC-OWL statement $(contextId, statement)$ corresponds to the "classic" [22, 23] context statement $ist(context, statement)$. But unfortunately, this "real" ist operator could not simply be made a first-class citizen of our language (which would allow for the nesting of context expressions), at least not without getting into trouble defining a semantics of the language, or without making the semantics as problematic as that of RDF-style reification. Instead, we allow for *bridge rules* and meta-axioms in order to interrelate social contexts.
The core idea underlying the following semantics of SOC-OWL is to group the axioms according to their social contexts, and to give each context its own interpretation function and domain within the model-based semantics, corresponding to the approach presented in [25]. In addition, we will provide meta-axioms (constraints) and bridge rules in order to state the relationships among the various communication attitudes (somewhat similarly to modal logic axiom schemes such as the well-known KD45 axioms of modal belief logic), and to allow for the interrelation of different attitudes, even across different contexts. E.g., we would like to express that a communication attitude such as ${}^{assertion}_{tina \longrightarrow tim,tom}(\neg Exploiter)(columbus)$ implies (intuitively) ${}^{publicIntention}_{tina}({}^{information}_{tim,tom \longrightarrow tina}(\neg Exploiter)(columbus))$, i.e., that Tina not only expresses her ostensible beliefs, but also ostensibly intends that others adopt her opinion.

Definition 5: Interpretation of SOC-OWL

A SOC-OWL *interpretation* is a pair $(I, \{e_{i,j}\}_{i,j \in Id})$ with $I = \{I_{id}\}$ being a set of *local interpretations* I_{id}, with each $I_{id} = \langle \triangle^{I_{id}}, (.)^{I_{id}} \rangle$, $id \in Id$. $e_{i,j} \subseteq \triangle^{I_i} \times \triangle^{I_j}$ is a relation of two *local domains* $\triangle^{I_{id}}$ ($e_{i,j}$ is required for the definition of bridge rules in B (Definition 4) as explained later in 4.1). $(.)^{I_{id}}$ maps individuals, concepts and roles to elements (respectively subsets or the products thereof) of the domain $\triangle^{I_{id}}$.
To make use of this interpretation, contextualized statements of SOC-OWL impose a grouping of the concepts, roles and individuals within the inner statements into sets C_{id}, R_{id} and c_{id} [25]. This is done in order to "localize" the names of concepts, individuals and roles, i.e., to attach to them the respective local interpretation function I_{id} corresponding to the social context denoted by $id \in Id$:
concretely, the sets C_{id}, R_{id} and c_{id} are defined inductively by assigning the concepts, individuals and role names appearing within the $statement$ part of each SOC-OWL axiom/fact $(context_{Id}, statement)$

to the respective set C_{id}, c_{id} or R_{id}. With this, the interpretation of concepts, individuals etc. is as follows:

$C^{I_{id}} =$ any subset of $\triangle^{I_{id}}$ for $C \in C_{id}$
$(C_1 \sqcap C_2)^{I_{id}} = C_1^{I_{id}} \cap C_2^{I_{id}}$ for $C_1, C_2 \in C_{id}$
$(C_1 \sqcup C_2)^{I_{id}} = C_1^{I_{id}} \cup C_2^{I_{id}}$ for $C_1, C_2 \in C_{id}$
$(\neg C)^{I_{id}} = \triangle^{I_{id}} \setminus C^{I_{id}}$ for $C \in C_{id}$
$(\exists R.C)^{I_{id}} = \{x \in \triangle^{I_{id}} : \exists y : (x, y) \in R^{I_{id}} \wedge y \in C^{I_{id}}$ for $C \in C_{id}, R \in R_{id}$
$(\forall R.C)^{I_{id}} = \{x \in \triangle^{I_{id}} : \forall y : (x, y) \in R^{I_{id}} \rightarrow y \in C^{I_{id}}$ for $C \in C_{id}, R \in R_{id}$
$c^{I_{id}} =$ any element of $\triangle^{I_{id}}$, for $c \in c_{id}$
(Interpretation of concrete roles T analogously)

Satisfiability and decidability

Given a SOC-OWL interpretation I, I is said to *satisfy* a (contextualized) statement ϕ ($I \models \phi$) if there exists an $id \in Id$ such that $I_{id} \models \phi$. A SOC-OWL ontology is then said to be "satisfied" if I satisfies each statement within the ontology (or statement set). $I_{id} \models (id, C_1 \sqsubseteq C_2)$ iff $C_1^{I_{id}} \subseteq C_2^{I_{id}}$, $I_{id} \models (id, R_1 \sqsubseteq R_2)$ iff $R_1^{I_{id}} \subseteq R_2^{I_{id}}$, $I_{id} \models (id, C(a))$ iff $a^{I_{id}} \in C^{I_{id}}$ etc., i.e., as in the semantics of $\mathcal{SHOIN}(D)$, but with socially indexed interpretations.
Note that using this extension, the inherited semantics and decidability of $\mathcal{SHOIN}(D)$ and C-OWL remain unaffected in SOC-OWL "within" each context, since the new interpretation function simply decomposes the domain and the set of concepts etc. into local "interpretation modules" corresponding to the contexts.

4.1. Bridge rules and cross-context mappings

According to Definition 4, a SOC-OWL ontology can optionally comprise bridge rules [25] B and various stronger relationships AF^i among classes, individuals and roles from different contexts. As an example, consider

$(context_i, x) \xrightarrow{\equiv} (context_j, y)$ in B, with x, y being concepts, individuals or roles.

Informally, such a bridge rule states that the x and y denote corresponding elements even though they belong to different contexts $context_i, context_j$.
With, e.g., $(^{assertion}_{tina}, columbus) \xrightarrow{\equiv} (^{assertion}_{tim,tom}, columbus)$ the interpretations of the "two Columbuses" would abstractly refer to the same object. Analogously, $\xrightarrow{\sqsubseteq}$ and $\xrightarrow{\perp}$ state that the first concept is more specific than the second, or that both concepts are disjoint, respectively. These relationships are given by the relation $e_{i,j}$ (Definition 5).
Formally: $I \models (context_i, x) \xrightarrow{\equiv} (context_j, y)$ iff $e_{i,j}(x^{I_i}) = y^{I_j}$) (resp. $e_{i,j}(x^{I_i}) \subseteq y^{I_j}$ and $e_{i,j}(x^{I_i}) \cap y^{I_j} = \emptyset$).
Please find details (which are out of the scope of this work) and analogously defined further bridge rules in [25].
A much stronger kind of relationship is stated by the syntax constructs where a concept, individual or role is directly indexed with a social context, as, e.g., in $(context_i, x) = (context_j, y)$, with x, y being concepts, individuals or roles.
Formally: $I \models (context_i, x) = (context_j, y)$ iff $x^{I_i} = y^{I_j}$ (analogously for $\sqsubseteq$ etc).

4.2. Meta-axioms

We state now some constraints regarding the social meaning of contexts, which will later be extended with (PMA5).

Actively asserting an opinion implies in our framework the intention of the source that the addressee(-s) adopt the asserted statement. With nested social contexts, we could formalize this using ${}^{assertion}_{s_1,\dots,s_n \longrightarrow a_1,\dots,a_m} \varphi \rightarrow$

$({}^{publicIntention}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_m}({}^{information}_{a_1,\ldots,a_m\longrightarrow s_1,\ldots,s_n}\varphi)$. But this "strong" and problematic nesting is not possible with our language.

The next meta-axiom simply demands that assertions include the attitude of informing the addressee:
(MA1) ${}^{assertion}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_m}\varphi \rightarrow {}^{information}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_m}\varphi$

In this work, we do not provide a full meta-theory corresponding to the KD(45) axioms of (e.g.) modal Belief-Desire-Intention logics (but see [39, 37]). Instead, we only demand that the inner statements of each context are mutually consistent (basic rationality):

(MA2) Each set a of statements such that for a specific $context$ all $(context, a_i), a_i \in a$ are axioms of the same SOC-OWL ontology, has an interpretation.

Furthermore, we demand - in accordance with many BDI-style logics - that the information/assertion contexts of a certain actor on the one hand and his intention context on the other do not overlap addressing the same set of addressees, i.e., an actor does not (ostensibly) intent what he (ostensibly) believes to be the case already:

(MA3) For each a such that $({}^{publicIntention}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_n}, a)$ is part of an SOC-OWL ontology o, no axiom/fact $({}^{information}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_n}, b)$, $b \vdash a$, is part of o (analogously for assertions).

The following constraints are *not* demanded, but could be helpful in application domains were mutual opinion consistency of subgroups is desired (we use $\bigwedge$ to abbreviate a set of SOC-OWL statements).
(MAx1) $({}^{attitude}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_n}\varphi) \leftrightarrow \bigwedge_{s\in 2^{\{s_1,\ldots,s_n\}}-\{\emptyset\}} {}^{attitude}_{s\longrightarrow a_1,\ldots,a_n}\varphi$
(MAx2) $({}^{attitude}_{s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_n}\varphi) \leftrightarrow \bigwedge_{a\in 2^{\{a_1,\ldots,a_n\}}-\{\emptyset\}} {}^{attitude}_{s_1,\ldots,s_n\longrightarrow a}\varphi$

But we can safely aggregate seemingly consented information in a separated fusion context:
(MA4) $\bigwedge_{s\in\{s_1,\ldots,s_n\}} (I_{\substack{information\\ s\longrightarrow a_1,\ldots,a_n}} \models \varphi) \rightarrow (I_{\substack{fusedInformation\\ s_1,\ldots,s_n\longrightarrow a_1,\ldots,a_n}} \models \varphi)$ (analogously for assertions). In general, such group opinions induce a ranking of multiple statements with the respective rank corresponding to the size of the biggest group which supports the statement (this can be used, e.g., for a majority voting on mutually inconsistent statements).

5. Social rating and social aggregation of subjective assertions

Building upon social contexts, the following extension of the previously presented logical framework is optional. It makes use of uncertainty reasoning and techniques from belief merging. They allow for i) the representation of gradual strengths of *uncertain opinions* held by individuals (corresponding to subjective probabilities) and social groups, and ii) the probabilistic *fusion* of semantically heterogeneous opinions held by different actors (basically by means of voting). Since the emergence of *folksonomies* can be seen as a social categorization process (social choice in form of collaborative tagging), ii) amounts to a generalization of folksonomies to social choice of ontology and knowledge base entries in general.
This feature is also useful in case traditional techniques to ontology integration fail, e.g., if the resulting merged ontology shall be accepted by all sources, but a consensus about the merging with traditional techniques to ontology mapping and alignment could not be found, or if the complexity of a high amount of heterogeneous information needs to be reduced by means of stochastic generalization. Probabilistic fusion is furthermore helpful in case statements shall be socially ranked, i.e., put in an order according to the amount of their respective social acceptance. In contrast to heuristical or surfer-behavior-related ways of information ranking or "knowledge ranking" such as those accomplished by most web search engines, the following approach is based on semantic opinion pooling [30].

In [29], the probabilistic extension $P - \mathcal{SHOQ}(D)$ of the $\mathcal{SHOQ}(D)$ description logic has been introduced. $\mathcal{SHOQ}(D)$ is very similar to $\mathcal{SHOIN}(D)$ and thus OWL DL, but does not have inverse roles, and is not restricted to unqualified number restrictions [46]. [29] shows that reasoning with $P - \mathcal{SHOQ}(D)$ is - maybe surprisingly - decidable. Instead of $P - \mathcal{SHOQ}(D)$, other probabilistic approaches (in the Bayesian sense) to semantic web and ontology languages could likely also be used as a basis for our approach, e.g., [27]. $P - \mathcal{SHOQ}(D)$ is now used to define a probabilistic variant of SOC-OWL.

Definition 6: P-SOC-OWL

A *P-SOC-OWL* ontology is defined to be a finite subset of $\{([p_l, p_u], id, a_i)\} \cup \{(id, a_i)\} \cup \{a_i\} \cup AF^i \cup B$, with $p_l, p_u \in [0, 1], id \in Id, a_i \in AF$, AF being the set of all well-formed $\mathcal{SHOQ}(D)$ ontology axioms, and B and AF^i as in the previous section. The syntax of $\mathcal{SHOQ}(D)$ can be obtained from that of $\mathcal{SHOIN}(D)$ by excluding inverse roles.
The $[p_l, p_u]$ are probability intervals. Non-interval probabilities p are syntactical abbreviations of $[p, p]$. If a probability is omitted, 1 is assumed.

Definition 7: Semantics of P-SOC-OWL

The semantics of a P-SOC-OWL ontology is simply given as a family of $P{-}\mathcal{SHOQ}(D)$ interpretations, each interpretation corresponding to a social context. Formally, a P-SOC-OWL interpretation is a pair $(PI, \{e_{i,j}\}_{i,j \in Id})$ with $PI = \{(PI_{id}, \mu_{id}) : id \in Id\}$ being a set of *local probabilistic interpretations* (each denoted as Pr_{id}), each corresponding to a probabilistic interpretation of $P - \mathcal{SHOQ}(D)$ and a social context with identifier id. $\mu_{id} : \Delta^{I_{id}} \rightarrow [0, 1]$ is a subjective probability function, and the $\Delta^{I_{id}}$ are the domains. The relation $e_{i,j}$ (required to state bridge rules) is defined analogously to SOC-OWL. When restricted to a certain context (using the respective interpretation), reasoning in P-SOC-OWL remains decidable. Individualistically assigned probabilities are constrained by the axioms of probability. But as with SSNs, no kind of (notoriously problematic) truth functionalism is provided beyond these axioms. E.g., even within a specific context, no rule is given as part of the semantics of P-SOC-OWL in order to infer the probability of $a \wedge b$ only from the probabilities of a and b.

Example:

$[0.5, 0.8]$: ${}^{assertion}_{tim,tom \longrightarrow tina} Exploiter(columbus)$ 0.7: ${}^{assertion}_{tina} Hero(columbus)$
0.9: ${}^{assertion}_{tim} Hero(columbus)$

This P-SOC-OWL ontology expresses inter alia that Tim and Tom (as a group, but not necessarily individually) hold the opinion that with a probability in $[0.5, 0.8]$, Columbus was an exploiter, while Tina does (publicly) believe he was a hero with strength 0.7, and Tim believes so with strength 0.9 (i.e., his private opinion disagrees with the public group opinion of him and Tom). In order to allow for a consistent fusion of opinions, we demand the following fusion meta-axiom, which effectively states how the probabilities of *social fusion contexts* are calculated. A social fusion context is a social context with more than one opinion source and a probability which pools the probabilities which subsets of the group assign to the respective statement. This allows to specify group opinions even if group members or subgroups do knowingly not agree with respect to this assertion. In this regard, we propose two versions of interpretation rules:

(PMA5') $(\bigwedge_{s_i \in \{s_1,...,s_n\}} (Pr_{\substack{information \\ s_i \longrightarrow addressees}} \models \varphi[p_i, p_i])) \rightarrow (Pr_{\substack{information \\ s_1,...,s_n \longrightarrow addressees}} \models \varphi[p, p])$
with $p = pool^{poolingType}((p_1, ..., p_n), extraKnowledge)$. At this, $Pr_{id} \models \varphi[l, u]$ attests φ a probability within $[l, u]$ in context id, and $extraKnowledge$ is any knowledge the pooling function might utilize in addition to the p_i (see below for examples). (Analogously for the attitude *assertion*.)

A problem with (PMA5') is that it can lead to unsatisfiability (due to inconsistencies) in case the derived probability p is different than a probability assigned explicitly by this group of people (remember that a group of agents is free to assign *any* truth value or probability to any statement, using *any* social choice procedure). A simple work around is to use a new kind of context with "attitude" $fusedInformation$, which is actually no speaker attitude of course, but a fusion operator type used by the observer who fuses opinions. Another possibility would be to introduce some kind of defeasible logic or priority reasoning which gives priority to explicitly assigned probabilities.

(PMA5) $(\bigwedge_{s_i \in \{s_1,...,s_n\}}(Pr_{\underset{s_i \longrightarrow addressees}{information}} \models \varphi[p_i, p_i])) \rightarrow (Pr_{\underset{s_1,...,s_n \longrightarrow addressees}{fusedInformation}} \models \varphi[p, p])$ (remainder as PMA5').

As for $pool^{poolingType}$, there are several possibilities: in the most simple case of "democratic" Bayesian aggregation given the absence of any opinion leader or "supra-Bayesian" [30], we define $pool^{avg}((p_1, ..., p_n), \emptyset) = \frac{\sum p_i}{n}$, i.e., $pool^{avg}$ averages over heterogeneous opinions. Using this aggregation operator, we could infer the following: 0.8: $^{assertion}_{tina,tim}Hero(columbus)$.
Social aggregation operators are traditionally studied in the field of *Bayesian belief aggregation* [30, 24] and *judgement aggregation*. The most common fusion operator extends $pool^{avg}$ with expert weights (e.g., stemming from factors such as the opinion holder's trustability or reputation, or social power degrees of the information sources):
$pool^{LinOP}((p_1, ..., p_n), (weight_1, ..., weight_n)) = \sum weight_i p_i$, with $\sum_{weight_i} = 1$. Also quite often, a geometric mean is used:
$pool^{LogOP}((p_1, ..., p_n), (weight_1, ..., weight_n)) = \kappa \prod_{i=1}^{n} p_i^{weight_i}$ (κ for normalization).
It is noteworthy that the operators given above do not deal with the problem of *ignorance* directly (e.g., by taking into account the evidence the information sources have obtained, as in Dempster-Shafer theory). But such ignorance could be modeled using the $weight_i$ of $pool^{LinOP}$ and $pool^{LogOP}$, and possibly using probability intervals instead of single probabilities. In case opinions with probability intervals $[p_i^l, p_i^u]$ shall be fused, the described fusion operators need to be accordingly applied to the interval boundaries.

One application of such rating in form of aggregated or individual probabilities is to take the probabilities (respectively, the mean values of the bounds for each interval) in order to impose an order (*ranking*) of the axioms of an ontology (T-Box as well as A-Box), so that inner statements can be directly ranked regard their degree of social acceptance, as in

0.8: $^{information}_{voters}innerStatement_1$ (highest social rating)
$[0.5, 0.8]$: $^{information}_{voters}innerStatement_2$
0.2: $^{information}_{voters}innerStatement_3$ (lowest social rating)

Again, such a ranking can also be easily used to transform inconsistent ordinary ontologies into consistent ontologies by a voting on the statements of the inconsistent ontology: in case there are inner statements which are mutually inconsistent, a ranking can be used to obtain a consistent ordinary (i.e., OWL DL) ontology by removing from each smallest inconsistent subset of inner statements the statements with the lowest rating until all remaining elements of each subset are mutually consistent.

6. Communication act semantics

Any data found in ontologies on web or any other content in open environments is in general, if taken "as is", neither knowledge nor private belief. Instead, such content needs to be interpreted as the content of communication acts. By means of this interpretation (which essentially unfolds the "social semantics" in the (semantic) web and other essentially communication-based environments), the recipient can then

possibly transform the data as knowledge (e.g., via applying (dis-)trust), and, usually in a previous step, as opinions and other social attitudes. The step from the respective speech act of communicating (asserting, requesting, denying...) content to its meaning in terms of social attitudes is specified in the following. Note that although we use an action notation, it would of course not be necessary to rewrite existing ontologies explicitly as speech acts, since the speech act meaning of public content is implicit. Some OWL document is essentially nothing else than the propositional content of an implicitly given assertive speech act.
But technically, the following acts could also be used more or less directly within a document as "social semantic links" to external content (similar to the import-directive of OWL).
The following rules (stemming from [36], where their modal logic equivalents were presented) provide both a semantics for typical communication acts in terms of their pre- and post-conditions, and at the same time an alternative (but under-specified) operational semantics for social attitudes in terms of those communication acts which are allowed to take place when certain attitudes hold. We achieve both at the same time by following the scheme used for the (mentalistic) FIPA-ACL semantics [53] and its social attitude-based version [38] i.e., by specifying for each act its so-called *feasibility precondition* (FP) and its *rational effect* (RE) [38]. The former denotes (in first order logic) what needs to hold in order to being able to perform the respective act, the latter denotes both guaranteed illocutionary and perlocutionary effect(s) of the respective act. The practical use of providing FPs is surely debatable, but at least FPs help to model *social commitment*: if an agent asserts that φ, the FPs of the $inform$-act demand i.a. that she will not assert $\neg\varphi$ subsequently without having retracted explicitly her formal commitment to φ.

Communication acts are denoted as

$id_1 : author.Performative(audience, \omega | id_2)$,
with ω being a logical statement (the so-called *content* of the act).

If $audience$ is omitted, the specific addressee of the respective act is either not known, or the act is addressed at an unspecific group of potential recipients (as it is usually in case when the locution is asynchronously provided in form of a web site, which addresses potentially the whole internet).
id_1 optionally assigns the act an unequivocal number (or a time stamp or a URI), or refers to another act. A helper function $content$ is used to map an identifier to the content of the respective act.
Note that in regard to our semantics of communication acts it is not of importance whether the acts are technically performed asynchronously with its potential reception (like it is the usual case on the web), or synchronously, e.g., as steps in an interactive argumentation process.

Furthermore, we assume that the observer of the agent discourse maintains a dynamic ontology Φ, either a conjunctive set of SOC-OWL or P-SOC-OWL axioms. This ontology is incrementally acquired and revised in the course of agent communication as follows. The ontology shall be logically consistent and might be non-empty initially, i.e., contain as further constraints a number of given statements (e.g., normative axioms). The subsequent states of this dynamic ontology are denoted as $\Phi_0, \Phi_1, \Phi_2, ...$ Let $s \vdash t :\Leftrightarrow I \models s \rightarrow I \models t$.

- $id : s.assert(A, \omega)$ (author s asserts ω towards the recipients A)

  FP: $(\neg\exists^{assertion}_{id \longrightarrow A} \psi \in \Phi_t : \psi \vdash \omega) \wedge (\neg\exists^{assertion}_{id \longrightarrow A} \psi \in \Phi_t : \psi \vdash \neg\omega) \wedge (\neg\exists^{information}_{A \longrightarrow id} \psi \in \Phi_t : \psi \vdash \omega)$

  RE: $\overset{assertion}{\underset{id \longrightarrow A}{}} \omega \in \Phi_{t+1}$

  RE(P): $\overset{assertion}{\underset{id \longrightarrow A}{}} \omega \in \Phi_t \wedge Pr_{\underset{s_1,...,s_n \longrightarrow A}{fusedInformation}} \models \omega[p,p]$, with $p = pool^{poolingType}((p_1, ..., p_n), \emptyset)$ and $\bigwedge_{s_i \in \{s_1,...,s_n\}} (Pr_{\underset{s_i \longrightarrow addressees}{information}} \models \omega[p_i, p_i])$.
  RE(Trust): $Bel^p_{observer}(\omega)$, with $p = pool^{poolingType}((p_1, ..., p_n), \emptyset)$

and $\bigwedge_{s_i \in \{s_1,...,s_n\}}(Pr_{\substack{information \\ s_i \longrightarrow addressees}} \models \omega[p_i, p_i])$ using some pooling operator as proposed in the previous section.

The feasibility precondition (FP) here, which is more exhaustive than the one provided in [38] due to the restrictions of the underlying logics (P-)SOC-OWL, expresses that i) in order to perform an assert act, neither the information source nor the audience have already publicly announced their alleged belief in ω already (otherwise the act would not make sense), and ii) that the information source does not contradict herself communicating with A (but she might expose a public assertion inconsistent with ω towards a different audience).

The postcondition RE expresses that it became known (to the observer) that s asserts ω (which includes the public intention of s to convince A that ω). In the case that no trustability or other expectations exist regarding the truth of ω and/or its provenance a_1, the postcondition makes the assert-act essentially a request do adopt a public assertion, with a more or less uncertain effect on the addressee. In case we use P-SOC-OWL, the observer shall also merge the observed content using RE(P). Of course, the observer is additionally free to choose the fused information as his own belief with strength p in case he trusts the sources $s_1, ..., s_n$ in RE(P) with respect to ω ($Bel^p_{observer}(\omega)$).

- $id : s.inform(A, \omega)$ (author s informs the recipients A that ω)

  FP, RE, RE(P): Analogously to act $inform$ (with context $^{information}_{...}$ instead)

  The inform act is simply a weaker form of the assert act in that the author does not necessarily aim to convince the receivers.

In order to provide convenient means for the agreement or disagreement with certain information, the following macro acts are proposed. We suppose they are particularly useful if the implemented language provides some sort of linkage facility, such as OWL's `owl:imports` and class or property descriptions in form of URI references. Interpreting the interlinked items (documents, class descriptions, meta-data etc.) as assertions, the graph formed from such items related by URL/URL references effectively maps to a communication process.

- $id_0 : s.agree(a, id_1) \equiv s.assert(\{a\}, content(id_1))$

- $id_0 : s.deny(a, id_1) \equiv s.assert(\{a\}, \neg content(id_1))$

We also propose the following *intend* and *request* acts, which (e.g.) allow to announce that a certain proposition is intended to be true, or to ask someone to publish a certain information. Note that our notion of intending includes as a special case *desiring* that another actor makes something true (like on request):

- $id : s.intend(A, \omega)$ (author s announces to an audience A that she intends that ω shall hold, e.g., by acting towards it, or asking other to do so)

  FP: $\neg\exists^{publicIntention}_{id \longrightarrow A} \psi \in \Phi_t : \psi \vdash \omega)$

  RE: $^{publicIntention}_{id \longrightarrow A}\omega \in \Phi_{t+1}$

6.1. Example

In order to demonstrate the properties and one possible application of our approach, this section presents a brief case study in form of a shortened *purchase negotiation* scenario (adapted from a scenario presented in [38, 36]), which should be quite typical for the semantic modeling of, e.g., seller/buyer platforms on the web using semantic web services.

The interaction roughly follows protocols for *purchase negotiation dialogue games*, but we omit some details which are not relevant for our purposes (e.g., specification of selling options). Although the example deals with negotiation, the approach is expected to be usable for the modeling of other types of interaction on the (semantic) web also (such as argumentation).

Our scenario consists of four participants $\{s_1, s_2, c_1, c_2\}$, representing potential car sellers and customers (implemented, e.g., in form of two seller web services and two agents of the customers). In the discourse universe exists two instances ϑ_1 and ϑ_2 of some car type ϑ (e.g., specimen of the Alfa Romeo 159).

The interaction course is presented as a sequence of steps in the following form. Note that the interaction course consists of multiple interlaced conversations among different sender/receiver pairs. In particular, c_2 is involved in two selling dialogues at the same time. The different dialogues shall be visible only for the participants (senders and receivers of the respective communication acts).

Utterance id. sender→receiver(-s): Descriptive act title
Message
Effect (optionally) gives the effect of the act in terms of social attitudes.

In contrast to *Effect*, *Private information (PI)* optionally unveils relevant mental attitudes before or after an act has been uttered and understood by the respective agents. The PIs are not determined by preceding communication acts, due to agent autonomy. They are also of course usually not available to observers on the web, and thus just given here for the reader's information.

In addition, we use a slightly abbreviated syntax in the following in order to make the protocol more readable. Some less important act effects are also omitted.

PI_{s_1}**:** $Bel_{s_1}\ discounts$

U1 $s_1 \rightarrow \{c_1, c_2\}$**: Information about discount**

$s_1.\mathsf{assert}(\{c_1, c_2\}, \neg discounts)$

Effect on context ontology:

$\Phi_1 = \Phi_0 \cup \{^{assertion}_{s_1 \longrightarrow c_1, c_2} noDiscounts\}$

Seller s_1 asserts that no discounts can be given while believing that the opposite is true (there might be the company policy that discounts should be given, but that might reduce the seller's individual profit).

Note that such a contradiction between private and public (communicated) beliefs or intentions could not be modeled using the omnipresent agent framework BDI or known semantic web languages (on the level of first-class language constructs).

U2 $s_1 \rightarrow \{c_2\}$**: Information about discount**

$s_1.\mathsf{assert}(\{c_2\}, discounts)$

Effect:

$\Phi_2 = \Phi_1 \cup \{^{assertion}_{s_1 \longrightarrow c_2} discounts\}$

While seller s_1 informed group $\{c_1, c_2\}$ that there would be no price discounts, he informs customer c_2 that this is not true (likely because s_1 thinks that c_2 is a valued customer whereas c_1 is not).

Such different, inconsistent assertions addressed to different (even nested) groups of addressees can also not be modeled using any current web semantics language (and also not by means of the BDI framework).

U3 $c_2 \rightarrow \{s_1\}$**: Query if car type has a high or a low accident rate**

$c_2.\mathsf{intend}(\{s_1\}, s_1.InformIfAccidentRate)$

Effect:

$\Phi_3 = \Phi_2 \cup \{^{publicIntention}_{c_1 \longrightarrow s_1} Done(s_1.InformIfAccidentRate)\}$

$PI_{s_1}: Bel_{s_1}\ accidentRateHigh(\vartheta)$

U4 $s_1 \rightarrow \{c_2\}$**: Information about accident rate**

$s_1.\mathsf{assert}(\{c_2\}, accidentRateLow(\vartheta))$

Effect:

$\Phi_4 = \Phi_3 \cup \{^{assertion}_{s_1 \longrightarrow c_2} accidentRateLow(\vartheta)\}$

Seller s_1 asserted $accidentRateLow(\vartheta)$ while thinking the opposite. Privately, c_2 believes this information and publicly agrees in the next step, but will revise her private (but not her public) belief later.

U5 $c_2 \rightarrow \{s_1\}$**: Expression of belief**

$c_2.\mathsf{inform}(\{s_1\}, accidentRateLow(\vartheta))$

Effect:

$\Phi_5 = \Phi_4 \cup \{^{information}_{c_2 \longrightarrow s_1} accidentRateLow(\vartheta)\}$

Since c_2 has himself asked s_1 to provide him the information uttered in the previous step, he publicly believes it.

$PI_{c_2} : Bel_{c_2} \neg accidentRateHigh(\vartheta)$

U6 $c_2 \rightarrow \{s_2\}$**: Query if car type has a high or a low accident rate**

$c_2.\mathsf{intend}(\{s_2\}, s_2.InformIfAccidentRate)$

To make sure, the potential buyer c_2 asks s_2 the same question.

U7 $s_2 \rightarrow \{c_2\}$**: Information about accident rate**

$s_2.\mathsf{assert}(\{c_2\}, accidentRateHigh(\vartheta))$

Effect: c_2 publicly believes the information facing s_2, and even trusts it for some reason privately more than the information given by seller s_1 earlier. Nevertheless, it remains true that he also still publicly believes the opposite towards the other seller.

$PI_{c_2} : Bel_{c_2}\, accidentRateHigh(\vartheta)$

U8 $c_2 \rightarrow \{s_2\}$**: Propose to buy at a low price**

$c_2.\mathsf{intend}(\{s_2\}, buy(\vartheta_2, 4000£))$

U9 $s_2 \rightarrow \{c_2\}$**: Accept proposal**

$s_2.\mathsf{intend}(\{c_2\}, sell(\vartheta_2, 4000£))$

Effect (together with the previous act):

$\Phi_9 = \Phi_8 \cup \{^{publicIntention}_{c_2 \longrightarrow s_2} buy(\vartheta_2, 4000£)\} \cup \{^{publicIntention}_{s_2 \longrightarrow c_2} sell(\vartheta_2, 4000£)\}$

(i.e., c_2 and s_2 are publicly committed to buy (or sell, respectively) ϑ_2 at the price of $4000£$ now).

7. Related work

The goal of this work is to provide a formal language for and a semantics of possibly contradictory assertions and other kinds of communication on the web. Doing so, we settle on the "opinion level" where neither beliefs are visible (due to the mental opaqueness of the information sources) nor criteria for the selection of useful knowledge or semantic mappings from/among heterogenous information exist initially. This is in strong contrast to the traditional aim of information integration and evolution for the determination of some consistent, reliable "truth" obtained from the contributions of multiple sources as in *multiagent belief representation and revision* (e.g. [52]) and approaches to ontology alignment, merging and mapping. Apart from the research field of knowledge and belief integration, the storage of heterogeneous information from multiple sources also has some tradition in the fields of *data warehousing* and *federated databases*, and view-generation for distributed and enterprise database systems [28], whereby such approaches do not take a social or communication-oriented perspective. *Opinions* are treated in the area of the (non-semantic) web (e.g., *opinion mining* in natural language documents) and in (informal) knowledge management (e.g., *KnowCat* [45]). The assignment of provenance information is mostly based on *tagging* and *punning* techniques, or makes use of the problematic reification facility found in RDF(S). Advanced approaches to provenance tracking (e.g., *named graphs* [50], the quoting facility of *Notation3* (*N3*) [6]) or social network-based approaches to provenance tracking [51] are already very useful if it is required to specify who contributed some information artifact, e.g., in order to assess information trustability. But these approaches either do not provide a proper formal semantics (concerning RDF-based approaches), or they allow to specify that someone "asserts" some information, but they do not actually lift statements in the sense of Section 2.1 and they do not tell semantically what asserting (requesting, denying...) actually means. The opinion calculus which we have presented in [39, 37, 36] provides such capabilities, but it is based on modal logic, which is often considered more complicated and less intuitive than context logics or distributed description logics. *Contextual reasoning*, which is underlying the framework presented in this paper, has been proposed for the modeling of individual beliefs of agents several times, starting with the original works on contexts by McCarthy [22]. Independently from web-related approaches, contexts have been widely used for the modeling of distributed knowledge and so-called *federated databases*, see e.g. [2, 47]. [26] provides an approach to the context-like grouping of RDF statements, but without taking into account the social (i.e., communicative) aspect, and in a more or less informal fashion. A mature approach, focusing on the aggregation of RDF(S) graphs using contexts, was presented in [23], and [25] provides a general formal account of contexts for OWL ontologies. But to our knowledge almost all traditional approaches to web semantics lack a specific social perspective wrt. formal semantics, i.e., they do not take into account on the level of semantics that providing information on the web *as-is* (i.e., before the application of trust or other judgement means) is a social act. Furthermore, the fact that public information can neither be classified directly as belief (since information sources on the web are mentally opaque) nor as knowledge (since there exists no commonly accepted knowledge authority), is still ignored by most approaches to web semantics. For an in-depth discussion of the differences between these epistemic modalities, please refer to [37, 38] instead.

8. Conclusion

There is an obvious and rapidly growing need for knowledge-based systems capable of running in open environments like the web with autonomous knowledge sources and users, given the increasing interoperability and inter-connectivity among computing platforms. Taking the key properties of multiagent systems, the semantic web and other open environments like actor autonomy and the emergence of meaning from interaction seriously wrt. the creation of ontologies is a great challenge: On the one hand, ontologies should provide a stable ground for user information, agent and user communication, and subsequent knowledge modeling, on the other hand, in open environments concept descriptions tend to be semantically inconsistent, they emerge from a possibly very high number of competing beliefs and goals, and a priori there might be no such thing as a commonly agreed “truth” (in the “real world”, not even a discursive trend towards such a thing can be assumed). To cope with these two seemingly contradictory aspects must be a core concern of the communication-oriented paradigm of knowledge modeling and representation, and is the basic motivation underlying the work described here. To this end, we believe to have proposed a crucial yet intuitive step towards the modeling and formal representation of knowledge heterogeneity for the semantic web. But to further explore and work out this new pragmatic and "social" perspective in the field of knowledge representation certainly constitutes a long-term scientific and practical endeavor of considerable complexity.

Acknowledgements

This work was funded by Deutsche Forschungsgemeinschaft (DFG) under grant BR 609/13-1 (research project *Open Ontologies and Open Knowledge bases*).

Thanks to the reviewers for many valuable suggestions.

References

[1] S. Agarwal, S. Handschuh, S. Staab. Annotation, composition and invocation of semantic web services. Journal on Web Semantics, 2(1):1–24, 2005.
[2] M. Bonifacio, P. Bouquet, R. Cuel. Knowledge nodes: the building blocks of a distributed approach to knowledge management. 2002.
[3] S. C. Bailin. Ontology negotiation in a community of agents. Lecture Notes in Computer Science Vol. 1871, Springer, 2001.
[4] S. Cayzer. Semantic blogging. `http://www.hpl.hp.com/personal/steve_cayzer/semblog.htm`

[5] M. Ehrig, P. Haase, F. van Harmelen, R. Siebes, S. Staab, H. Stuckenschmidt, R. Studer, C. Tempich. The swap data and metadata model for semantics-based peer-to-peer systems. In M. Schillo, M. Klusch, J.P. Müller, H. Tianfield (eds.), Procs. of MATES-2003. First German Conference on Multiagent Technologies, LNAI 2831. Springer-Verlag, 2003.
[6] Tim Berners-Lee. Notation3. `http://www.w3.org/DesignIssues/Notation3.html`
[7] D. Fensel. Ontologies: a silver bullet for knowledge management and electronic commerce. Springer-Verlag, 2001.
[8] A. Gangemi, P. Mika, M. Sabo, D. Oberle. An ontology of services and service descriptions. Technical report, Laboratory for Applied Ontology (ISTC-CNR), 2003.
[9] A. Gómez-Pérez, M. Fernández-López, O. Corcho. Ontological engineering. Springer-Verlag, 2004.
[10] T.R. Gruber. Toward principles for the design of ontologies used for knowledge sharing. International Journal of Human-Computer Studies, 41:399–424, 1995.
[11] J. Bao, D. Caragea, V. Honavar. Towards collaborative environments for ontology construction and sharing. International Symposium on Collaborative Technologies and Systems (CTS 2006), pages 90Ű108. IEEE Press, 2006.
[12] T.R. Gruber. A translation approach to portable ontologies. Knowledge Acquisition, 5(2):199–220, 1995.
[13] S. Handschuh, S. Staab. Annotation of the shallow and the deep web. In S. Handschuh. S, Staab (eds.), Annotation for the Semantic Web, Frontiers in Artificial Intelligence and Applications, vol. 96. IOS Press, 2003.
[14] H.-J. Müller, R. Dieng (eds.). Computational Conflicts - Conflict Modeling for Distributed Intelligent Systems. Springer-Verlag, 2000.
[15] M Paolucci, T. Kawamura, T. R. Payne, K. P. Sycara. Semantic matching of web services capabilities. Procs. of the First International Semantic Web Conference. Springer-Verlag, 2002.
[16] M. Schlosser, M. Stinek, S. Decker, W. Nejdl. A scalable and ontology-based p2p infrastructure for semantic web services. Procs. of the 2nd International Conference on Peer-To-Peer Computing. IEEE Computer Scociety, 2002.
[17] S. Staab, R. Studer (eds.). Handbook on Ontologies in Information Systems. Springer-Verlag, 2003.
[18] V. Tamma. An Ontology Model Supporting Multiple Ontologies for Knowledge Sharing. PhD thesis, The University of Liverpool, 2002.
[19] M. Uschold. Building ontologies: towards a unified methodology. Procs. of Expert Systems '96, 1996.
[20] M. Uschold, M. Gruninger. Ontologies: principles, methods and applications. The Knowledge Engineering Review, 11(2):93–155, 1996.
[21] J. Heflin, J. A. Hendler. Dynamic Ontologies on the Web. Procs. of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence, 2000.
[22] J. L. McCarthy. Notes on formalizing context. In IJCAI'93, pages 555-562, 1993.
[23] R. V. Guha, R. McCool, R. Fikes. Contexts for the Semantic Web. Procs. of the Third International Semantic Web Conference, 2004.
[24] M. Richardson, P. Domingos. Building Large Knowledge Bases by Mass Collaboration. Technical Report UW-TR-03-02-04, Dept. of CSE, University of Washington, 2003.
[25] P. Bouquet, F. Giunchiglia, F. van Harmelen, L. Serafini, H. Stuckenschmidt. C-OWL: Contextualizing Ontologies. Procs. of the Second International Semantic Web Conference, LNCS 2870, Springer-Verlag, 2003.
[26] G. Klyne. Contexts for RDF Information Modelling. http://www.ninebynine.org/RDFNotes/RDFContexts.html, 2000.
[27] P. Costa, K.B. Laskey, K.J. Laskey. PR-OWL: A Bayesian Framework for the Semantic Web. Procs. of the First Workshop on Uncertainty Reasoning for the Semantic Web (URSW-05), 2005.
[28] J. Ullmann. Information Integration Using Logical Views. Procs. of the 6th International Conference on Database Theory. Springer-Verlag, 1997.
[29] R. Giugno, Th. Lukasiewicz. P-SHOQ(d): A Probabilistic Extension of SHOQ(d) for Probabilistic Ontologies in the Semantic Web. Procs. of JELIA '02 - European Conference on Logics in Artificial Intelligence. Springer-Verlag, 2002.
[30] R.M. Cooke. Experts in Uncertainty: Opinion and Subjective Probability in Science. Oxford University Press, 1991.
[31] R. Parikh. Towards a theory of social software. Procs. of Workshop DEON'02. Imperial College, 2002.
[32] H. Wache, T. Vögele, U. Visser, H. Stuckenschmidt, G. Schuster, H. Neumann, S. Hübner". Ontology-based integration of information — a survey of existing approaches. Procs. of the IJCAI-01 Workshop on Ontologies and Information Sharing, 2001.
[33] P. Shvaiko, J. Euzenat. A Survey of Schema-based Matching Approaches. Journal on Data Semantics (JoDS), IV, LNCS 3730, Springer-Verlag, 2005.
[34] Cyc, `http://www.cyc.com`
[35] T. Froehner, M. Nickles, G. Weiss. Towards Modeling the Social Layer of Emergent Knowledge Using Open Ontologies. Procs. of the ECAI-04 Workshop on Agent-Mediated Knowledge Management (AMKM-04), 2004.
[36] M. Nickles. Modeling Social Attitudes on the Web. Procs. of the 5th International Semantic Web Conference, 2006.
[37] F. Fischer, M. Nickles. Computational Opinions. Procs. of the 17th European Conference on Artificial Intelligence (ECAI'06), 2006.
[38] B. Gaudou, A. Herzig, D. Longin, M. Nickles. A New Semantics for the FIPA Agent Communication Language based on Social Attitudes. Procs. of the 17th European Conference on Artificial Intelligence (ECAI'06), 2006.
[39] M. Nickles, F. Fischer, G. Weiss. Communication Attitudes: A Formal Approach to Ostensible Intentions, and Individual and Group Opinions. Procs. of the 3rd International Workshop on Logic and Communication in Multiagent Systems (LCMAS'05), 2005.
[40] M. Nickles, R. Cobos. Social Contexts and the Probabilistic Fusion and Ranking of Opinions: Towards a Social Semantics for the Semantic Web. Procs. of the Second Workshop on Uncertainty Reasoning for the Semantic Web (URSW'06), 2006.
[41] M. Nickles, G. Weiss. A Framework for the Social Description of Resources in Open Environments. Procs. of the Seventh International Workshop on Cooperative Information Agents (CIA'03). LNCS Volume 2782. Springer, 2003.

[42] M. Nickles, Towards a Multiagent System for Competitive Website Ratings. Research Report FKI-243-01, Technical University Munich (TUM), 2001.
[43] M. Nickles, T. Froehner, R. Cobos Pérez, G. Weiss. Multi-Source Knowledge Bases and Ontologies with Multiple Individual and Social Viewpoints. In Procs. of The 2005 International Conference on Web Intelligence (WI'05), IEEE Computer Society Press, 2005.
[44] J. van Benthem, J. van Eijck, B. Kooi. Logics of communication and change. Information and Computation 204/11. Academic Press, 2006.
[45] R. Cobos. Mechanisms for the Crystallisation of Knowledge, a Proposal Using a Collaborative System. Ph.D. thesis. Universidad Autonoma de Madrid, 2003.
[46] I. Horrocks, P. F. Patel-Schneider. Reducing OWL entailment to Description Logic Satisfiability. Journal of Web Semantics, Vol. 1(4), 2004.
[47] A. Farquhar, A. Dappert, R. Fikes, W. Pratt. Integrating Information Sources using Context Logic. Procs. of the AAAI Spring Symposium on Information Gathering from Distributed Heterogeneous Environments, 1995.
[48] Open Ontologies and Open Knowledge Bases, `http://www.openontologies.org`
[49] J. F. Sowa. Semantic Networks. In S. C. Shapiro (Ed.), Encyclopedia of Artificial Intelligence, Wiley, 1987.
[50] J. Carroll, Ch. Bizer, P. Hayes, P. Stickler. Named Graphs, Provenance and Trust. Procs. of the 14th International World Wide Web Conference, 2005.
[51] P. Mika. Ontologies Are Us: A Unified Model of Social Networks and Semantics. Procs. of the 4th International Semantic Web Conference, 2005.
[52] A. Dragoni, P. Giorgini. Revisining Beliefs Received from Multiple Sources. In Frontiers in Belief Revision, H. Roth and M. Williams, Eds. Kluwer Academic Publisher, 431–444, 2001.
[53] FIPA-ACL, `http://www.fipa.org/repository/aclspecs.html`