

# An Approach to Description Logic with Support for Propositional Attitudes and Belief Fusion<sup>\*</sup>

Matthias Nickles<sup>1</sup> and Ruth Cobos<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Bath,  
Bath, BA2 7AY, United Kingdom, [m.l.nickles@cs.bath.ac.uk](mailto:m.l.nickles@cs.bath.ac.uk)

<sup>2</sup> Departamento de Ingeniería Informática, Universidad Autónoma de Madrid,  
28049 - Madrid, Spain, [ruth.cobos@uam.es](mailto:ruth.cobos@uam.es)

**Abstract.** In the (Semantic) Web, the existence or producibility of certain, consensually agreed or authoritative knowledge cannot be assumed, and criteria to judge the trustability and reputation of knowledge sources may not be given. These issues give rise to formalizations of web information which factor in heterogeneous and possibly inconsistent assertions and intentions, and make such heterogeneity *explicit* and manageable for reasoning mechanisms. Such approaches can provide valuable meta-knowledge in contemporary application fields, like open or distributed ontologies, social software, ranking and recommender systems, and domains with a high amount of controversies, such as politics and culture. As an approach to this, we introduce a lean formalism for the Semantic Web which allows for the explicit representation of controversial individual and group opinions and goals by means of so-called *social contexts*, and optionally for the probabilistic belief merging of uncertain or conflicting statements.

Doing so, our approach generalizes concepts such as provenance annotation and voting in the context of ontologies and other kinds of Semantic Web knowledge.

*Keywords:* Semantic Web, OWL, Knowledge Integration, Context Logic, Voting, Provenance Annotation

## 1 Introduction

Information found in open environments like the web can usually not be treated as objective, certain knowledge directly, and also not as truthful beliefs (due to the mental opaqueness of the autonomous information sources). Only a few approaches to the semantic modeling of what could be called subjective opinions, ostensible beliefs or “public assertions”, which are neither truthful beliefs nor objective knowledge, exist so far [11, 12]. In contrast, most prevalent formal approaches to knowledge representation and reasoning for the Web handle logical

---

<sup>\*</sup> This work is a revised and extended version of a paper published in the Proceedings of the Second Workshop on Uncertainty Reasoning for the Semantic Web (URSW-06), 2006.

inconsistencies and information source controversies mostly as something which should be avoided or filtered.

Against that, we argue that making (meta-)knowledge about the social, heterogeneous and controversial nature of web information *explicit* can be extremely useful - e.g., in order to gain a picture of the opinion landscape in controversial domains such as politics, for subsequent decision making and conflict resolution, for the acquisition and ranking of information from multiple, possibly dissent sources, and not at last for tasks like the learning whom (not) to trust. Such knowledge is especially crucial in domains with a strong viewpoint competition and difficult or impossible consensus finding like politics, product assessment and culture, and in current and forthcoming Semantic Web applications which support explicitly or implicitly people interaction, like (semantic) blogging, discussion forums, collaborative tagging and folksonomies, and in social computing in general. Approaching this issue, this work presents a lean approach to the formal representation of semantical heterogeneity by means of *social contexts* and the probabilistic weighting and fusion of inconsistent opinions.

The remainder of this paper is structured as follows: the following section defines the two most important concepts underlying our approach, namely social contexts and *social ontologies*. Section 3 introduces a formal, C-OWL based framework for the modeling of social contexts, and Section 4 shows how the formerly presented formal framework can be extended in order to allow for the fusion and probabilistic weighting of competing statements. Section 5 concludes with a discussion of related works.

## 2 Integration of Divergent Viewpoints and Intentions Using Social Contexts

In the following, we describe the main concepts underlying our approach. First we introduce a so-called *social ontology* of social entities and structures. This ontology is then used to obtain a certain type of logical contexts (called *social contexts*) which allow for the modularization of (ordinary) ontologies w.r.t. the addressee-dependent propositional attitudes of actors or organizations towards the axioms and facts in these ontologies.

A more in-depth exploration of these concepts can be found in [22].

### 2.1 Social Ontologies

Technically, our approach is based on implementing an interrelationship of a *social ontology* 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 (*opinions*) on the other hand. Instances of the social ontology represent the knowledge sources which contribute these opinions. Special terms which are assembled using names from

the social ontology then identify social contexts for the contextualization and optionally the fusion of semantically heterogeneous statements. The social ontology can thus be seen as a meta-ontology which is used to provide elements which are used to annotate facts and axioms of other ontologies (the ontologies which contains the opinions). The contextualization itself (independent of the social ontology) corresponds to the context-driven partitioning of a knowledge space, analogously to the approach presented in [2, 4].

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 interacting agents and social groups, such as “Author”, “Publisher” or “Reader”, or, most basic, “Actor”. We believe that information sources shall be seen as active, autonomous and - most important - communicating (i.e., *social*) actors, as well as the recipients of the information. A mere conceptualization of the (Semantic) Web as a kind of huge distributed document or knowledge base containing passive information fragments would be highly inadequate [23]. We see the Semantic Web rather as a place where actively pursued opinions and intentions will either compete against or strengthen each other interactively [24]. This viewpoint is independent from the concrete ways such interaction is technically performed (directly or indirectly, synchronously or asynchronously...).

The following example ontology fragment will do for the purpose of this work:

**Definition 1: Social ontology *SO* (example)**

*Actor(person<sub>1</sub>), Actor(person<sub>2</sub>), Actor(person<sub>3</sub>)*  
 ...  
*Communication(com<sub>1</sub>), Communication(com<sub>2</sub>), Communication(com<sub>3</sub>),*  
*Communication(com<sub>4</sub>)*  
 ...  
*Source(com<sub>1</sub>, person<sub>2</sub>), Addressee(com<sub>1</sub>, person<sub>3</sub>)*  
*Source(com<sub>2</sub>, person<sub>1</sub>), Addressee(com<sub>2</sub>, person<sub>2</sub>)*  
 ...  
*Content(com<sub>1</sub>, “a\_reified\_statement”)*  
  
*DegreeOfCertainty(com<sub>1</sub>, 0.75)*  
*DegreeOfCertainty(com<sub>7</sub>, 0)*  
 ..  
*SocialGroup(group<sub>1</sub>), SocialGroup(group<sub>2</sub>)*  
 ...  
*hasMember(group<sub>1</sub>, person<sub>1</sub>), hasMember(group<sub>2</sub>, person<sub>1</sub>)*  
*Actor(group<sub>1</sub>)*  
 ...  
*Actor(organization<sub>1</sub>)*  
 ...

*Source(com<sub>4</sub>, group<sub>1</sub>), Addressee(com<sub>4</sub>, organization<sub>1</sub>)*  
 ...  
*CA(assertion), CA(publicBelief), CA(publicIntention)*  
*Attitude(com<sub>1</sub>, publicBelief), Attitude(com<sub>2</sub>, assertion),*  
*Attitude(com<sub>3</sub>, publicIntention)*  
 ...  
*Aggregation(fusedPublicBelief)*

At this, *Actor* is the category of the participating actors, whereby these can be any kind of information sources or addressees, like persons, organizations, documents, web services, as well as the holder of a so-called public intention or goal (cf. below). *Communication* is the category of elementary communication acts, described by the properties *Source*, *Addressee*, *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.

Information sources and addresses 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 source. 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 judgment aggregation [27]. But we think that no single particular way of group opinion settlement should be statically fixed. Instead, we will later introduce a special aggregation operator (informally denoted as *fusedPublicBelief* 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.

At a first glance, it might seem that on the Semantic Web, the addressee of information is always the general public and thus a fine grained modeling of communication addressees would not be required. This is untrue at least for two reasons: firstly, Semantic Web technologies are also useful in environments where the set of recipients of some information is limited, such as in closed web communities. Secondly, even if some information is in principle visible to everybody, it is nevertheless usually targeted at some specific audience (although it might be difficult to obtain this kind of meta knowledge).

In this work we support the modeling of three public propositional attitudes: *assertion*, *publicBelief*, and *publicIntention*, all subsumed in the ontology under *CA* (“Communication Attitude”).

*assertion* means that a certain statement is ostensibly believed *and* that the speaker (author) has the ostensible intention to make the addressee(-s) adopt the same attitude towards the respective statement also (e.g., “This product is the best-buy!”). This corresponds more or less to the communication act semantics which we have introduced in [11, 5, 12], and to Grice’s conceptualization of speech acts as communications of intentions. *publicBelief* means here more or less the same as *assertion*, but in distinction from the latter *publicBelief* is a passive stance and does not necessarily comprise the person’s intention to make the addressees approve the respective statement but merely that a person agrees with some statement (but note that it is not possible to communicate an information  $p$  without the implicit assertion that  $p$  is indeed an information...). We could likewise have called *publicBelief* *belief* instead, but avoid the latter in order to be able to distinguish between mental (truthful) beliefs and opinions. Both *publicBelief* and *assertions* are sometimes called “opinions” in this work. The pragmatic status of *publicBelief*, being a kind of “weak assertion”, is somewhat unclear and mainly introduced for compatibility reasons w.r.t. [13], and we believe that *assertion* is sufficient to model most cases of information dissemination on the (Semantic) Web.

*publicIntention* finally is the communication attitude of ostensibly intending that a statement shall become true (i.e., an intention or goal of the actor to change the world appropriately). The attitude of *requesting* something from another actor is a subtype of *publicIntention*. As a simplification, we consider the attitude of *denial* as identical with the positive attitude towards the negation of the denied statement. This would perhaps be 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.

*assertion*, *publicBelief* and *publicIntention* are no propositional attitudes in the usual mentalistic sense but *public* propositional attitudes, 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 - stances taken on statements in the course of social interaction. As a consequence, they can not be treated like their mental counterparts. E.g., an actor might hold the opinion  $\phi$  towards addressee one and at the same time  $\neg\phi$  informing addressee two (while believing neither  $\phi$  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. Even more, *all* information on the web is “just” opinion, simply due to the absence of a commonly accepted truth assessment authority.

*fusedPublicBelief* will be described later. It is used in place of communication attitudes, but it actually stands for the merging of opinions by some observer.

## 2.2 Social Contexts

Contexts (aka *microtheories*) have been widely used in AI since the early nineties, originally intended by McCarthy as a replacement of modal logic. [1, 2] propose a context operator  $ist(context, statement)$  which denotes that *statement* is true (“ist”) within *context*. Building upon general approaches to contexts (specifically [2, 4]), and earlier works on social reification [24], 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: firstly, 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 which express the particular attitudes web authors have towards statements. Although “propositional attitude” is traditionally a psychological concept, we use this term here for attitudes reported communicatively.

Secondly, 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 among statements [27].

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.

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 logical context which allows to model the communicated attitudes associated with information on the web.

The idea is to use parts of the descriptions of individual elementary communications as defined in *SO* as identifiers of contexts. That is, we maintain two ontologies: first *SO*, and second a dynamic context ontology, with context identifiers created from certain instances of *SO*. But for some applications, it will be sufficient to actually create and maintain only the latter ontology, whereas *SO* is given only implicitly in form of the context identifiers.

## 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 (see the following 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 *Source*, *Addressee* 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:

$$\begin{array}{l} attitude \\ source \longrightarrow adresse \end{array}$$

This term is obtained from a  $SO$  fragment

$$Source(com, source), Addressee(com, adresse), Attitude(com, attitude)$$

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

$$\begin{array}{l} attitude \\ source_1, \dots, source_n \longrightarrow adresse_1, \dots, adresse_n \end{array}$$

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

As an abbreviation, we define  $\begin{array}{l} attitude \\ source_1, \dots, source_n \end{array} =_{source_1, \dots, source_n} \begin{array}{l} attitude \\ source_1, \dots, source_n \end{array} \longrightarrow Actor$ , 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 public web page. If the sources, addressees and the attitude are unspecified, for both sources and addressees the extension of  $Actor$  is assumed, and *publicBelief* as the attitude.

At this, it is important to see that - like in real life - a certain source can hold mutually inconsistent attitudes even towards different members or subgroups of  $Actor$  at the same time (but not towards the same addressee).

*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 *Source*. The syntax of fusion context identifiers is

*fusedPublicBelief*  
 $source_1, \dots, source_n \longrightarrow addressee$

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

*fusedPublicBelief*  
 $source_1, \dots, source_n \longrightarrow addressee_1, \dots, addressee_n$

A question in this regard is how the information required in order to create social contexts (i.e., information source, addressee(-s), attitude) can be obtained. Basically, the answer is analogous to the answer to the question where other Semantic Web data such as RDF or OWL documents shall come from: they need to be manually created or automatically generated. Other somewhat applicable analogies are the process of quotation, referencing, the provision of *named graphs* [20] and provenance annotation (but note that named graphs and all kinds of annotation are significantly weaker concepts compared to logical contexts). For example, authors could provide social contexts with their own statements on the web. Other knowledge workers or ontology creators could use social contexts in order to integrate statements provided by different people. As long as the authors of these statements are known (or at least URIs), at least the most simple kinds of social context identifiers can be easily generated. In contrast to techniques such as ontology mapping or trust assessment, social contextualization, if seen as a technical approach to quotation, is a simpler means to create correct and mutually consistent statements from inconsistent or dubious source statements (but of course it might require the recursive application of social contextualization...). Although social contexts only “wrap” the general problem of limited trustability on the web, they can be useful in order to integrate information on the fly, especially if no trust information is available. This functionality is shared with *RDF reification*, but the use of the long established context logic and its Semantic Web versions such as C-OWL appears to be a cleaner and better researched approach.

### 3 A Description Logic with Support for Social Contexts

We introduce now a formal language based on C-OWL [4] for the representation of ontologies with social contexts.

We settle on the *SHOIN(D)* description logic (over data types *D*), because ontology entailment in the current quasi-standard OWL DL can be reduced to *SHOIN(D)* knowledge base satisfiability [16]. 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), rule languages, or first-order logic.

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

The context-free grammar of  $\mathcal{SHOIN}(D)$  concepts  $C$  is as follows. Please find detailed information about the syntax and semantics of  $\mathcal{SHOIN}(D)$  in [16, 17].

$$\begin{aligned}
C &\rightarrow A \mid \neg C \mid C_1 \sqcap C_2 \mid C_1 \sqcup C_2 \mid \exists R.C \mid \forall R.C \\
&\quad \mid \geq nS \mid \leq nS \mid \{a_1, \dots, a_n\} \mid \geq nT \mid \leq nT \mid \exists T_1, \dots, T_n.D \mid \forall T_1, \dots, T_n.D \\
D &\rightarrow d \mid \{c_1, \dots, c_n\}.
\end{aligned}$$

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* [16], 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* (or *knowledge base*) is then a finite, non-empty set of TBox axioms and ABox axioms (“facts”)  $C_1 \sqsubseteq C_2$  (inclusion of concepts),  $\text{Trans}(R)$  (transitivity),  $R_1 \sqsubseteq R_2$ ,  $T_1 \sqsubseteq T_2$  (role inclusion for abstract respectively concrete roles),  $C(a)$  (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 [16].

**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 [4]. While the syntax of SOC-OWL can be seen as a defined subset of the syntax of C-OWL, and SOC-OWL essentially shares with C-OWL the interpretation of concepts, individuals and roles, SOC-OWL satisfiability is constrained by meta-axioms (cf. 3.2) which go beyond C-OWL and put SOC-OWL somewhat close to BDI-style modal logics [11]. Essentially, SOC-OWL adds a kind of “S-Box” (“social box”, i.e., social contexts) to a formal ontology language. In contrast to the mere *annotation* of axioms or facts with provenance information or other meta data, these contexts provide separate (but bridgeable) spheres of reasoning.

In the next section, the language P-SOC-OWL will be introduced, which also allows for uncertainty reasoning.

A SOC-OWL ontology 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  $\mathcal{SHOIN}(D)$  TBox and ABox axioms,  $AF^i$  being such axioms but with concepts, individuals and roles directly indexed with social contexts (i.e.,  $AF^i = \{(id_i, C_h) \sqsubseteq (id_j, C_k), (id_i, a_h) = (id_j, a_k), \dots : id_i, id_j \in Id\}$ ), and  $B$  being a set of *bridge rules* (see 3.1). 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  $(id, s)$  are called *inner statements* which are said to “be *true* (or *intended* in case of *publicIntention*) within the respective context”.

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

$InfluentialPainter(FrankFrazetta)$   $InfluentialPainter \sqsubseteq Painter$   
 $\overset{assertion}{tina \rightarrow tim, tom} InnovativeArtist(FrankFrazetta)$   
 $\overset{assertion}{tim, tom \rightarrow tina} (\neg InnovativeArtist)(FrankFrazetta)$   
 $\overset{assertion}{tim, tom \rightarrow tina} TrashArtist(FrankFrazetta)$   
 $\overset{assertion}{tom} (\neg InnovativeArtist)(FrankFrazetta)$

$ControversialWikipediaArticle \sqsubseteq WikipediaArticle$   
 $NeutralWikipediaArticle \sqsubseteq WikipediaArticle$   
 $\overset{assertion}{tina} WikipediaArticle \sqsubseteq NeutralWikipediaArticle$   
 $ControversialWikipediaArticle(ArticleAboutFrankFrazetta)$   
 $\overset{assertion}{tim, tom \rightarrow tina} (\neg NeutralWikipediaArticle)(ArticleAboutFrankFrazetta)$

This SOC-OWL ontology (modeling as a whole a sort of neutral point of view, like taken by an ideal Wikipedia article) expresses that the information sources Tim and Tom hold the opinion towards Tina that the painter Frank Frazetta is not an innovative artist but a trash artist, while Tina does allegedly believe that the opposite is true. But there is consensus of the whole group that Frazetta is an influential painter. Furthermore, Tina believes that all Wikipedia articles present a neutral point of view.

Notice that without explicit further constraints, bridge rules or meta-axioms, different social contexts are logically fully separated. Also, using only the above ontology it could *not* be inferred that  $\overset{publicBelief}{tina \rightarrow tim} InfluentialPainter(FrankFrazetta)$ , because

$InfluentialPainter(FrankFrazetta)$  as an abbreviation of

$\overset{publicBelief}{tina, tim, tom \rightarrow tina, tim, tom} InfluentialPainter(FrankFrazetta)$

in the example above is uttered/addressed *exactly* by/to the social group of all participants and not by/to any subgroup or individual. Consensus is always bound to a concrete social group and does not necessarily propagate to social subgroups. 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 inner statements which indicate that a certain actor is insincere (as it would be the case with  $\overset{assertion}{tina \rightarrow tim} (\neg C)(x)$  and  $\overset{assertion}{tina \rightarrow tom} C(x)$  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 ob-

tain such SOC-OWL knowledge.

Obviously, each SOC-OWL statement  $(contextId, statement)$  corresponds to the “classic” [1, 2] context logic 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 the need for a considerably more complicated semantics. As a further serious restriction compared to real context logic, it is not possible to relate contextualized statements freely with logical connectives like in  $ist(c_1, s_x) \vee ist(c_2, s_y) \rightarrow ist(c_1, s_z)$ .

Instead of these features, 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 [4]. 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

$\overset{assertion}{tina \rightarrow tim, tom}(\neg TrashArtist)(FrankFrazetta)$  implies (intuitively)

$\overset{publicIntention}{tina}(\overset{publicBelief}{tim, tom \rightarrow tina}(\neg TrashArtist)(FrankFrazetta))$ , 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  $I_{id} = \langle \Delta^{I_{id}}, (\cdot)^{I_{id}} \rangle$ ,  $id \in Id$ .  $e_{i,j} \subseteq \Delta^{I_i} \times \Delta^{I_j}$  is a relation of two *local domains*  $\Delta^{I_{id}}$  ( $e_{i,j}$  is required for the definition of bridge rules in  $B$  (Definition 4) as explained later in 3.1).  $(\cdot)^{I_{id}}$  maps individuals, concepts and roles to elements (respectively subsets or the products thereof) of the domain  $\Delta^{I_{id}}$ .

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  $c_{id}$  [4]. 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:

$$C^{I_{id}} = \text{any subset of } \Delta^{I_{id}} \text{ for } C \in C_{id}$$

$$(C_1 \sqcap C_2)^{I_{id}} = C_1^{I_{id}} \cap C_2^{I_{id}} \text{ for } C_1, C_2 \in C_{id}$$

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

### Satisfiability and decidability

Given a SOC-OWL interpretation  $I$ ,  $I$  is said to *satisfy* a (contextualized) statement  $\phi$  ( $I \models \phi$ ) if there exists an  $id \in Id$  such that  $I_{id} \models \phi$ , with  $I_{id} \in I$ . A SOC-OWL ontology is then said to be “satisfied” if  $I$  satisfies each statement within the ontology (or statement set) and the ontology observes the meta-axioms listed below.

$I_{id} \models (id, C_1 \sqsubseteq C_2)$  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.

With this configuration, the inherited semantics and decidability of  $\mathcal{SHOIN}(D)$  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.

### 3.1 Bridge Rules and Cross-Context Mappings

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

$(context_i, x) \underline{\equiv} (context_j, y)$  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.,  $(\overset{assertion}{\underset{tina}{}}, FrankFrazetta) \underline{\equiv} (\overset{assertion}{\underset{tim,tom}{}}, FrankFrazetta)$  the interpretations of the “two Frank Frazettas” would abstractly refer to the same object. Analogously,  $\underline{\sqsubseteq}$ , and  $\perp$ , 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:  $I \models (context_i, x) \underline{\equiv} (context_j, y)$  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 [4]. Also, reasoning in the presence of bridge rules follows that with C-OWL.

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).

### 3.2 Meta-Axioms

We state now some constraints, which will later be extended w.r.t. a different formal language with meta-axiom (PMA5). All so-called meta-axioms are in fact either entailment rules (which could not be formulated using SOC-OWL axiom schemes because the language is not expressive enough), or they put constraints regarding its integrity on an ontology which is sliced into social contexts. Although a practical reasoner could possibly take advantage of the latter kind of meta axioms (since these exclude certain constellations such as inconsistent contexts), they don't demand special reasoning procedures.

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

$\overset{assertion}{s_1, \dots, s_n \rightarrow a_1, \dots, a_m} \varphi \rightarrow (\overset{publicIntention}{s_1, \dots, s_n \rightarrow a_1, \dots, a_m} (\overset{publicBelief}{a_1, \dots, a_m \rightarrow s_1, \dots, s_n} \varphi))$ . But this "strong" and problematic nesting is not possible in our language.

The next meta-axiom simply demands that assertions include the attitude of informing the addressee:

(MA1)  $\overset{assertion}{s_1, \dots, s_n \rightarrow a_1, \dots, a_m} \varphi \rightarrow \overset{publicBelief}{s_1, \dots, s_n \rightarrow a_1, \dots, a_m} \varphi$

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 [5, 11]). 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, is satisfiable (ensuring the consistency of one's opinions).

Furthermore, we demand - in accordance with many BDI-style logics - that the approval/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  $(\overset{publicIntention}{s_1, \dots, s_n \rightarrow a_1, \dots, a_n}, a)$  is part of an SOC-OWL ontology  $o$ , no axiom/fact  $(\overset{publicBelief}{s_1, \dots, s_n \rightarrow a_1, \dots, a_n}, b), b \vdash a$ , is part of  $o$  (analogously for assertions).

The following constraints are *not* demanded, but could be helpful in application domains where mutual opinion consistency of subgroups is desired (we use  $\bigwedge$  to abbreviate a set of SOC-OWL statements).

(MAx1)  $(\overset{attitude}{s_1, \dots, s_n \rightarrow a_1, \dots, a_n} \varphi) \leftrightarrow \bigwedge_{s \in 2^{\{s_1, \dots, s_n\}} - \{\emptyset\}} \overset{attitude}{s \rightarrow a_1, \dots, a_n} \varphi$

(MAx2)  $(\overset{attitude}{s_1, \dots, s_n \rightarrow a_1, \dots, a_n} \varphi) \leftrightarrow \bigwedge_{a \in 2^{\{a_1, \dots, a_n\}} - \{\emptyset\}} \overset{attitude}{s_1, \dots, s_n \rightarrow a} \varphi$

But we can safely aggregate seemingly consented information in a separated fusion context:

(MA4)  $\bigwedge_{s \in \{s_1, \dots, s_n\}} (I_{\substack{\text{publicBelief} \\ s \rightarrow a_1, \dots, a_n}} \models \varphi) \rightarrow (I_{\substack{\text{fusedPublicBelief} \\ s_1, \dots, s_n \rightarrow a_1, \dots, a_n}} \models \varphi)$  (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).

## 4 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). 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 [15].

In [10], the probabilistic extension  $P\text{-}\mathcal{SHOQ}(D)$  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 [16]. [10] shows that reasoning with  $P\text{-}\mathcal{SHOQ}(D)$  is - maybe surprisingly - decidable. Instead of  $P\text{-}\mathcal{SHOQ}(D)$ , other probabilistic approaches to Semantic Web and ontology languages could likely also be used as a basis for our approach, e.g., [7].  $P\text{-}\mathcal{SHOQ}(D)$  is now used to define a probabilistic variant of SOC-OWL.

### Definition 6: P-SOC-OWL

A  $P\text{-SOC-OWL}$  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 given as a family of  $P\text{-}\mathcal{SHOQ}(D)$  interpretations, each interpretation corresponding to a certain 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  $P\text{-}\mathcal{SHOQ}(D)$  and a social context with identifier  $id$ .

$\mu_{id} : \Delta^{Id} \rightarrow [0, 1]$  is a subjective probability function, and the  $\Delta^{Id}$  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, since “within” this context, no bridge rules or meta-axioms need to be observed and thus P-SOC-OWL behaves in this case just like  $P\text{-}\mathcal{SHOQ}(D)$ .

Individualistically assigned probabilities are constrained by the axioms of probability.

Example:

$[0.5, 0.8]: \overset{\text{assertion}}{\underset{\text{tim,tom}}{\longrightarrow \text{tina}}} \text{TrashArtist}(\text{FrankFrazetta})$   
 $0.7: \overset{\text{assertion}}{\underset{\text{tina}}{\longrightarrow}} \text{InnovativeArtist}(\text{FrankFrazetta})$   
 $0.9: \overset{\text{assertion}}{\underset{\text{tim}}{\longrightarrow}} \text{InnovativeArtist}(\text{FrankFrazetta})$

This P-SOC-OWL ontology expresses inter alia that Tim and Tom (as a group, but not necessarily separately) hold the opinion that with some probability in  $[0.5, 0.8]$ , Frank Frazetta is a trash artist, while Tina does (publicly) believe he is an innovative artist 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')  $(\bigwedge_{s_i \in \{s_1, \dots, s_n\}} (Pr_{s_i \rightarrow \text{addressees}}^{\text{publicBelief}} \models \varphi[p_i, p_i])) \rightarrow (Pr_{s_1, \dots, s_n \rightarrow \text{addressees}}^{\text{publicBelief}} \models \varphi[p, p])$   
 with  $p = \text{pool}^{\text{poolingType}}((p_1, \dots, p_n), \text{extraKnowledge})$ . At this,  $Pr_{id} \models \varphi[l, u]$  at-  
 tests  
 $\varphi$  a probability within  $[l, u]$  in context  $id$ , and  $\text{extraKnowledge}$  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 in-  
 consistencies) in case the derived probability  $p$  is different than a probability  
 assigned explicitly by this group of people - a group of agents is free to assign  
*any* truth value or probability to any statement, using *any* social choice pro-  
 cedure. A simple workaround is to use a new kind of context with aggregating  
 “attitude” *fusedPublicBelief*, which is actually no speaker attitude of course, but  
 a belief merging operator 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)  $(\bigwedge_{s_i \in \{s_1, \dots, s_n\}} (Pr_{s_i \rightarrow \text{addressees}}^{\text{publicBelief}} \models \varphi[p_i, p_i])) \rightarrow (Pr_{s_1, \dots, s_n \rightarrow \text{addressees}}^{\text{fusedPublicBelief}} \models \varphi[p, p])$  (remainder as PMA5').

As for  $\text{pool}^{\text{poolingType}}$ , there are several possibilities: in the most simple case  
 of “democratic” Bayesian aggregation given the absence of any opinion leader  
 or so-called “supra-Bayesian” [15], we define  $\text{pool}^{\text{avg}}((p_1, \dots, p_n), \emptyset) = \frac{\sum p_i}{n}$ , i.e.,  
 $\text{pool}^{\text{avg}}$  averages over heterogeneous opinions. Using this aggregation operator,  
 we could infer the following:

0.8:  $\text{fusedPublicBelief}_{\text{tina, tim}} \text{InnovativeArtist}(\text{FrankFrazetta})$ .

Social aggregation operators are traditionally studied in the field of *Bayesian  
 belief aggregation* [15, 3].

The most common fusion operator extends  $\text{pool}^{\text{avg}}$  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):

$\text{pool}^{\text{LinOP}}((p_1, \dots, p_n), (\text{weight}_1, \dots, \text{weight}_n)) = \sum \text{weight}_i p_i$ , with  $\sum \text{weight}_i =$   
 1. Also quite often, a geometric mean is used:

$\text{pool}^{\text{LogOP}}((p_1, \dots, p_n), (\text{weight}_1, \dots, \text{weight}_n)) = \kappa \prod_{i=1}^n p_i^{\text{weight}_i}$  ( $\kappa$  for normal-  
 ization).

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  $\text{weight}_i$  of  $\text{pool}^{\text{LinOP}}$  and  $\text{pool}^{\text{LogOP}}$ , 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 (TBox as well as ABox), so that inner statements can be directly ranked regard their degree of assumed social acceptance. The following is an example for how such a top-k list of socially preferred statements looks like.

0.8:  $\underset{\text{voters}}{\text{fusedPublicBelief}} \text{statement}_1$  (highest social rating)  
 $[0.5, 0.8]: \underset{\text{voters}}{\text{fusedPublicBelief}} \text{statement}_2$   
 ...  
 0.2:  $\underset{\text{voters}}{\text{fusedPublicBelief}} \text{statement}_3$  (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.

What could also be generated quite easily are rankings w.r.t. of the degrees of certainty assigned to the same statement by different voters or groups of voters:

0.8:  $\underset{\text{actor}_1}{\text{publicBelief}} \text{statement}_1$   
 $[0.5, 0.8]: \underset{\text{group}_3}{\text{publicBelief}} \text{statement}_1$   
 0.4:  $\underset{\text{actor}_1, \text{actor}_4}{\text{fusedPublicBelief}} \text{statement}_1$   
 ...  
 0.1:  $\underset{\text{actor}_2}{\text{publicBelief}} \text{statement}_1$

## 5 Related Works and Conclusion

The goal of this work is to provide a social semantics of possibly contradictory assertions on the web, i.e., to state their amount of social support, their communicative emergence and dissemination, and the consensus or dissent they give rise to. Doing so, we settle on the “opinion level” where neither true 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 heterogeneous information exist initially. This is both in contrast to the traditional aim of information integration and evolution for the determination of some consistent, reliable “truth” obtained from contributions of multiple sources as in traditional

*multiagent belief representation and revision* (e.g., [21] - although this direction has still much in common with ours) 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 [9], 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* [14]). The assignment of provenance information is mostly based on *tagging* and *punning* techniques, or makes use of the semantically problematic reification facility found in RDF. Meta knowledge modeling and reification techniques for the purpose of adding certain “slots” for provenance and statement identification data, and other useful meta information to Semantic Web languages can be found in [20, 25, 25]. These approaches, with *named graphs* [20] being currently the most popular representative, leave the original semantics of the underlying language more or less untouched and “merely” annotate traditional language constructs with some optional meta-information. In contrast, our approach aims at a truly social semantics and language.

[6] provides an approach to the grouping of RDF statements using contexts (including contexts for provenance and speech act performatives). Another related approach focusing on contexts including contexts for the aggregation of RDF graphs, was presented in [2], and [4] provides a general formal account of contexts for OWL ontologies. Independently from web-related approaches, contexts have been widely used for the modeling of distributed knowledge and federated databases, see, e.g., [18, 19].

To further explore and work out the new “social” perspective on uncertain information on the web modeled using contexts certainly constitutes a long-term scientific and practical endeavor of considerable complexity, with this work hopefully being a useful starting point.

**Acknowledgements:** This work was partially funded by the German National Research Foundation DFG (Br609/13-1, research project “Open Ontologies and Open Knowledge Bases”) and by the Spanish National Plan of R+D, project no. TSI2005-08225-C07-06.

## References

1. J. L. McCarthy. Notes on formalizing context. In IJCAI, pages 555-562, 1993.
2. R. V. Guha, R. McCool, R. Fikes. Contexts for the Semantic Web. Procs. of the Third International Semantic Web Conference (ISWC-04), 2004.

3. 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.
4. P. Bouquet, F. Giunchiglia, F. van Harmelen, L. Serafini, and H. Stuckenschmidt. C-OWL: Contextualizing Ontologies, Second International Semantic Web Conference (ISWC-03), LNCS vol. 2870, Springer Verlag, 2003.
5. M. Nickles, F. Fischer, G. Weiss. Communication Attitudes: A Formal Approach to Ostensible Intentions, and Individual and Group Opinions. Procs. of the 3rd Intl. Workshop on Logic and Communication in Multiagent Systems (LCMAS 2005), 2005.
6. G. Klyne. Contexts for RDF Information Modelling. <http://www.ninebynine.org/RDFNotes/RDFContexts.html>, 2000.
7. P. Costa, K. B. Laskey, K. J. Laskey. PR-OWL: A Bayesian Framework for the Semantic Web. In Procs. First Workshop on Uncertainty Reasoning for the Semantic Web (URSW-05), 2005.
8. T. Froehner, M. Nickles, G. Weiss. Towards Modeling the Social Layer of Emergent Knowledge Using Open Ontologies. In Proceedings of The ECAI-04 Workshop on Agent-Mediated Knowledge Management (AMKM-04), 2004.
9. J. Ullmann. Information Integration Using Logical Views. Proc. 6th Int'l Conference on Database Theory. Springer, 1997.
10. R. Giugno, Th. Lukasiewicz. P-SHOQ(d): A Probabilistic Extension of SHOQ(d) for Probabilistic Ontologies in the Semantic Web. In JELIA '02: Procs. of the European Conference on Logics in Artificial Intelligence. Springer, 2002.
11. F. Fischer, M. Nickles. Computational Opinions. Procs. of the 17th European Conference on Artificial Intelligence (ECAI'06), 2006. To appear.
12. 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. To appear.
13. M. Nickles. Modeling Social Attitudes on the Web. In *Proceedings of the 5th International Semantic Web Conference (ISWC)*, 2006.
14. R. Cobos. Mechanisms for the Crystallisation of Knowledge, a Proposal Using a Collaborative System. Ph.D. thesis. Universidad Autonoma de Madrid, 2003.
15. R. M. Cooke. Experts in Uncertainty: Opinion and Subjective Probability in Science. Oxford University Press, 1991.
16. I. Horrocks, P. F. Patel-Schneider. Reducing OWL entailment to Description Logic Satisfiability. *Journal of Web Semantics*, Vol. 1(4), 2004.
17. P. Haase, B. Motik. A Mapping System for the Integration of OWL-DL Ontologies. In Proceedings of the First International Workshop on Interoperability of Heterogeneous Information Systems. ACM Press, 2005.
18. M. Bonifacio, P. Bouquet, R. Cuel. Knowledge Nodes: The Building Blocks of a Distributed Approach to Knowledge Management. *Journal for Universal Computer Science*, Vol. 8/6, 2002.
19. A. Farquhar, A. Dappert, R. Fikes, W. Pratt. Integrating Information Sources using Context Logic. In Procs. of the AAAI Spring Symposium on Information Gathering from Distributed Heterogeneous Environments, 1995.
20. J. Carroll, Ch. Bizer, P. Hayes, P. Stickler. Named Graphs, Provenance and Trust. In Procs. of the 14th International World Wide Web Conference, 2005.
21. 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.

22. M. Nickles. Social Acquisition of Ontologies from Communication Processes. *Applied Ontology*. IOS Press. Volume 2, Number 3-4, 2007.
23. M. P. Singh. The Pragmatic Web: Preliminary thoughts. In *Proceedings of the NSF-OntoWeb Workshop on Database and Information Systems Research for Semantic Web and Enterprises*, 82–90, 2002.
24. M. Nickles, G. Weiss. A Framework for the Social Description of Resources in Open Environments. *Proc. 7th International Workshop on Cooperative Information Agents (CIA)*. Springer LNCS, 2003.
25. D. T. Tran, P. Haase, B. Motik, B. C. Grau, I. Horrocks. Metalevel Information in Ontology-Based Applications. In *Procs. of the 23th AAAI Conference on Artificial Intelligence (AAAI'08)*, 2008.
26. B. Schueler, S. Sizov, S. Staab, D. T. Tran. Querying for Meta Knowledge. In *Procs. of the 17th International Conference on the World Wide Web (WWW'08)*, 2008.
27. G. Pigozzi, S. Hartmann. Judgment aggregation and the problem of truth-tracking. *Procs. of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2007)*, 2007.