

Modeling Social Attitudes on the Web

Matthias Nickles

AI/Cognition Group, Department of Computer Science,
Technical University of Munich
D-85748 Garching b. München, Germany
`nickles@cs.tum.edu`

Abstract. This paper argues that in order to allow for the representation, comparison and assessment of possibly controversial or uncertain information on the web, the semantic web effort requires capabilities for the social reasoning about web ontologies and other information acquired from multiple heterogeneous sources. As an approach to this, we propose formal means for the representation of possibly controversial opinions of groups and individuals, and of several other social attitudes regarding information on the web. Doing so, we integrate concepts from distributed artificial intelligence with approaches to web semantics, aiming for a *social semantics* of web content.

Keywords: Semantic Web, Description Logic, Information Integration, Agent Communication, Modal Logic.

1 Introduction

Social aspects of the web have attracted increased attention in the field of semantic web research recently. This development is certainly driven in part by the tremendously growing interest in web-based collaboration by means of *social software* (e.g., for blogging, collaborative tagging, wiki creation etc.). But the increasing interest in sociality also seems to be stemming from the general insight, that the semantic web can never become some kind of huge distributed knowledge base in the traditional sense. Instead, it will in our opinion become a more and more realistic emergent image of the current "non-semantic" web, i.e., an open environment with heterogeneous groups of information sources and users, both with different and often conflicting viewpoints and interests, and with a high amount of personal interaction. This development will likely accelerate with new interaction-oriented developments like *semantic blogging*.

Approaches like *emergent semantics* [1], *dynamic ontologies* [14] and advances in the field of information integration in general and ontology mapping and merging specifically (e.g., [13,2]) already provide strong responses to some of the challenges posed by open information environments. But nevertheless, to their major part, current approaches to web semantics are concerned with the modeling of homogenous information, or the assessment and filtering of heterogeneous information in terms of trustability and suitability. What is still widely missing

are formal means for the modeling of (knowledge-related) sociality on the web itself, especially the simultaneous and comparative representation of heterogeneous and possibly inconsistent viewpoints of multiple information sources. Such means would not only allow for a rich modeling of social (i.e., communication) structures of web information (providing meta-information useful for, e.g., a subsequent resolution of conflicts and credibility issues). They would also allow for the social reasoning about the social meaning of information contributions on a logical level (as opposed to semi-formal approaches like *social networks* or *provenance information*). In this regard, the annotation of web information with meta-information denoting their provenance already received significant attention, but provenance modeling is to its main part strongly tailored to the problem of trustability, and it usually provides "only" meta-data (i.e., identifiers of the information sources), not an integrated logical model suitable for social reasoning.

As a response to the described issues, the main contribution of this paper is a general approach to a social semantics for the web by introducing a formal framework for a "social" multi-modal description logic, as informally outlined in the next section, and formally presented in Section 3 (including the semantics and decidability results). 3.4 introduces a social semantics of web publishing acts building on the formal framework, and Section 4 demonstrates our approach by means of a case study. Section 5 concludes.

2 A Communication-Oriented Model of Web Semantics

The basic concept underlying our approach is that of the integration of ("first-level") information with ("second-level") information about the social meaning of the former. This can in principle be done by the assignment of appropriate second-level meta-data to first-level information artifacts, using various techniques (e.g., higher-order logic, modal logic, or even RDF's reification). Henceforth, this way of reifying information is called *social reification* (or *social higher-order modeling*, if a higher-order logic shall be used). In demarcation from more informal ways of annotating information with meta-data (and the problematic and semantically extremely weak reification facility of RDF), we focus in this work on the formal integration of first-level and second-level information items which correspond (as a compound pair) to formulas and expressions of logic languages, e.g., axioms and facts represented in description logic (DL). E.g., a simple form of such a socially reified statement could look like "Frank *informs*: Sheep are pink" (or "Frank informs us that sheep are pink"), in contrast to the un-reified and hopefully highly controversial first-level statement "Sheep are pink".

Supposedly, it is easier for agents (including humans) to agree on a statement which quotes the opinion of someone else, compared to the less likely agreement with that opinion content itself ("agree to disagree" probably, so to say). Thus we believe that the (selective) social reification of web statements would achieve a great deal of increased semantic consistency of the web. In addition, social

reification provides in many cases a rather safe way of semantic linkage among different information sources or documents (like different ontologies), and is thus expected to provide a means for information integration in case there exists not enough meta-knowledge (like trust) in order to decide about the alignment and merging of the information into a consistent set of un-reified axioms and facts. Although, e.g., OWL already allows for the inclusion of foreign ontologies and class descriptions, this works on a syntactical, constraint-less level, possibly leading to inconsistencies. Information integration using social reification in contrast allows to integrate foreign information in a relatively "safe" manner even if some reified first-level information is mutually inconsistent, and later integration steps can use the meta-information about the social meaning of the reified information for tasks like conflict resolution or credibility assignment.

Most important, social reification enables social reasoning about web information *within* a logic-based knowledge representation framework itself, not needing the help of "external", non-logical approaches (although we expect that a combination of such approaches - especially that of social networks - with social reification could be very fruitful).

So, the question is how such meta-knowledge about the social meaning of web information should look like in detail. In this work, we claim that - independently of technical information authorship or message-passing means - authoring information on the web is implicitly and unavoidably a *communication* performed by an autonomous, self-interested source, namely an assertive or informational speech act performed in order to express a subjective opinion (probably about some formerly published opinion). Thereby, it is not even necessary to make either the authorship or the propositional attitude towards the act content explicit. But we find a formal account to such second-order information highly useful, namely to have a means for its representation using semantic web languages.

To this end, the direct inclusion of speech act locutions as with agent communication languages would not be adequate for languages like OWL or RDF since speech acts are on a different conceptual level. Therefore, we propose a logical means for the modeling of asserted information. Intuitively, we want to express that, after uttering, the knowledge source (human, agent, web service...) is committed to his assertion. To represent such states, we have chosen the approach introduced in [7,4] and the more or less equivalent approach presented in [6,5] as a starting point for the semantics of web communication. Our approach demarcates itself strongly both from the well-known BDI agent model and from multiagent belief modeling (including dynamic epistemic models such as *public announcement logic*, which is not concerned with opinions in our sense but with the effects of announcements on beliefs [3]). At this, we distinguish sincere individual and group beliefs (which might be not visible on the web) from publicly visible *alleged* beliefs and subjective claims (the latter also to be distinguished from objective knowledge), and from public intentions.

Observe in this respect that it would not be sufficient to extend web knowledge representation languages with modal operators for belief or epistemic knowledge

modalities. The former denotes a mentalistic concept, which cannot be extended to information publishing since information sources in open environments like the web are autonomous actors with opaque beliefs and intentions. A certain publisher might assert some statement a to audience A , assert at the same time $\neg a$ to another audience B , while believing neither a nor $\neg a$. In contrast, the modality of knowing would provide epistemic introspection, but would not be particularly useful in regard to the mentioned issues stemming from opinion controversies and the absence of an authoritative "truth" in the web.

In order to model public¹ attitudes of web actors towards information (and thus assign second-order information about social meaning), we introduce the modalities *public assertion and belief* and *public intention* as communication-level pendants to the mental attitudes belief and intention, and thus lift mental attitudes to the "social stage".

Examples demonstrating the high expressivity of a DL language with social attitudes are:

- $\Box_{\{s_1, s_4, s_5\}, \{a_1, a_7\}} \exists \text{loves.} \top = \top$
(The group of web actors $\{s_1, s_4, s_5\}$ (e.g., bloggers) publicly express facing $\{a_1, a_7\}$ the opinion that everybody loves somebody)
- $\Box_{\{a_1, a_5\}} \text{Asserts}_{\text{Frank} \rightarrow \text{all}} (\exists \text{loves.} \top = \top)$
(Group $\{a_1, a_5\}$ publicly believes that Frank asserts that everybody loves someone)
- $\text{PInt}_{\text{Frank} \rightarrow \{a_1, a_5\}} (\exists \text{loves.} \top = \top)$
(Frank publicly intends towards group $\{a_1, a_5\}$ that everybody shall love someone)
- $\text{PInt}_{\text{Frank} \rightarrow \{a_1\}} (\text{Asserts}_{a_1 \rightarrow \text{all}}^{\text{expr}} (\exists \text{loves.} \top = \top))$
(Frank publicly intends that a_1 publicly announces that everybody shall love someone (while probably privately intending the opposite))
- $\Box_{\{\text{Sarah}\}} \text{Customer} = \Box_{\{\text{Frank}\}} \text{Customer} \wedge \exists \text{has.Money}$
(Sarah's customers are Frank's customers, but only those with money. Modalities for social attitudes are thus also attachable to concept names (and role names), not only to axioms.)

Social attitudes will be given an intuitive but precise formal semantics which resembles the modalities of an actor's belief and intention in many ways [9]. They are nevertheless cleanly separated from mental attitudes (as in the BDI agent model and related multiagent belief frameworks), and can thus be used together with these without any interference. An agent might, e.g., reason simultaneously about the "real" beliefs of another agent and about the information this other agent gave to the public communicatively, which is not possible using BDI. Note in this respect that our approach is to its main part settled on a different conceptual level than related fields like multi-agent belief revision and information integration: Whereas these are mainly concerned with the determination of correct, consistent and useful information, the primary purpose of social attitudes is to represent communicatory properties of and differences among semantically heterogeneous attitudes, possibly preceding their assessment in terms of trustability and reliability.

¹ We use the term "public" not necessarily in the sense of "everyone attending", but to refer to communicatively disclosed information within specific closed or open groups (including "the web public"), as well as single persons (in form of singleton groups).

What mainly distinguishes an agent's social attitude of *public* belief from her mental attitude of belief is that the former is an *ostensible* belief expressed communicatively (maybe restricted to a specific *audience*), and triggered and revised by social conditions (in our context namely web publishing and reception). Different groups and even subgroups of some groups can hold different public beliefs without causing logical inconsistencies. Public assertions in addition aim at ostensibly convincing their addressees. While public assertions and beliefs might as well reflect the true beliefs of benevolent, trustworthy information sources, this should be considered a special case for autonomous, self-interested sources in open environments. Also, uttering a public assertion doesn't necessarily mean that the agent truthfully intends to make someone adopt this public assertion as a mentalistic belief (that would be unrealistic), but as a *public* belief. Essentially this denotes that the release of information on the web is understood to be an important part as a request which asks the readers of the information to show a positive attitude regarding the released information (i.e., implicit or explicit approval).

The general provision of meaning for data found on the web in terms of social attitudes as described in the next sections is henceforth called *social web semantics* (or just *social semantics*).

In a nutshell, we see the main benefits of modeling social attitudes in the provision of means for:

- Representation of and reasoning about public opinions (with sub-types such as agreement and disagreement), as opposed to mental beliefs, and also as opposed to ("objective") knowledge. We claim that virtually *any* kind of information published on the (semantic) web *initially* falls in the category of opinion, since such information is initially (from an observer's point of view without any meta-knowledge such as about reliability) neither known as sincere subjective belief of the publisher nor knowledge. But quite surprisingly, to our knowledge no explicit, sufficiently powerful formal means for the modeling of opinions existed so far (i.e., with a required expressivity higher than BDI or other Kripke-style multi-belief frameworks).
- Representation of and reasoning about public intentions (like the intentions behind web publishing acts, e.g., the intention to make others agree with the respective claims)
- Modeling of different audiences. E.g., someone (using different nicknames) might utter inconsistent information depending on the addressees of these opinions. Such issues will likely become increasingly important with the rise of social software and the use of semantic web technology for the representation of discourses like in web blogs or web-based negotiation platforms.

All social attitudes introduced in this paper can be restricted to specific audiences (part-publics, so to say).

- A social semantics of web publishing acts (and some other internet-relevant communication acts such as *request*) in terms of social attitudes.

2.1 Related Works

Apart from the related research field of ontology integration, the storage of heterogeneous information from multiple sources also has some tradition in the fields of *data warehousing* and view-generation for distributed and enterprise database systems [15,2], whereby such approaches do not take a social or communication-oriented perspective. *Contexts* are also used for the integration of heterogeneous information [12], but contexts in this sense originate from McCarthy's *truth contexts*, as opposed to the essentially pragmatic "social contexts" implicitly used for social reification. The assignment of provenance information is mostly based on annotation, or makes use of the reification facility found in RDF, which also lacks a social semantics of course. Approaches to provenance are already very useful if it is required to specify who contributed some information artifact (which is also done with a similar intent on the basis of social networks [11]), but they do not provide a logic model of the meaning of being an opinion source. Precisely, they allow to specify that someone asserts some information, but they do not handle what asserting (requesting, denying...) actually means, in contrast to the semantics introduced in this paper.

3 A Description Logic with Social Modalities

This section presents a description logic enhanced with modalities for the previously introduced social attitudes public assertion, public belief and public intention. The latter modality is mainly introduced for the purpose of providing a web publishing semantics later in this work; since we omit the specification of cross-modality axioms in this paper, it could be safely removed from the language in case one only wants to model public assertions and beliefs.

Our language is based on the "standard" DL \mathcal{ALC} (Attributive concept description Language with Complements) [8], with modal extensions in the style of [9,10]. A further extension with additional features found in $\mathcal{SHOIN}(D)$ (the description logic equivalent to OWL-DL) is omitted here as being not relevant in the context of this work, but should be completely straightforward, since they do not affect the model-based semantics of our modalities (see below).

3.1 The Language $\mathcal{S} - \mathcal{ALC}$

Definition 1. The language $\mathcal{S} - \mathcal{ALC}$ (*Social ALC*) is defined as follows. The syntax allows both to specify terminological knowledge about concepts and roles ("TBox"), and assertional knowledge² about their instances ("ABox"). Modalities can be attached not only to formulas, but also to roles and concepts.

Atomic concepts: $C = \{C_0, C_1, \dots\}$

Atomic Roles: $R = \{R_0, R_1, \dots\}$

Individuals: $I = \{o_0, o_1, \dots\}$

² The double meaning of "assertion" in this paper is unfortunate, but we wanted to stick with the usual DL terminology.

Inductively, we define (compound) *concepts* and *roles* now. Let R be a role and C and D concepts. Then $\Box_i R$, $\Box_i C$ and atomic roles are roles, and \top , $C \cap D$, $\neg D$, $\exists R.C$, $\Box_i C$, $\Box_i C$ and atomic concepts are concepts ($i \in \mathbb{N}$). Observe that in our framework \Box is not the dual of \Box .

Formulas (in other DLs often called *axioms*) are either the atomic formulas \top , $C = D$, aRb , $a : C$ (with $a, b \in I$), or compound formulas $\Box_i \phi$, $\Box_i \phi$, $\neg \phi$, $\phi \wedge \psi$, with ϕ and ψ being formulas.

For convenience, we also define $\phi \rightarrow \psi = \neg(\phi \wedge \neg\psi)$.

Let *authors* = $\{s_1, \dots, s_n\}$ be a finite set of information sources (web actors, publishers, web sites, peers...), and *addressees* = $\{a_1, \dots, a_m\}$ the set of recipients (possibly overlapping or identical with *authors*). Let *actors* = *authors* \cup *addressees* be the set of all participants.

Then we define the following *social attitudes*, where $\theta : 2^{\text{authors}} \times 2^{\text{addressees}} \rightarrow \mathbb{N}$ maps elements of the cartesian product of the powerset of authors and the powerset of addressees unequivocally to a multi-modality index number (e.g., $(\{s_5, s_7\}, \{a_1, a_6\}) \mapsto 570160$), $s \in \text{authors}$, $S \subseteq \text{authors}$, and $A \subseteq \text{addressees}$. Possibly empty or singleton subsets of *authors* \cup *addressees* are called *groups*.

Opinion / Public weak assertion:

$$\Box_{S,A} \phi = \Box_{\theta((S,A))} \phi \text{ (analogously for } \Box_{S,A} C \text{ and } \Box_{S,A} R \text{).}$$

This attitude denotes that group S holds towards group A the opinion that ϕ is true. Opinions need not to be honest (thus they are also called *ostensible beliefs* [7]), and a certain author or a group can hold mutually inconsistent opinions (precisely: inconsistent propositional contents of the resp. opinions) facing different addressees.

In case S and A are identical in the definition of opinion, we use the following abbreviation:

Ostensible group belief:

$$\Box_I \phi = \Box_{I,I} \phi,$$

with $I \subseteq \text{actors}$, $|I| > 1$, denoting that in a group I ϕ is the ostensibly accepted group belief (while it is possible that a member or a subgroup of I ostensibly or sincerely believes $\neg\phi$ - an example for such a subgroup would be some politically dissident group which can articulate its true beliefs only in the underground). This important special case is close to the notion of *grounding* [5,6].

(Analogously for $\Box_I C$ and $\Box_I R$)

Public intention: $PInt_{s,A}(\phi) = \Box_{\theta(\{s\},A)} \phi$

This modality denotes that s publicly (i.e., facing A) intends that ϕ becomes true. A public intention is also ostensible only, and might have nothing in common with any true (i.e., mental) intention of s . Again, "public" means group A plus s .

Public assertion:

$$\text{Asserts}_{s,A}(\phi) = \Box_{\{s\},A} \phi \wedge PInt_{s,A}(\Box_{A,\{s\}} \phi)$$

(analogously for classes and roles instead of ϕ .)

Thus, informally, a public assertion modality states that an actor *ostensibly* believes some content, and *ostensibly* intends other actors to adopt his viewpoint in this regard (not necessarily explicitly).

The difference of assertion (*Asserts*) and weak assertion / opinion is simply that the latter attitude does by itself not include convincing the addressee from its propositional content but the information source just expresses herself. For a simplified notation for the expression of *disagreement* with some given information cf. 3.4.

Upon these definitions, various constraints extending the well-known KD45 axioms scheme could be imposed optionally, which is omitted here for lack of space (but see [4,7] for an - non-terminological - approach to this issue).

Private belief: $Bel_s(\phi) = \Box_{\{s\},\{s\}}\phi$
(analogously for classes and roles instead of ϕ .)

Private intention: $Int_s(\phi) = PInt_{s,\{s\}}(\phi)$.

Maybe surprisingly at a first glance, the ordinary mental (i.e., private) belief Bel_s of a single actor can be modeled as a special case of the former public belief, namely the ostensible belief of a singleton group. Informally, one can imagine that actor s always expresses her honest private belief if she is within a group consisting of herself only (talking to herself, so to say). Although we don't encourage the use of this operator without being prefixed with another modality (denoting, e.g., that it is publicly believed that someone privately believes something), it allows $\mathcal{S} - \mathcal{ALC}$ to model multi-agent beliefs also.

Note that while the *Asserts* operator distinguishes information source and addressee, this differentiation is not required for the semantics of the public belief operator \Box_I . The latter just expresses that a certain proposition (concept, role) is *ostensibly* (but not necessarily sincerely or correctly) believed by a certain group of agents (not to be confused with the traditional concepts of multiagent *group belief* and *common belief*, which have a different meaning - recall that ostensible group belief does *not* entail that any of the group members or subgroups believes the respective statement). E.g., $Asserts_{s_3,\{a_1,a_5\}}(\phi)$ denotes that ϕ is asserted by source s_3 , addressing recipients a_1 and a_5 , which is equivalent to saying that i) $\Box_{\{s_3\},\{a_1,a_5\}}\phi$ (i.e., s_3 ostensibly believes ϕ) and ii) $PInt_{s_3,\{a_1,a_5\}}(\Box_{\{a_1,a_5\},s_3}\phi)$. We will later describe in detail the semantics of *making* public assertions (i.e., *announcing* ostensible beliefs) in a speech act -sense.

3.2 Model-Based Semantics

The following provides a model-based semantics for $\mathcal{S} - \mathcal{ALC}$, with an integration of multi-modalities \Box_i and \Box_i using a multiple-world approach (a.k.a. Kripke-style semantics) enhancing [9]. For details on modal description logics in general please refer to [9,10].

Definition 2. A model (of $\mathcal{S} - \mathcal{ALC}$) is a pair $M = (F, I)$ with $F = (W, \triangleright_{PI}, \triangleright_{PB})$. W is a non-empty set of worlds, $\triangleright_{PI}^i \subseteq W \times W$ and $\triangleright_{PB}^i \subseteq W \times W$ are a so-called binary *accessibility relations* for public intentions and beliefs respectively, with each element linking one world to another, $i \in \mathbb{N}$. Each \triangleright_{PB}^i shall be serial, transitive and euclidian, \triangleright_{PI}^i shall be serial. Moreover, for each \triangleright_{PB}^i shall exist

one irreflexive predecessor element (thus this relation shall reflect the axioms of a KD45 modal logic - cf. [4] for details on these axioms).

I is an interpreting function which associates with each world in W an \mathcal{ALC} model $I(w) = (D, R_0^{I,w}, \dots, C_0^{I,w}, \dots, o_0^{I,w}, \dots)$. At this, D is the *domain* of the model, $R_i^{I,w} \subseteq D \times D$, $C_i^{I,w} \subseteq D$, and $o_i^{I,w} \in D$ (i.e., the o_i 's are objects within the domain).

Definition 3. The values of concepts and roles, and truth-relation $(M, w) \models \phi$ for formulas are defined as follows:

1. $\top^{I,w} = D$, $C^{I,w} = C_i^{I,w}$, $R^{I,w} = R_j^{I,w}$ for $C = C_i$, $R = R_j$
2. $x(\Box_i R)^{I,w} y$ iff $\forall v \triangleright_{PI}^i w : xR^{I,w} y$
3. $x(\Box_i R)^{I,w} y$ iff $\forall v \triangleright_{PB}^i w : xR^{I,w} y$
4. $(C \sqcap D)^{I,w} = C^{I,w} \cap D^{I,w}$
5. $(\neg C)^{I,w} = D - C^{I,w}$
6. $x \in (\Box_i C)^{I,w}$ iff $\exists v \triangleright_{PI}^i w : x \in C^{I,w}$
7. $x \in (\Box_i C)^{I,w}$ iff $\exists v \triangleright_{PB}^i w : x \in C^{I,w}$
8. $x \in (\exists R.C)^{I,w}$ iff $\exists y \in C^{I,w} : xR^{I,w} y$
9. $(M, w) \models C = D$ iff $C^{I,w} = D^{I,w}$
10. $(M, w) \models a : C$ iff $a^{I,w} \in C^{I,w}$
11. $(M, w) \models aRb$ iff $a^{I,w} R^{I,w} b^{I,w}$
12. $(M, w) \models \Box_i \phi$ iff $\forall v \triangleright_{PI}^i w : (M, v) \models \phi$
13. $(M, w) \models \Box_i \phi$ iff $\forall v \triangleright_{PB}^i w : (M, v) \models \phi$
14. $(M, w) \models \phi \wedge \psi$ iff $(M, w) \models \phi$ and $(M, w) \models \psi$
15. $(M, w) \models \neg \phi$ iff $(M, w) \not\models \phi$

3.3 Decidability

A formula ϕ is *satisfiable* w.r.t. the semantic above if there exists a pair (M, w) such that $(M, w) \models \phi$.

Theorem 1. *The satisfaction problem for $\mathcal{S} - \mathcal{ALC}$ formulas is decidable.*

As shown in [9], the satisfaction problem (i.e, whether there exists a model and a world such that $(M, w) \models \phi$) is decidable for \mathcal{ALC}_M , which apart from the multi-modalities has an identical semantics. \mathcal{ALC}_M uses arbitrary models, as well as such models which have two constrained accessibility relations corresponding to the modal logics $S5$ and $KD45$, the latter commonly used to model agent belief (as well as multi-agent belief when using multi-modalities \Box_i ($KD45_n$)). Since our accessibility relations $\triangleright_{PI/PB}^i$ observe $S5_n$ respectively $KD45_n$, and we are effectively mapping agent and group beliefs to single pseudo-agent beliefs (of pseudo-agents representing groups and ostensible attitudes) using θ , $\mathcal{S} - \mathcal{ALC}$ is decidable as well.

Related to this, it can be also easily seen that in the case we would allow only singletons for indexing the multi-modalities \Box_i , we would gain multi-agent belief modalities (in the sense of [10]). If we would additionally do the same with \Box_i ,

and also remove all constraints on the accessibility relation \triangleright , $\mathcal{S} - \mathcal{ALC}$ would "deflate" to a syntactic variant of $\mathcal{ALC}_{\mathcal{M}}$. The former is expressed with

Theorem 2. *Public singleton group belief corresponds to private individual belief.*

Thus, mental propositional attitudes can be written as a special case of social attitudes.

3.4 Social Semantics of Web Publishing

The data found on web pages or any other content on the web is in general, if taken "as is", neither knowledge nor private belief. Instead, web content needs to be interpreted as the content of communication acts. By means of such an interpretation (which essentially unfolds the semantics in the semantic web), the recipient can then classify the web artifact as knowledge (e.g., via trust), and, usually in a previous step, as opinions and other social attitudes. The step from the respective speech act of publishing (asserting, requesting, denying...) web 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 web pages as speech acts, since the speech act meaning of publishing web content is implicit. A web page claiming that "sheep are pink" is essentially nothing else than the description of an assertive speech act.

But technically, the following acts could also be used more or less directly within a document as "social semantics links" to external content (similar to the import-directive of OWL).

The following provides both a semantics for typical publishing 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 acts which are allowed to take place when certain attitudes hold. We achieve both by using the formal style (but not the content) of mentalistic semantics of agent communication languages, i.e., by specifying for each act its so-called *feasibility precondition* (FP) and its *rational effect* (RE) [6]. The former denotes what needs to hold (in terms of social attitudes in this work), the latter denotes both guaranteed illocutionary and perlocutionary effect(s) of the respective act.

Publishing acts are denoted as

$identifier_1 : author.Performative(audience, \phi | C | R | identifier_2)$, with $author \in authors$, $audience \subseteq addressees$, and $identifier \in \mathbb{N}$.

If *audience* is omitted, the respective act is either not addressed at someone, or at an unspecified group of potential recipients (as it is usually in case when the locution is asynchronously given in form of a web site, which addresses potentially the whole internet).

identifier optionally assigns the act an unequivocal number (or a time stamp or an URI), or refers to another act. The helper function *content* maps an identifier to the propositional or terminological content of the respective act.

The content of an act (class, role or axiom) is denoted as ω .

Note that in regard to the semantics of speech 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.

- $id : s.assert(A, \omega)$ (author s asserts ω towards the recipients A)
 FP: $\neg Asserts_{s,A}(\omega)$ and $\neg Asserts_{s,A}(\neg\omega)$ and $\neg \Box_{A \cup \{s\}} \Box_A \omega$
 RE: $Asserts_{s,A}(\omega)$

The feasibility preconditions (FP) here express that 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). $\neg Asserts_{s,A}(\neg\omega)$ ensures 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 expresses that it became public 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 are existing in regard to 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.

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

FP: $\neg \Box_{\{s\},A}(\omega)$ and $\neg \Box_{\{s\},A}(\neg\omega)$
 RE: $\Box_{\{s\},A}(\omega)$

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

Note that we can not simply define similar publishing acts for the utterance of group belief corresponding to \Box_{group} (except in the case that the group is in fact a single actor on the web, like an organization - but this is already covered with *assert* and *inform*). Uttering such a group belief would require some judgement aggregation procedure (like voting), and can for principle reasons not be successful in all cases.

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)
 FP: $\neg PInt_{s,A}(\omega)$
 RE: $PInt_{s,A}(\omega)$
- *id* : $s.request(a, \omega) \equiv s.intend(\{a\}, PInt_{a,\{s\}}(\omega))$
- *id* : $s.requestEach(A, \omega) \equiv \forall a \in A : s.intend(\{a\}, PInt_{a,\{s\}}(\omega))$ The latter two acts express requests directed to another agent (or a group thereof) to make some desired state come true. The act types *agree* and *deny* can be used to utter positive or negative replies to such requests, by asserting to intend resp. not to intend the requested act/state.

These makro acts are not unproblematic, since they request a potentially insincere intention ($PInt_{a,\{s\}}$). Instead, we could write $PInt_{a,\{a\}}$ to demand a sincere intention, but this would also be problematic.

4 Case Study

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 [6]), which should be quite typical for the semantic modeling of, e.g., seller/buyer platforms on the web.

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.

PI_{s_1} : *Bel*_{s₁} *discounts*

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

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

Effect: $\Box_{\{s_1\}, \{c_1, c_2\}} \neg\text{discounts}$

$\wedge PInt_{s_1, \{c_1, c_2\}} \Box_{\{c_1, c_2\}, \{s_1\}} \neg\text{discount}$

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 BDI or known semantic web languages, although being, as already pointed out, crucial for the semantic web as a public opinions platform.

Intentions can also not be modeled with any current web semantics framework known to us, including the highly relevant ostensible public intentions (*PInt*...).

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

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

Effect: $\Box_{\{s_1\}, \{c_2\}} \text{discounts} \wedge PInt_{s_1, c_2} \Box_{\{c_2\}, \{s_1\}} \text{discount}$

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 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 high accident rate**

$c_2.\text{request}(\{s_1\}, \text{InformIfAccidentRateHigh})$

Effect: $PInt_{c_2, s_1} \text{Done}(s_1 : \text{InformIfAccidentRateHigh}) \wedge \dots$, with

$\text{InformIfAccidentRateHigh} \stackrel{\text{def}}{=}$

$s_1.\text{inform}(\{c_2\}, \text{accidentRateHigh}(\vartheta)) \vee s_1.\text{inform}(\{c_2\}, \neg\text{accidentRateHigh}(\vartheta))$

$PI_{s_1} : \text{Bel}_{s_1} \text{accidentRateHigh}(\vartheta)$

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

$s_1.\text{assert}(\{c_2\}, \neg\text{accidentRateHigh}(\vartheta))$

Effect: $\Box_{\{s_1\}, \{c_2\}} \neg\text{accidentRateHigh}(\vartheta)$

Seller s_1 asserted $\neg\text{accidentRateHigh}(\vartheta)$ while thinking the opposite. Privately, c_2 believes this information (see PI_{c_2} below) and publicly agrees in the next step, but will revise her private (but not her public) belief this later.

U5 $c_1 \rightarrow \{s_2\}$: **Expression of belief**

$c_2.\text{inform}(\{s_1\}, \neg\text{accidentRateHigh}(\vartheta))$

Effect: $\Box_{\{c_2\}, \{s_1\}} \neg\text{accidentRateHigh}(\vartheta)$

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

$PI_{c_2} : \text{Bel}_{c_2} \neg\text{accidentRateHigh}(\vartheta)$

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

$c_2.\text{request}(\{s_2\}, \text{InformIfAccidentRateHigh})$

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

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

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

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 (i.e., that $\Box_{\{c_2\}, \{s_1\}} \neg\text{accidentRateHigh}(\vartheta)$).

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

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

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

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

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

Effect (together with the previous act):

$PInt_{c_2, s_2} buy(\vartheta_2, 4000\text{£}) \wedge PInt_{s_2, c_2} sell(\vartheta_2, 4000\text{£})$ (i.e., c_2 and s_2 are publicly committed to buy resp. sell ϑ_2 at the price of 4000£ now).

5 Conclusion

This paper argued that in order to allow for the logical representation of possibly controversial or uncertain web information on the web, current formal representation frameworks for web knowledge need to be enhanced for the modeling of the social meaning of information. To this end, we proposed a socially-enhanced description logic for the foundational modeling of socially acquired knowledge on the web, and a social semantics of web publishing acts in terms of the dynamics of social (i.e., communication) attitudes. Next steps will concentrate on a practical evaluation, the enhancement of other relevant languages such as *SHOIN(D)* with operators for social attitudes, and a full axiomatization of the modal logic.

Acknowledgements. This work is funded by Deutsche Forschungsgemeinschaft (DFG) (research project *Open Ontologies and Open Knowledge Bases*, contract BR609/13-1). I would also like to thank the anonymous reviewers for their very valuable comments.

References

1. A. Maedche, F. Nack, S. Santini, S. Staab, L. Steels. Emergent Semantics. IEEE Intelligent Systems, Trends & Controversies, 17(2), 2002.
2. D. Calvanese, G. De Giacomo, M. Lenzerini. Ontology of Integration and Integration of Ontologies. Procs. of the 2001 Description Logic Workshop (DL 2001), 2001.
3. J. A. Plaza. Logics of Public Communications. In M. L. Emrich et al (eds.). Procs. of the 4th International Symposium on Methodologies for Intelligent Systems, 1989.
4. F. Fischer, M. Nickles. Computational Opinions. Procs. of the 17th European Conference on Artificial Intelligence (ECAI-06), 2006.
5. B. Gaudou, A. Herzig, D. Longin. Grounding and the expression of belief. Procs. of the 10th Intl. Conf. on Principles of Knowledge Representation and Reasoning (KR 2006), 2006.
6. 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.
7. 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-05), 2005.

8. M. Schmidt-Schau, G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48(1):1-26, 1991.
9. F. Wolter, M. Zakharyashev. Modal description logics: Modalizing Roles. *Fundamenta Informaticae*, v.39 n.4, p.411-438, 1999.
10. A. Laux. Beliefs in multi-agent worlds: A terminological approach. In *Procs. of the 11th European Conference on Artificial Intelligence*, 1994.
11. J. Golbeck, B. Parsia, J. Hendler. Trust networks on the semantic web. *Procs. of the 7th International Workshop on Cooperative Information Agents (CIA)*, 2003.
12. 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.
13. V.A.M. Tamma. An Ontology Model Supporting Multiple Ontologies for Knowledge Sharing. PhD Thesis, The University of Liverpool, 2002.
14. 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.
15. R. Hull. Managing semantic heterogeneity in databases: A theoretical perspective. In *Procs. of PODS-97*, 1997.