ISWC 2006 Workshop on Uncertainty Reasoning for the Semantic Web
November 5th 2006, Athens (GA), USA

Social Contexts and the Probabilistic Fusion and Ranking of Opinions: Towards a Social Semantics for the Semantic Web

Matthias Nickles[1] Ruth Cobos[2]

[1]AI / Cognition Group
Department of Computer Science
Technical University of Munich
nickles@in.tum.de

[2]Department of Informatics
Universidad Autónoma de Madrid (UAM)
ruth.cobos@uam.es

Introduction

- Web publishers are autonomous, possibly insincere actors
- Heterogeneous and controversial viewpoints on the (Semantic) Web are not just design problems, but inevitable and sometimes even useful in social environments
- Almost all traditional approaches to inconsistency and insincerity on the Web rely on some kind of filtering (by criteria such as trust or preferences)
- No sufficient means yet for
    - the formal representation of "the Web" *itself,* in terms of *subjectivity, competing viewpoints, divergent opinions, intentionality, group knowledge…*
    - the probabilistic/voting-based aggregation and representation of *group beliefs*

  = ***Social Semantics***
- This work introduces a Description Context Logic for the representation of so-called certain and uncertain *Social Attitudes* (= public beliefs and intentions of groups and individuals)

Overview

- Issues
- Approach outline
- Features
- Social Attitudes
- Syntax (non-probabilistic)
- Semantics (non-probabilistic)
- P-SOC-OWL (probabilistic)
- Fusion and ranking of opinions
- Conclusion

Issues

- Semantic heterogeneity with impossible (or too expensive) alignment, subjectivity and irresolvable belief conflicts
- Inconsistent information simultaneously addressed to different parties and information published by/for specific social groups (e.g., in web blogs or intranets)
- Complexity caused by semantic heterogeneity
- Modeling of intentionality “behind” information
- OWL provides no suitable means for formalization of subjectivity
- RDF not suitable due to shallow / non-existent semantic wrt. subjectivity and sociality
- Traditional approaches to *provenance* only semi-formal or non-formal
- Modal logic usable (partially), cf. ISWC 2006 talk on *Social Attitudes*

Approach outline

- ***Social Reification***:
  Subjective statements are lifted to the social level.

  E.g., “Sheep are red” → “Actor/Group *s* asserts: ‘Sheep are red’”
- Social Reification makes an inconsistent knowledge base or ontology consistent in case each of the sources provides consistent knowledge
- Technically: Contexts for individual and group beliefs and intentions
- Formal languages *SOC-OWL* (based on SHOIN(D) and C-OWL [Bouquet et al]) and *P-SOC-OWL* (based on C-OWL and P-SHOQ(D) [Giugno and Lukasiewicz])
- Modeling of social propositional attitudes (ostensible beliefs and intentions of individuals and groups)
- Private (mental) attitudes covered as a special case
- Optional: Differentiation of addressees (for the modeling of different “publics” or closed groups)
- Optional: Fusion operators for the aggregation of multiple opinions to a single group opinion

Features

- Simple and intuitive extensions of standard Description Logics / OWL-DL
- Consistent representation of inconsistencies acquired from dissenting sources (by means of “agree to disagree”)
- Allows for the modeling of both public and private beliefs and intentions
- Allows for the differentiation of several addresses
- Allows to aggregate multiple dissenting assertions via belief fusion operators
- BDI (standard framework for the reasoning about mental attitudes in agent research) is essentially a special case

Social Attitudes (1)

- SAs model individual and group attitudes regarding information (approval, denial, desire…) logically
- ***Opinion***: The public (i.e., possibly insincere) attitude of an agent regarding the truth of a certain statement
- ***Ostensible intention***: The public (i.e., possibly insincere) attitude of intending that a certain statement (or description, given DL) shall become true
- Triggered by implicit or explicit communication acts (e.g., web publishing, HTML links, web service assertions, discourses by blogging …)
- Allow, e.g., to model the "web personalities" of mentally opaque actors
- Multiple (possibly inconsistent) "web personalities" of the same actor are possible, even simultaneously
- Mental attitudes (i.e., personal, opaque beliefs and intentions) as a special case

Social Attitudes (2)

- Held by actors (individual or social group) towards a group of addressees
- Represented using ***Social Contexts*** (each indexed by attitude type $\sqcup$ source(-s) $\sqcup$ addressee(-s))
- Information ("passive" opinion): $\mathbf{Information(a_1, a_2, \varphi)}$
- *Public intention*: $\mathbf{PublicIntention(a_1, a_2, \varphi)}$
- *Assertion* ("active" opinion): $\mathbf{Assertion(a_1, a_2, \varphi)}$ = Information(a_1, a_2, φ) ∧ PublicIntention(a_1, a_2, Information(a_2, a_1, φ)) (i.e., includes the intention to *convince* the addressees)
- *Mental beliefs and intentions* can be formulated as special cases:
  - $\mathbf{Bel(a_1, \varphi)}$:= Information(a_1, a_1, φ)
  - $\mathbf{Int(a_1, \varphi)}$:= PublicIntention(a_1, a_1, φ)
  - Internal group beliefs: Information(g, g, φ)

Syntax of *SOC-OWL*

(slightly simplified, cf. paper)

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

TBox and ABox:

$$C_1 \sqsubseteq C_2,\ Trans(R),\ R \sqsubseteq S,\ T \sqsubseteq U,\ C(a), R(a,b),\ a = b,\ a \neq b,$$

"SBox":

$$\overset{attitude}{source_1,...,source_n \rightarrowtail addresse_1,...,addressee_n}\ s$$

attitude: $assertion, information, publicIntention$

(+ bridge rules, socially contextualized concepts, individuals and roles, …)

Example

$ControversialPerson(columbus)$

$\underset{tim,tom \rightarrowtail tina}{\overset{assertion}{}}(\neg Hero)(columbus)$

$\underset{tina \rightarrowtail tim,tom}{\overset{assertion}{}}Hero(columbus)$

$\underset{tim,tom \rightarrowtail tina}{\overset{assertion}{}}Exploiter(columbus)$

Note:

- Group beliefs do not constrain beliefs of group members (i.e., no *common knowledge* here, but cf. *fusedInformation* in paper Sect. 3.2)

- A certain source can hold mutually inconsistent opinions towards different addressees

- Nesting of contextualized statements not allowed (unlike in “real” Context Logic):

$\underset{tina}{\overset{publicIntention}{}}(\underset{tim,tom \rightarrowtail tina}{\overset{information}{}}(\neg Exploiter)(columbus))$

Semantics of *SOC-OWL*

(simplified, cf. paper & C-OWL)

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

Some meta-axioms:

- $\overset{assertion}{s_1,\ldots,s_n \rightarrowtail a_1,\ldots,a_m} \varphi \rightarrow ((\overset{publicIntention}{s_1,\ldots,s_n \rightarrowtail a_1,\ldots,a_m} \sqcup \overset{assertion}{a_1,\ldots,a_m \rightarrowtail s_1,\ldots,s_n}, e) = (\overset{assertion}{s_1,\ldots,s_n \rightarrowtail a_1,\ldots,a_m}, e))$
- $\overset{assertion}{s_1,\ldots,s_n \rightarrowtail a_1,\ldots,a_m} \varphi \rightarrow \overset{information}{s_1,\ldots,s_n \rightarrowtail a_1,\ldots,a_m} \varphi$
- Mutual consistency of all statements within a specific context
- […]

P-SOC-OWL

SOC-OWL, plus contextualized statements with optional ***probability intervals*** for SBoxes:

$$[p_l, p_u]\ {}^{attitude}_{source_1,\ldots,source_n \rightarrowtail addresse_1,\ldots,addressee_n}\ s$$

Interpretations extended with subjective probability functions μ_{id} (cf. P-SHOQ(D)) for each context id:

$$PI = \{(PI_{id}, \mu_{id}) : id \in Id\}$$

Examples:

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

$$0.9:\ {}^{assertion}_{tim} Hero(columbus)$$

$$0.7:\ {}^{assertion}_{tina} Hero(columbus)$$

…plus additional context type for aggregation: ***fusedInformation***

Fusion and ranking of opinions (1)

Two approaches to derive social group opinions from other opinions:

$$(\bigwedge_{s_i \in \{s_1,...,s_n\}} (Pr_{\substack{information \\ s_i \rightarrowtail addressees}} \models \varphi[p_i, p_i])) \rightarrow (Pr_{\substack{information \\ s_1,...,s_n \rightarrowtail addressees}} \models \varphi[p, p])$$

with pooling result $p = pool^{poolingType}((p_1, ..., p_n), priorKnowledge)$

Problem: Possible inconsistencies in case freely defined group opinions exist in the KB/ontology.

Therefore:

$$(\bigwedge_{s_i \in \{s_1,...,s_n\}} (Pr_{\substack{information \\ s_i \rightarrowtail addressees}} \models \varphi[p_i, p_i])) \rightarrow (Pr_{\substack{fusedInformation \\ s_1,...,s_n \rightarrowtail addressees}} \models \varphi[p, p])$$

Fusion and ranking of opinions (2)

Standard pooling operators:

Averaging:

$$pool^{avg}((p_1, ..., p_n), \varnothing) = \frac{\sum p_i}{n}$$

Linear pooling with weights (e.g., to consider *trust degrees* or *social power*):

$$pool^{LinOP}((p_1, ..., p_n), (weight_1, ..., weight_n)) = \sum weight_i p_i,$$

$$\text{with } \sum_{weight_i} = 1$$

Logarithmic pooling:

$$pool^{LogOP}((p_1, ..., p_n), (weight_1, ..., weight_n)) = \kappa \prod_{i=1}^{n} p_i^{weight_i}$$

Fusion and ranking of opinions (3)

Example using $pool^{avg}((p_1, ..., p_n), \varnothing) = \frac{\sum p_i}{n}$

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

$\vdash \quad 0.8$: $^{assertion}_{tina,tim} Hero(columbus)$

Induced rankings:

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

Conclusion

- Presented a context-based Semantic Web language for the formal modeling of the “social dimension” of the Web and the Semantic Web
- Related developments:
  - FOL/dynamic logics for opinions and other social attitudes (cf. e.g. ECAI 2006)
  - Alternative DL-based approach using modal logic instead of contexts (cf. ISWC 2006 paper on Social Attitudes)
  - Further works: **`http://www.openontology.info`**
- Future work:
  - Issues:
    - Nested contexts
    - Open domains of actors (sources and addressees)
  - Implementation of a reasoner
  - Identification of emergent groups (knowledge communities) on the Web

Thank you very much for your attention!