The Business of Ontology calls for a Formal Pragmatics

Hans AKKERMANS

The Network Institute
Informatics Department
VU University Amsterdam
www.thenetworkinstitute.eu; www.e3value.com

Abstract. I develop a few suggestions how formal ontology can meet industry and practice. This is not just a matter of useful ontology-based applications. That is a necessary but also rather trivial idea that perpetuates the gap between fundamental and applied research and keeps alive associated outdated linear process ideas of innovation. Instead, I suggest that for further progress the ontology field is to move from a semantic to a pragmatic approach in the communication-theoretic sense. Formal pragmatics stands a better chance to provide a rational and scientific foundation for an integrated ontology theory as well as practice than the deductivist formal semantics approach.

Introduction

The research tradition that goes under the name of "ontology" within Computer Science (CS) and, more specifically, Information and Knowledge Systems (IKS) is now about twenty years old. Whatever the adopted measure – strictly academic or in terms of socio-economic or industrial usefulness – the field of ontology has been extremely successful, much more so than the first generation of CS "ontologists" and "ontology engineers" (including the present author) had anticipated.

After twenty years, it is a good time to take stock, not only of what has been achieved, but (more importantly and interestingly) also what still needs to be done, and especially what has to be changed in the direction we are going if we undertake to make further significant scientific as well as practical-applicational progress. In the following I develop (continuing work in [1]) a few observations, positions, and proposals on this.

1. Ontologies are us (or them?)

To cite the very recent (December 2007) definition in the Handbook of Knowledge Representation [2]:

> "An ontology is an explicit specification of a shared conceptualization that holds in a particular context" [2].

The (slow and gradual) changes of the definition of what an ontology is are interesting in their own right. Gruber's original definition takes ontology to be an explicit specification of a conceptualization [3]. The introduction of the adjective *shared* occurred some years later, especially in the work of Borst et al. [4]. The reference to validity in a particular *context* is again a later addition, and attempts to encapsulate lessons learned in building and using ontologies in practical applications, especially in knowledge management [5] and on the Web [6].

Increasingly, the realization now is that context first of all concerns *social* context. Ontologies as conceptualizations express social constructions and play a role as such in Web communities of practice and social networks, an idea aptly formulated by Péter Mika as "ontologies are us" [7]. Thus, the original idea of ontologies is that of a consciously engineered artefact to express intended (and as far as possible universalistic) meaning through specifying the formal semantics of concepts. In the latter conception, ontologies figure as social network-created, even emergent, "meaning productions" inseparable from the context of the community in which they are created.

Many authors have noted that these two different approaches to ontologies do not necessarily exclude each other. There are many different types of ontologies that have demonstrated practical usefulness in applications, in both approaches. Also, both approaches need critical scrutiny. In the engineering approach, an important critical issue is the distance (often expressed in terms of intuitiveness and perceived complexity) between the engineered ontology and its intended user or application community. In the social community approach, on the other hand, a critical issue is that in ontology extracted as a matter-of-fact empirical (i.e. non-engineered) social phenomenon ultimately "anything goes": shared meaning simply becomes what any community happens to believe, without any further rational foundation or justification. Just one example I am thinking specifically of is the xenophobia that has in recent years become a plague in my own country and others, and that also has found its way to the Internet in blog communities. It would be an easy task to extract a really shared ontology here (certainly when applying the maxim or slogan [2] "smaller ontologies are better"), but it will be one that is much less harmless than the conceptualizations shared by, say, the Semantic Web research community. Artefacts should not be approached uncritically, but neither should communities. Rational foundations and justifications continue to be necessary in both cases.

2. Formal pragmatics (rather than formal semantics)

Nevertheless, a fundamental realization that is here to stay is (i) theoretically, the inherently *social* nature of ontologies as shared conceptualizations; (ii) practically, the inherently *contextual* nature of ontologies as useful and valid support tools for information and knowledge sharing. So, as a next step, ontology as a scientific and computational engineering theory of conceptualization needs to account for this in a principled way.

A starting point to do so (cf. [8]) is to reconsider the, underlying and in part implicit, conception of *meaning* itself as it has been employed e.g. in ontology engineering. To date, much of ontology engineering derives its foundational ideas from the standard truth-conditional view: essentially, the meaning of a sentence or

proposition or utterance coincides with the conditions for its truth. This is a view that goes back to logicians such as Frege and Tarski and is still held by many modern model-theoretic logicians, formal semanticists, theoretical computer scientists, and knowledge representation researchers.

However, the above characterization of ontology as inherently social and contextual in nature leads us to entertain a different foundational notion of meaning: the meaning of a sentence or proposition or utterance lies in its actual *use* in communication. This is a view that goes back to the later Wittgenstein, and has been developed by a (wide) variety of philosophers such as Dummett, Grice, Austin, Searle, Habermas, linguists such as Sperber and Wilson, and argumentation theorists such as Van Eemeren & Grootendorst, and Walton.

I contend that this pragmatic view of meaning will show to be the more fruitful one for the field of ontology engineering and the Social Semantic Web. Rather than formal semantics, in ontology we need a both conceptual and computational theory of formal pragmatics. To come to such a theory, I further contend that there are several research lines, also in computer science and artificial intelligence, with relevant results already available that are just waiting to be taken up. Below, I will suggest just a few components of this.

3. Problem-Solving Methods: what happened to them?

As quite an elementary point of context, ontologies are formal conceptualizations not made *l'art pour l'art*, but to help achieve a goal or task by an actor. Often that task involves knowledge-intensive reasoning. The conceptual distinctions that we make in our attempts to understand the world are not just static and descriptive domain characterizations, but they are made to serve practical purposes of action by someone or something in that world. We cannot detach knowledge from action.

Task and actor characteristics are key parts of operational context definition [9], but left out or kept implicit in much current Web ontology research.

Two decades of knowledge engineering (cf. [2] and references therein) have delivered a wealth of evidence that there are recurring patterns or stereotypes in the structuring and use of knowledge as an instrument in tasks that involve reasoning and computing. These recurring knowledge stereotypes are variously referred to as inference schemas, task templates, strategy patterns, or problem-solving methods (PSMs). PSMs are heuristic and stereotypical in the sense that they do not guarantee to solve a given knowledge-intensive problem in general (unlike a normal algorithm). But, they do have demonstrated pragmatic value in solving common, typical, or average cases of knowledge-intensive tasks that can, moreover, be reused in many different situations. The 80-20 rule applies here: that PSMs as stereotypical schemas only handle the typical or average case is often seen as a kind of algorithm bug in traditional computer science, but it is in fact a useful feature on the Web as an open dynamic system of unknown but certainly very high complexity (cf. [10]).

As an example, consider PSMs such as Propose-and-Revise or Propose-Critique-Modify that have been originally developed for configuration design tasks in engineering. Such tasks have customer requirements as input, and an artefact

description as output. The specific task type of configuration design presupposes a generic domain ontology in which predefined functional components, "hard" constraints, and "soft" partial preference orderings are key concepts, on top of which the PSMs can do their work. Concepts from the domain ontology are then invoked by the PSM as a specific knowledge role (i.e. way of use) in a certain phase of the reasoning process.

This is clearly a pragmatic, use-oriented view of ontologies. Part of the context surrounding the normal static ontologies is computationally defined by explicating the reasoning task and methods in which ontologies find their use. Another part of context is explicating the task from the (often non-technical) perspective of the agent or actor who carries out that task, such as users or customers. In the family of networked business ontologies (such as *e*^{3}*value*) that our Business Informatics group at Amsterdam has developed, we employ configuration PSMs for automated service bundling, based on a service ontology that does not take – as is usually the case – the technical Web Service perspective, but defines services as an economic/commercial entity in business supplier-customer terms [11]. Other PSMs (planning, parametric design) are also candidates for provisioning of service bundles over the Web.

Two decades of knowledge engineering has provided a library of useful task decomposition and inference method schemas. Semantic Web and Web Service research has tended to overlook much of older PSM research, with the result that handling different task and actor contexts is dealt with in an unnecessarily limited (i.e. unintelligent) way. Instead, and on their turn, PSMs could themselves be made available as services on the Web. For that to work, they also need to be supplied with background knowledge how to use them as a service, i.e. with some kind of task/method ontology that explicates in what kind of context (typical or average case) they work and what the requirements (knowledge roles of domain ontology) for their practically successful functioning are.

4. The dialectical turn: argumentation in dialogue

Another relevant strand of research in pragmatics that has recently made great theoretical and computational progress is argumentation theory (for a survey, see [12]) and the related fields of informal logic and critical thinking. Analogous to PSMs, it has produced a sizable library of argument schemas that model often occurring prototypical arguments. It has two important characteristics relevant to ontology-based reasoning and semantic approaches to the Web (even though it does not have much to say about the issue of conceptualization *per se*):

- It breaks away from the deductivism that is still prevalent also in formal ontology and semantics. Namely, there are many arguments and argument schemas that are not deductively valid, are defeasible instead, but still are to be considered as good and acceptable (typically because they withstand scrutiny by thorough critical questioning). These non-deductive, "presumptive" arguments naturally occur in settings of commonsense everyday reasoning, are irreducible to deductive reasoning, but still can be explicated in a formalized and computational way.

- The dialectical turn. It views reason and argument as an inherently social process: argumentation concerns a certain type of (rational) communication and dialogue that is subject to certain social rules and principles. Critical questioning is the dialectical mechanism that tests the quality of an argument and its claim in contexts where deductive certainty is impossible to come by (because of bounded resources, time, lack of information, etc.). Clearly, this is the context that constitutes the normal condition of the Web.

As an example, argumentation theory and informal logic have developed reusable schemas for practical reasoning (answering the question: what should we do if we want to achieve goal G, under conditions of uncertainty). That has some promise to be applicable when reasoning about service provisioning over the open dynamic system environment of the Web. One might consider constructing this as a kind of PSMs with built-in critical tests for their applicability and validity.

Very first attempts to introduce an argumentative approach to the Semantic Web go under the flag of the claim web [13] and the argument web [14]. Independently of the value of such proposals, however, argumentation theory has several important things to say in general to the field of ontology engineering and the semantic side of the Web.

As a theory of (a certain area of) pragmatics, argumentation theory is able to cover several aspects of context, especially concerning the communicative relationships between actors. This is definitely of importance in talking about a Social Semantic Web. It furthermore introduces ontological distinctions between different types of dialogue, which in fact provides a characterization of different kinds of communicative goals and tasks of actors. Rather than a purely empirical emergent semantics [15] based e.g. on gossip algorithms as a mechanism modelling communication between agents, it could provide more structured and critically-rationally justifiable alternative mechanisms, in the form of dialectical-argumentative schemas which are still social and communicative in nature. Thus it not only widens the range of possible mechanisms to achieve forms of self-organization (more generally, self-* properties) on the Web, but also might give them a better rational foundation that is moreover communicatively sharable.

5. Formal Pragmatics and discursive rationality

A natural extension of my proposed idea to introduce results of theories of pragmatics into ontology engineering and semantic-based approaches to the Web, concerns various brands of speech act and communicative action theory [16, 17]. In a basic form, Searle's speech act theory has found its way into agent communication languages. But, in current semantic approaches to ontology-based Web (service) engineering this plays at most a peripheral role.

However, speech act theory is important to the Web because it extends – much more radically than ontology-based PSMs and argumentation schemas – the types of communicative acts that can be treated computationally. Formal semantics, also in ontology engineering, is still very much focused on those utterances that express propositions or assertions (true or false beliefs about the world). But there are many other types of speech acts that play an important role in an open Web environment that

is of a social nature. In e-business, for example, the majority of speech acts relates to negotiating and transacting and so is definitely not of a propositional nature. This is a phenomenon not adequately accounted for in current Web-semantic based approaches. In a semantic-based Service Science [18] and ICT-based innovation that is to be practically as well as scientifically adequate the non-propositional character of much Web communication cannot stay in the periphery of attention, as is now in fact the case. However, as indicated above, by employing and integrating various strands of existing research results relevant to the pragmatics of the Web and of Web intelligence, significant progress is in my opinion clearly within reach.

Further developments related to speech act, communication and social network theory have taken place but are still beyond the horizon of ontology and semantic-based technology. For example, Searle in his critique of the standard model of rationality in [16] argues that the intentions, desires, and goals of agents are not input to reasoning but should be part of reasoning. Habermas's communicative action theory sees *all* speech acts (not just propositions that assert something) as inherently making validity claims (regarding relevance, truth, rightness, truthfulness) that are *discursively* criticizable; and in this fact also lies their rational justification [17]. If we develop a corresponding knowledge-level computational framework for this, we will have not a Web of data, not a Web of enriched semantic annotations or of community chat and gossip, but a *reflective* Web. Truly reflective open systems are currently still out of scope (cf. [19]). But reaching beyond semantics to pragmatics is not; neither scientifically nor practically.

References

[1] J.M. Akkermans and J. Gordijn: *Ontology Engineering, Scientific Method, and the Research Agenda*, in S. Staab and V. Svatek (Eds.): Managing Knowledge in a World of Networks (Proc. 15th Int. Conf. on Knowledge Engineering and Knowledge Management EKAW-2006), pages 112-125, Lecture Notes in Computer Science LNCS-4248, Springer-Verlag, Berlin, D, 2006.

[2] G. Schreiber: *Knowledge Engineering*, in F. van Harmelen, V. Lifschitz, B. Porter (Eds.): *Handbook of Knowledge Representation*, Ch. 25, pages 929-946, Elsevier, Amsterdam, NL, 2007.

[3] T. Gruber: A translation approach to portable ontology specifications, *Knowledge Acquisition* **5** (1993), 199-220.

[4] P. Borst, H. Akkermans, and J. Top: Engineering ontologies, *Int. J. Human-Computer Studies* **46** (1997), 365-406.

[5] P. Mika and J.M. Akkermans, Towards a New Synthesis of Ontology Technology and Knowledge Management, *Knowledge Engineering Review* **19**:4 (2004), 317-345

[6] See e.g. the Special Issue on Contexts and Ontologies, *Knowledge Engineering Review* **23**:1 (2008).

[7] P. Mika: Ontologies are us: A unified model of social networks and semantics, *J. Web Semantics* **5** (2007) 5-15.

[8] C. Bianchi and N. Vassallo: Meaning, Contexts, and Justification, in B. Kokinov et al. (Eds.): CONTEXT 2007, pp. 69-81, LNAI 4635, Springer-Verlag, Berlin, D, 2007.

[9] A. Zimmerman, A. Lorenz, and R. Oppermann: An operational definition of context, in B. Kokinov et al. (Eds.): CONTEXT 2007, pp. 558-571, LNAI 4635, Springer-Verlag, Berlin, D, 2007.

[10] F. van Harmelen: *Semantic Web Research anno 2006: main streams, popular fallacies, current status and future challenges*, in M. Klusch, M. Rovatsos, and T. Payne (Eds.), Proc. 10th Int. Workshop on Cooperative Information Agents (CIA-2006), pages 1-7, LNAI 4149, Springer-Verlag, Berlin, D, 2006.

[11] For publications and tools on the e^3 family of networked business ontologies, see www.e3value.com.

[12] T.J.M. Bench-Capon and P.E. Dunne: Argumentation in artificial intelligence, *Artificial Intelligence* **171** (2007), 619-641.

[13] T. Groza et al.: SALT: Weaving the claim web, in K. Aberer et al., Proc. ISWC/ASWC 2007, LNCS 4825, pp. 197-210, Springer-Verlag, Berlin, D, 2007.

[14] I. Rahwan, F. Zablith, and C. Reed: Laying the foundations for a World Wide Argument Web, *Artificial Intelligence* **171** (2007), 897-921.
[15] K. Aberer et al.: *Emergent semantics principles and issues*, in: Database Systems for Advanced Applications Ninth International Conference, LNCS 2973, pp. 25–38, Springer-Verlag, Berlin, D, 2004.
[16] J. Searle: *Rationality in action*, MIT Press, Cambridge, MA, 2001.
[17] J. Habermas: *On the pragmatics of communication*, MIT Press, Cambridge, MA, 1998.
[18] J. Spohrer and D. Riecke (Eds.): Special Issue on Service Science, *Communications of the ACM* **49**:7 (2006), 30-87.
[19] M. Reinders, E. Vinkhuyzen, A. Voss, J.M. Akkermans, J.R. Balder, B. Bartsch-Spörl, B. Bredeweg, U. Drouven, F. van Harmelen, W. Karbach, Z. Karssen, A.Th. Schreiber, and B.J. Wielinga: A Conceptual Modelling Framework for Knowledge-Level Reflection, *AI Communications* **4** (1991), 74-87.