Patterns for the Pragmatic Web[1]

Aldo de Moor
STARLab, Vrije Universiteit Brussel,
Pleinlaan 2, 1050 Brussels, Belgium
ademoor@vub.ac.be

Abstract. The Semantic Web is a significant improvement of the original World Wide Web. It models shared meanings with ontologies, and uses these to provide many different kinds of web services. However, shared meaning is not enough. If the Semantic Web is to have an impact in the real world, with its multiple, changing, and imperfect sources of meaning, adequately modeling context is essential. Context of use is the focus of the Pragmatic Web and is all-important to deal with issues like information overload and relevance of information. Still, great confusion remains about how to model context and which role it should play in the Pragmatic Web. We propose an approach to put ontologies in context by using pragmatic patterns in meaning negotiation processes, among other meaning evolution processes. It then becomes possible to better deal with partial, contradicting, and evolving ontologies. Such an approach can help address some of the complexities experienced in many current ontology engineering efforts.

1. Introduction

The World Wide Web has profoundly changed the way people collaborate. Whereas e-mail has lowered the threshold for interpersonal communication by providing a medium for fast, cheap, ubiquitous and global communication, the Web has become the metaphor and technology for doing the same with respect to linking and sharing knowledge resources. Even for the computing community, used to fast technological progress, the speed with which the Web has evolved from initial prototype to a foundation of daily life has been dazzling. It was only in 1991 that the following was announced by a then unknown employee from CERN:

> "The WorldWideWeb application is now available as an alpha release in source and binary form from info.cern.ch. WorldWideWeb is a hypertext browser/editor which allows one to read information from local files and remote servers. It allows hypertext links to be made and traversed, and also remote indexes to be interrogated for lists of useful documents. Local files may be edited, and links made from areas of text to other files, remote files, remote indexes, remote index searches, internet news groups and articles … This project is

[1] Invited paper, Proc. of the 13th International Conference on Conceptual Structures (ICCS 2005), Kassel, Germany, July 2005, LNAI. Springer Verlag, Berlin, pp.1-18.

experimental and of course comes without any warranty whatsoever. However, *it could start a revolution in information access* [my italics]"[2].

The rest, as they say, is history.

The rise of the World Wide Web has led to many benefits to society. Documents, news, and results to queries can be obtained 24 hours a day from all over the world. The Web has given a huge boost to research, education, commerce and even politics. An interesting example of how deeply the Web has become embedded in the fabric of our globalizing society is the significant role web sites play in political reforms in less-than-democratic countries [17]. Still, not all is good. One serious consequence of the explosion of Web-accessible information resources is information overload. It is not uncommon to get hundreds, thousands, or even millions of hits when looking for a certain piece of information. Increasingly, the problem shifts from making information accessible, to delivering *relevant* information to the user.

The Semantic Web plays an important role in making the Web more relevant. Berners-Lee, et al. [1] present a cogent view of how the Semantic Web will structure meaningful content and add logic to the Web. In this web, data and rules for reasoning about data are systematically described, after which they can be shared and used by distributed agents. Granted, many of the basic theoretical ideas were already conceived by the AI community in the 1970s and 80s. The added value of the Semantic Web, however, is that this theory is finally being put into large scale-practice. The main components implementing this Web vision include techniques such as XML, for adding arbitrary structures to documents; RDF, to express meaning by simple statements about things having properties with values; and ontologies, to formally describe concepts and their relations. A typical ontology, in the sense of being an explicit specification of a conceptualization [10], consists of a taxonomy with a set of inference rules. Ontologies can be used to improve the accuracy of, for instance, Web search and service discovery processes. Ultimately, such an approach should lead to the evolution of human knowledge by scaling up collaboration from individual efforts to large, joint endeavors. Multiple ontologies then come into play. By selecting the right ontology for the right task, knowledge exchange, at least in theory, could become more effective and efficient.

In practice, however, the Semantic Web comes with its own set of problems. Voices are increasingly being heard that there is a need not only for explicitly taking into account the semantics, but also the pragmatics of the Web, e.g. [25,26,13,7,29,22]. Still, ideas and proposals are preliminary and sketchy and need further elaboration and integration. With this paper, we hope to contribute to the further maturation of thought on this important subject. We have two main objectives: finding out (1) what are fundamental conceptual elements of the Pragmatic Web and (2) how to use these elements in making meaning represented in semantic resources more relevant. In Sect. 2, we outline some contours of the Pragmatic Web that are becoming visible at the moment. This analysis results in a conceptual model of the Web in Sect. 3, outlining how the Semantic and the Pragmatic Web are interrelated. In Sect. 4, we focus on pragmatic patterns as a way to operationalize the pragmatics of

[2] Tim Berners-Lee, *comp.sys.next.announce* newsgroup, Aug.19, 1991.

the Web. In Sect. 5, we present a scenario of how a Pragmatic Web could look in practice. We end the paper with a discussion and conclusion.

2. Contours of the Pragmatic Web

The Semantic Web, with all its (potential) benefits, still poses a number of difficult challenges, both with respect to the ontologies which contain the shared meanings and the services in which these are used.

Unlike data models, ontologies contain relatively generic knowledge that can be reused by different kinds of applications. Ontologies should therefore not be too tightly linked to a specific purpose or user group [30]. To select the right (parts of) ontologies, the communicative situation needs to be taken into account. To this purpose, a "mindshaking procedure" needs to be developed, in which a formal language for information exchange is determined (syntax), and a synchronisation of the meaning of concepts (semantics) takes place on the basis of a particular context, such as purpose, time, date, or profile [29]. An example of a (typically) manual version of such a procedure is described in [9]. There, a conceptual model supervisor regularly creates reports of existing classes. If concepts seem to be in conflict, and the conflicts are important enough, the model supervisor starts and controls a discussion among stakeholders, who can be either modelers or representatives from the involved departments. If the conflict remains unresolved, both concepts remain in the model marked with their own namespaces.

Ontologies are not an end in themselves. One of the major functions of the Semantic Web is to provide access to web services. These are often described and invoked through central registries. However, for describing, discovering, and composing web services, a semantic approach is not enough. Services cannot be described independently of how they are used, because communities of practice use services in novel, unexpected ways. Social mechanisms are therefore needed for evaluating and discovering trustworthy providers and consumers of services, taking into account contexts and interactions in the composition of service applications [25-26].

Clearly it is not sufficient to model semantics to resolve such issues related to the use of ontologies. Contextual elements like the community of use, its objectives and communicative interactions are important starting points for conceptualizing the pragmatic layer. These elements are combined in a conceptualist perspective. In such a view, meanings are elements of the internal cognitive structures of language users, while in communication, the conceptual structures of different views become attuned to each other [13]. We can therefore make a distinction among shared semantic resources, such as ontologies; individual pragmatic resources, i.e. the internal conceptual models of users applying the semantic resources to their own purposes; and common pragmatic resources, in which joint *relevant* meanings have been established through communication. In communication between users aiming at achieving joint objectives, concepts that are part of individual and common pragmatic resources are selected, defined, aligned, and used. Finding out how such a meaning

negotiation process works is essential to understanding the pragmatics of the Web, and to developing (partially) automated support processes for meaning negotiation.

Developing sound and complete pragmatic perspectives, models, and methods can shed light on the confusing debates raging in the ontology and Semantic Web research communities. One fundamental question, for example, is whether the way to go is to develop large, detailed, standard ontologies such as Cyc[3] or myriad independent, domain-specific, micro-ontologies, one for each application. The answer is not either/or, but a mix of both approaches. A major reason why such a hybrid point of view cannot be easily adopted and defended, is that the real issues underlying these debates are not semantic, but pragmatic. The focus of many of these debates has thus been the wrong one, without the ontological engineering community making any significant progress on resolving the underlying issues.

Before further examining the Pragmatic Web, we first take a closer look at some of the finer details of pragmatics.

2.1 A Primer of Pragmatics

A traditional source of problems, often found in traditional conceptual modelling approaches, is to try and produce THE description of a joint reality. If members of a particular community disagree, the modellers, in the best case, keep negotiating explicit meanings until everybody agrees. If no agreement can be reached (or is not even sought) modellers often impose a meaning by choosing an ontology definition or system specification themselves.

A pragmatic approach, on the other hand, should allow for contradictions, different importance weights of information and subtle cultural differences [9]. Such differences, however, create problems of their own if not handled properly. Collaboration often fails, not because participants do not want to collaborate, but because pragmatic errors lead to the breakdown of the social and contextual components of a discourse [14]. To become successful, a pragmatic approach thus needs to acknowledge and adequately handle ambiguity and consequences of (differences in) semantics.

Facts only get their 'ultimate meaning' in their human context of use, and are always ambiguous. Such *ambiguities* are about shades of differences in meaning. The extent to and way in which ambiguities need to be resolved, depend on the context, including the points of view of the communicating agents, i.e. utterer and interlocutor, their common understanding of each other, and their (partially) shared goals [18].

But how to decide which ambiguities need to be resolved? A semantic approach, even when accepting different sources of meaning (i.e., ontologies), does not explicitly acknowledge the *consequences* of semantic choices. A pragmatic approach, on the other hand, assumes there are always conditions of difference, dependence, and novelty, and recognizes the need for an overall process for transforming existing knowledge to deal with negative consequences for community members [3]. We

[3] http://www.cyc.com/

would argue that, in addition, the community should also examine the *positive* consequences, such as opportunities for action.

In a pragmatic approach, control over representation should shift from the information producer to the information consumer [22]. More precisely, we think control over how to *use* meaning representations should shift to the user, from which controlling representations follows.

The need to accept a necessary amount of ambiguity by communities of users assessing the consequences of semantic choices in a particular pragmatic context, implies that there needs to be some user-controlled *selection* process of semantic representations. In such a process, members of the community, using the knowledge for a particular purpose, are actively involved, and aim to reach agreement only on *relevant* knowledge issues. Pragmatically established changes in the implicit meaning of representations should in the end also lead to changes the *representation* of those meanings in ontologies. For instance, if users always ask for concepts that are not, or only insufficiently, described in an ontology, it may be worthwhile to add this concept to the ontology. Meaning selection and representation processes, however, do not occur in isolation, but are driven by a meaning *negotiation* process in a specific community of users. In such a process, stakeholders arrive at the requisite (as determined by their shared goals) amount of agreement on shared concepts.

3. A Conceptual Model of The Web

Summarizing the previous discussion, we consider 'The Web' to consist of a Syntactic, a Semantic, and a Pragmatic web (Fig.1).

The *Syntactic Web* consists of interrelated syntactic information resources, such as documents and web pages linked by HTML references. These resources describe many different domains.

The *Semantic Web* consists of a collection of semantic resources about the Syntactic Web, mainly in the form of ontologies. The ontologies contain semantic networks of concepts, relations, and rules that define the meaning of particular information resources.

The *Pragmatic Web* consists of a set of pragmatic contexts of semantic resources. We consider a *pragmatic context* to consist of a common context and a set of individual contexts. A *common context* is defined by the common concepts and conceptual definitions of interest to a community, the communicative interactions in which these concepts are defined and used, and a set of common context parameters (relevant properties of concepts, joint goals, communicative situation, and so on). Each community member also has an *individual context*, consisting of individual concepts and definitions of interest and individual context parameters. Common and individual context parameters are not discussed further in this paper, as we will focus on the meaning negotiation process in which these contexts play a role.

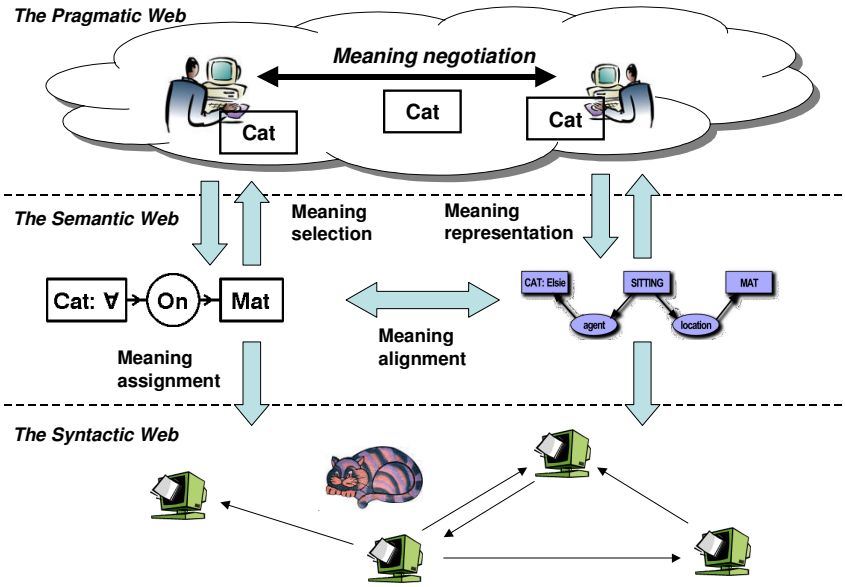

Fig. 1. A Conceptual Model of 'The Web'

Meaning plays a central role in connecting the various Webs. *Meaning assignment* takes place when syntactic resources are semantically enriched, such as by XML-tags being added to HTML-pages. *Meaning alignment* has to do with interoperability between ontologies: to what extent do their semantic models agree? How can (parts of) ontologies be meaningfully linked? How to deal with definitions that partially overlap in meaning? Much recent work addresses these – very hard – issues, e.g. [24,4]. Such meaning alignment problems mostly focus on modeling representational and evolutionary aspects of ontologies. However, as we have seen what needs separate attention are issues of ontology *use*. In other words: how can the process of *meaning negotiation* be improved? Meanings evolve not in the ontologies themselves, but in the pragmatic contexts where they are being used. Thus, a strong involvement of the community in ontology engineering processes is required, ensuring that individual and community changes in meaning are represented adequately in the ontologies.

3.1 The Complexity of Contexts

Our conceptual model allows us to examine a wide range of pragmatic contexts in the real world, and to identify commonalities and differences in problems with modelling, sharing and (re)using semantic resources such as ontologies.

Note that the sheer number of elements to analyze decreases as we move from the Syntactic to the Semantic Web, but strongly increases again when moving from the Semantic to the Pragmatic Web. There may be many thousands of (syntactic) information resources for a particular domain. In general, there will be many fewer ontologies defining the meanings of those resources. However, of pragmatic contexts

there can be an infinite number. There are many dimensions of pragmatics to be taken into account, such as purposes, communicative situations, organizational norms, individual values, and so on. These contextual parameters lead to a great variety of contexts. The multiple pragmatic contexts are even harder to formalize and standardize than the semantics of the concepts they interpret. Individual context views may agree with each other, or differ. Community members may use different ontologies to define the meaning of a particular concept. Many concepts have rich tacit meanings for individuals that can, nor should, always be made explicit in collaborative situations [21]. To assess the consequences of meaning choices, fully-automated negotiation processes will therefore never be sufficient. *Augmentation*, not automation of human meaning negotiation processes is required, in the sense proposed by Doug Engelbart [27].

One strategy to deal with this pragmatic complexity is to only model those pragmatic constructs that are essential to reach joint objectives. The meaning negotiation process should be a consensus seeking process, balancing individual and common requirements. Different individual views on the meaning of common concepts should be allowed, as long as they do not endanger the quality of the communicative interaction. For example, in a business transaction, it is essential that both parties have the same view of crucial parts of the definition of their contract, such as legal obligations. Where and how to store copies of the contract internally does not need to be part of a common meaning, however, and can thus be left as a degree of freedom. If differences in meaning are inhibiting the accomplishment of common goals, however, meaning negotiation has to proceed until the necessary amount of consensus has been reached.

How to proceed? What is a scalable way to operationalize such a pragmatic approach? If pragmatic contexts are unique and very different, how to systematically support meaning negotiation and related processes like meaning selection and representation? What is a requisite amount of consensus? The approach we propose in this paper is to base meaning negotiation on a set of fundamental pragmatic patterns, which can made available in a *meta-ontology*. These formal patterns can be used to define pragmatic constraints on processes in which explicit meanings are being defined and applied in contexts of use. Such an approach can help to better understand the potential uses and limitations of particular ontology engineering efforts, by clarifying the 'meaning of those meanings' for particular contexts of use.

4. Pragmatic Patterns

In [6], we presented a method for collaboratory improvement. Collaboratories are evolving socio-technical systems of people and tools aimed at providing environments for effective and efficient collaboration. About collaboratories often only partial knowledge of different degrees of specificity is or can be represented. The method uses ontology-grounded *improvement patterns* to capture various levels of socio-technical context knowledge about information and communication processes in collaboratories, including knowledge about workflows, design processes and

improvement processes. We view collaboratory improvement as a Peircean pragmatic inquiry process in which hypotheses about socio-technical improvements of the collaboratory are continuously constructed and tested in the community. This process, properly supported, should lead to more effective and efficient collaboratory evolution. Such an inquiry process could be a major driver of meaning selection in a community and hence form an important constituent of the Pragmatic Web [7].

A collaboratory improvement process is a good example of a community using patterns to evolve specification knowledge about its own socio-technical system. In the current paper, we want to develop a broader perspective. Instead of using patterns just to improve collaboratories, we intend to use patterns to 'improve semantics'. Given our conceptual model of the Web, what kind of patterns do we need? How do we represent them? How can we use them to deal with some of the problems inhibiting the progress of the Semantic Web?

4.1 Patterns

Humans use patterns to order the world and make sense of things in complex situations [15]. Patterns are often used in the construction of complex systems. An influential definition of patterns in architecture, also useful for information systems, was given by Christopher Alexander: "A pattern is a careful description of a perennial solution to a recurring problem within a building context, describing one of the configurations which brings life to a building (Alexander, et al., 1977, in [23]". A pattern thus contains elements of a solution to a problem, and applies within a particular context. Important is to focus on the words *recurring* problem and *perennial* solutions, indicating that the pattern definition of problems and solutions must be generic enough to cover a range of problem situations which in reality are always subtly different from the ideal, while being specific enough to offer useful solutions for the particular problem at hand.

Patterns are another view on domain models stored in ontologies. Developing ontologies for open environments like the Semantic Web is difficult, since more rules make ontologies less generic, while light-weight ontologies are not very useful [30]. This problem of finding the right degree of semantic specificity of ontologies to address problems in the domains they were created for, is not going to go away. The problem is not technical, but philosophical. If the types and number of applications of an ontology are infinite, and cannot be known beforehand, it will not work to try and produce the 'ultimate ontology' of semantic patterns. The usefulness of an ontology is always in the eye of the beholder, or more precisely, the eyes of many beholders: the many communities and individuals within communities using the ontology for their particular, changing, collaborative purposes.

Accepting this reality of eternal semantic partiality, conflict, and confusion, there is another, potentially more rewarding way to go. It consists of (1) making a strict conceptual separation between *modelling* and *using* ontologies, (2) identifying meta-patterns, i.e. *pragmatic patterns* that can (3) be used in *meaning evolution* processes in communities of users in order to make existing ontologies more useful and easier to

change[4]. These processes include what we referred to in the previous section as meaning representation, assignment, selection, alignment, and negotiation. Only by tackling these pragmatic issues head-on can the vision of the Semantic Web assisting the evolution of human knowledge as a whole [1], be realized in practice.

4.2 Core Pragmatic Patterns

To operationalize our vision of the Pragmatic Web, we need some core pragmatic patterns. We do not formalize the patterns in this article, but will outline some and describe their possible role in the scenario presented in the next section. Using conceptual graphs, it should be relatively easy to structure and reason about their (meta)semantics.

For a particular community, core pragmatic patterns include:

- *Pragmatic context:* a pattern that defines the speakers, hearers, type of communication, and identifiers of the individual and common contexts of a community.
- *Individual context:* a pattern that defines an individual community member, individual context parameters and an identifier of the individual context ontology.
- *Common context:* a pattern that defines the common context parameters and an identifier of the common context ontology of a community.
- *Individual pragmatic pattern*: a meaning pattern relevant to an individual community member. An individual context ontology consists of the total set of meaning patterns relevant to that individual.
- *Common pragmatic pattern*: a meaning pattern relevant to the community as a whole. The common context ontology consists of the total set of common meaning patterns relevant to the community.

Pragmatic patterns are template definitions that can be used as the basis of conceptual definitions used in meaning negotiation and other meaning evolution processes. These patterns can be refined and extended by communities if and when necessary.

Pragmatic patterns have a normative status, being either required, permitted, or forbidden. In the case of a pattern being required, this implies that the pattern must be satisfied in the process where it is used. If it is forbidden, it may not be matched in such a process. If permitted, it may be applicable, but not necessarily so. Such normative matching processes can provide powerful guidance of meaning evolution processes.

Earlier, we said that there is a much larger number and diversity of pragmatic contexts than of the ontologies which they use. Still, the number of pragmatic patterns, if chosen at the right level of specificity, can be relatively small. These patterns should

[4] These processes concern the evolution of *explicated* meanings. Many meanings are implicit, in people's heads. Although they may, and probably should change as well, understanding these are more the focus of, for instance, psychological studies.

not include the infinite number of details that make each pragmatic context unique, but only those that contribute to improving the effectiveness and efficiency of meaning evolution, with a focus on meaning negotiation. Of course, in this paper, we do not claim to solve the pragmatic puzzle. We will not provide the ultimate reference set of pragmatic patterns to be used in optimizing meaning evolution on the Semantic Web. Our aim is much more modest: showing proof of principle about what pragmatic patterns are and the role they could play in dealing with some of the meaning evolution issues mentioned. To this purpose, we introduce a hypothetical case very relevant to the conceptual graphs community: getting the famous cat its mat.

5. Using Pragmatic Patterns: How to Get a Mat for the Cat?

The mat producing company MatMakers wants to explore new markets. The grapevine has it that an interesting niche exists of cat lovers wanting nothing but the best for their furry friends. Its marketing officer Charles is commissioned to find new customers who will appreciate MatMaker's high-quality mats for their cats.

Charles decides to look for potential customers using the WYO=WTW (WhatYouOffer-is-WhatTheyWant) e-business broker. This broker is a web service that maximizes precision of advertising by using the latest Pragmatic Web-technologies. In particular, it mediates in meaning negotiation between sellers and prospective buyers by intelligent use of pragmatic pattern matching. The following type hierarchy is part of the WYO=WTW community context ontology (Fig.2):

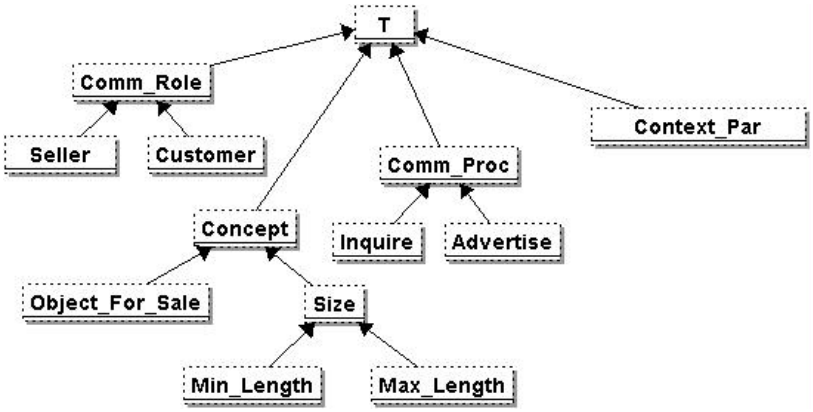

Fig. 2. The WYO=WTW Community Context Ontology

A most relevant concept in any advertising process is the *object for sale*. One important property of these objects, which is often discussed in the business negotiations of this particular community, is the *size* of the object being offered. Two important size indicators are the *minimum* and the *maximum length* of the object. Two communication roles in an e-business transaction are the *seller* and the *consumer*,

referring to the parties who can play the speaker or hearer-roles. The community using the WYO-WTW service distinguishes two types of communication processes: *inquiring* about objects for sale, initiated by customers, and *advertising* objects, initiated by producers.

MatMakers has its own corporate ontology, from which Charles imports the Mat and Size-concepts (including their positions in the type hierarchy) into the individual context ontology of MatMakers for the WYO=WTW service. He also adds the Cat-concept, since that is what he wants to focus his particular potential customer search on. Since the maximum length of the mats produced by MatMakers is one meter, Charles adds to his individual context ontology the required pattern that to be of interest for an advertisement any cat for which the mat is bought should be at most one meter long (Fig.3):

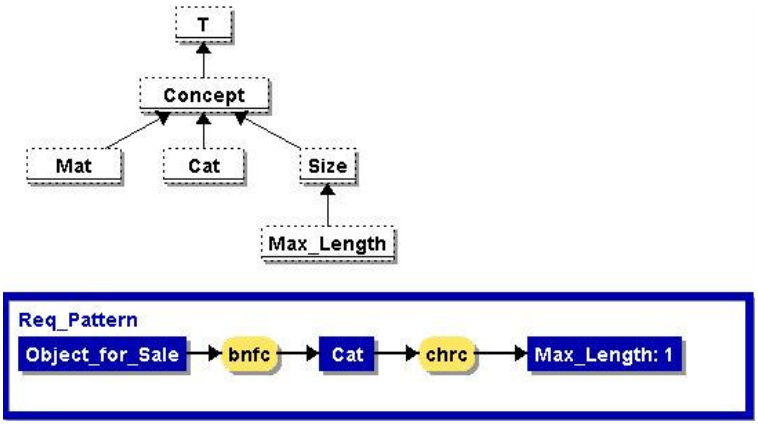

Fig. 3. The Individual Context Ontology of MatMakers for the WYO-WTW Service

The Cat-Lovers-Association-of-the-World (CLAW) is a worldwide virtual community of amateurs crazy about cats. They have interest groups studying not only small cats, like street cats and Siamese cats, but also large cats, like lions and tigers. The database of member addresses of such a highly motivated global community is of high potential value to corporations. In principle, CLAW is not adverse to their members being offered products for their pets. However, they are not interested in offers of products for large cats, since their members are amateurs only, not zoo owners. Therefore, they demand that any sales offer in an advertisement concerns small cats only (Fig.4):

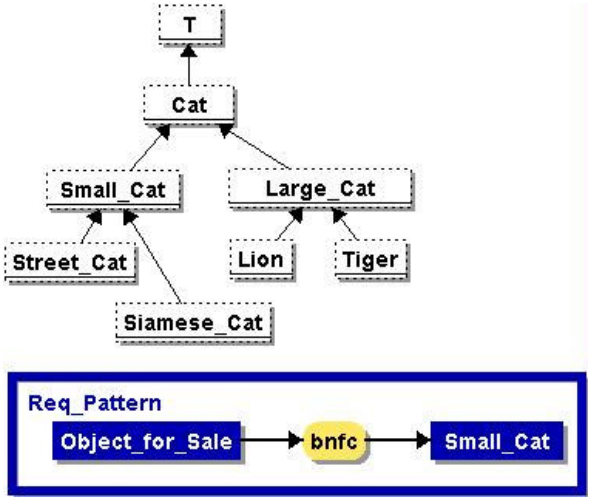

Fig. 4. The Individual Context Ontology of CLAW for the WYO-WTW Service

To find potential customers, Charles first sends a query identical to his required pattern to the web service (Fig.5):

Fig. 5. The Initial Query

WYO-WTW searches the individual context ontologies of all registered members of its services for a match with this required pattern, by projecting the required pattern on the individual context ontologies of the various members[5]. Nothing matches. Charles realizes that his query could have been too specific, not because no customers share his interest, but because their meanings have not yet been sufficiently specified in the ontologies they use with respect to Charles' purpose. He decides to relax the query by only looking for potential customers who are interested in products for cats, and try to find out about the length of their animals later. He therefore sends the following generalization of his required pattern (Fig.6):

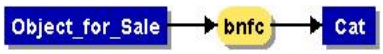

Fig. 6. The Generalized Query

[5] We do not go into the details of dealing with labels like Req_Pattern here, these can be dealt with by syntactically rewriting the queries.

WYO-WTW projects this generalization again on the various individual context ontologies. It now matches with all (i.e. the only) required pattern of the CLAW context ontology, returning the following result (Fig.7):

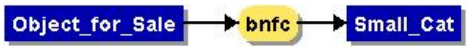

Fig. 7. The Result of the Generalized Query

CLAW's only (and all) pragmatic requirements on any seller have now been satisfied, and the association in principle is open to being sent the advertisement. However, MatMaker's own required pattern has not been satisfied yet. To see if it could, WYO-WTW goes on the Semantic Web, projecting MatMakers *core semantic pattern* (i.e. the essence of MatMakers' required pattern adapted to the semantic constraints of the potential customer party's required patterns) on the public interfaces of various ontologies that are in CLAW's list of *trusted semantic resources*[6]. It does this to see if these ontologies can be useful in enriching CLAW's ontology sufficiently for it to match with MatMaker's required pattern. The core semantic pattern in this case is the part that follows the bnfc-relation, since this indicates what MatMakers requires from its customers for them to be eligible candidates for advertisement. The Cat-concept is thereby specialized to Small_Cat, since that specialization is demanded by CLAW's required pattern. Furthermore, any instances are left out, since values may have to be calculated by inference rules, instead of being stored directly in ontologies. This would lead to queries failing, even though semantically they should match with an ontology. Thus, the WYO-WTW service sends out the following core semantic pattern query to trusted ontologies on the Semantic Web (Fig.8):

Fig. 8. The Core Semantic Pattern Query to the Semantic Web

Again nothing matches. WYO-WTW now automatically starts to look for similar concepts. It first tries to find synonyms for the *Small_Cat*-label by contacting the Cyc-URI (Uniform Resource Indicator) service. This Semantic Web-service finds *Felinae* as a synonym. It resends the query, but this time with *Small_Cat* replaced by *Felinae*. It turns out that this query matches with the ontology of Animal Diversity Web (ADW), a university zoological taxonomy server[7]. A part of this ontology is the following (Fig.9):

---

[6] If a semantic resource like an ontology is trusted, a user is willing to accept the ontological commitments implied by the definitions of that resource.

[7] In fact, the taxonomy server on which this (hypothetical) ontology is based really exists. It is hosted by the University of Michigan: http://animaldiversity.ummz.umich.edu/site/. Such relatively stable reference ontologies could play an important role in optimizing meaning negotiation processes on the future Pragmatic Web.

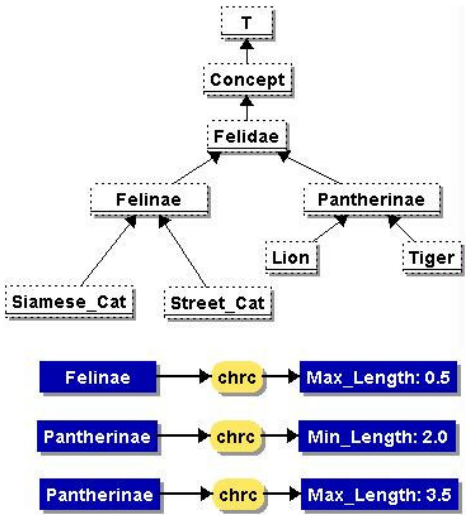

Fig. 9. ADW's Ontology on the Semantic Web

The fact in ADW's ontology that matches with the query (i.e. is a specialization) is the following (Fig.10):

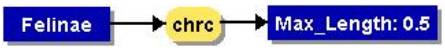

Fig. 10. The Result of the Revised Core Semantic Pattern Query

The following fact is automatically added to the WYO-WTW common context ontology (since *Felinae* is equivalent to the *Small_Cat*-label, and the latter is the terminology used by at least one of the community members) (Fig.11):

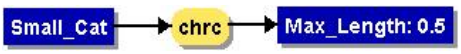

Fig. 11. The Common Pragmatic Pattern

This *common pragmatic pattern* forms the basis for starting the actual advertising process. It means that seller and customer share an interest in beginning an advertising process about objects for sale for small cats which have a maximum length of half a meter. The pattern is - necessarily - a specialization of the parts of the required patterns of both communicating parties that define the properties of the beneficiary of the object for sale.

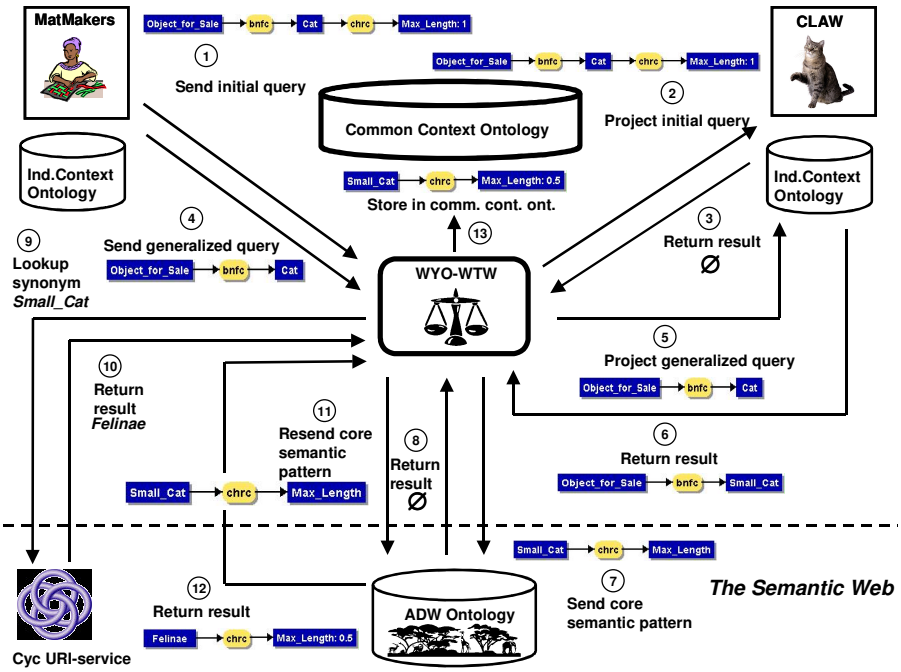

Fig 12. Meaning Negotiation on the Pragmatic Web

Fig. 12 summarizes the meaning negotiation process. This already complex scenario was just a simplified example of a realistic meaning negotiation process. Still, it should demonstrate the power of a combination of a conceptually clearly separated, yet interdependent Semantic and Pragmatic Web.

6. Discussion

The purpose of this paper was to expand current thinking on the Pragmatic Web by identifying some issues and presenting a sketch of one possible approach for its operationalization. There are many directions in which this work should be expanded, however. For example:

- The Pragmatic Web and the Semantic Web are strongly interdependent. Many open issues of the Semantic Web still need to be resolved, before a robust Pragmatic Web can be constructed. One example concerns useful and widely adopted URI (Uniform Resource Identifier) schemes. Still, it would already be very useful to systematically examine current and projected components of the Semantic Web through a pragmatic lens, in order to discover new applications. Vice versa, insights about the Pragmatic Web may help address some of the thorny issues currently blocking progress in the Semantic Web community.

Many semantic approaches, for example, already have pragmatic components. However, semantics and pragmatics are often mixed up in confusing ways. Our approach could help disentangle some of these conceptual knots, allowing for optimizations with the right focus, e.g. with respect to either modelers' or users' needs in particular cases.

- We have mostly stuck to a rather practical and shallow interpretation of pragmatics. Philosophically, pragmatics is a very complex idea, however. Insights from philosophers with a strong focus on evolution of meaning, such as Peirce's pragmaticism [2] and Habermas' theory of communicative action [11] could be very useful in strengthening the theory of the Pragmatic Web.
- We only defined meaning negotiation and selection processes informally as sequences of graph projections. The (meta)-semantics of the various pragmatic patterns is still quite fuzzy. How to formalize individual and common contexts and pragmatic patterns? What role should they play in the various meaning evolution processes? When and how should recurring pragmatic patterns stored in meta-ontologies be included in domain ontologies on the Semantic Web? Conceptual graphs research can also make important contributions here, both in terms of advanced theoretical research like context modelling [20] and architectures for pragmatic graph application systems [28,5]. Also, the normative status of patterns is a complex issue. In the scenario, we only used required patterns. In realistic applications, these may conflict with prohibited and permitted patterns. Deontic logic is one theoretical field that help clarify some of these issues [19].
- Human communication is crucial in meaning negotiation on the Pragmatic Web. Conceptual approaches such as proposed in this paper can only augment, not automate human meaning interpretation and negotiation processes. A theoretical foundation for modelling more complex and realistic communicative interactions is the Language/Action Perspective, which stresses the coordinating role of language. This perspective has led to various proposals for human/agent communication-based collaborative models and systems e.g. [16,28,31,12]. Another rich source of ideas for designing pragmatic systems supporting human communication is (business) negotiation theory (e.g. [8]).
- Ontologies play a crucial role, at both the Semantic and Pragmatic Web levels. The ontologies presented in the scenario were exceedingly simple, since the focus was on proof of concept, not on the finer semantic details. Much ontology research focuses on these representation and reasoning issues. Although valuable and necessary, ontology research on the Pragmatic Web level should also focus much more on ontology *methodology* issues. These include the (partially) human processes of modeling, selecting, using and changing meanings for collaborative purposes. The DOGMA-methodology being developed at STARLab consists of a set of methods, techniques, and tools to arrive at scalable ontologies that can actually be made useful in practice [29]. One of our projects in which a strong focus will be on exploring the relations between ontologies and pragmatics is the CODRIVE project[8]. Our aim is to develop a methodology for negotiating a

[8] EU Leonardo da Vinci project BE/04/B/F/PP-144.339

common competence ontology by key stakeholders in the European labor market. These parties, representing the educational sector, public employment agencies, and industry have a need for a common competency ontology that can be used for collaborative applications such as doing job matches and developing individual training pathways. Given the widely varying interests and definitions of competence concepts, this should be a very interesting test case to further develop theory and practice of the Pragmatic Web.

7. Conclusion

The Pragmatic Web is the next phase in the evolution of the Web. Most research attention currently focuses on the Semantic Web. However, for the Semantic Web to truly realize its potential, much more work needs to be done on its pragmatics aspects. This entails that the context of use of explicated meanings that are stored in ontologies need to be much better understood. The driver of the Pragmatic Web are meaning negotiation processes. These processes are connected to the Semantic Web by meaning selection and representation processes.

In this paper, we have explored the contours and some fundamental concepts of the Pragmatic Web. By means of a scenario we have explored what the Pragmatic Web in a few years time might look in practice. The aim of this paper was not to solve existing problems, but to help open up an exciting new territory for intellectual and practical exploration. Moving the research focus from semantics to pragmatics, from representing to *using* meaning, is the next step on the way to network applications that help communities of people realize their full collaborative potential.

References

1. Berners-Lee, T., Hendler, J. and Lassila, O. (2001), The Semantic Web, *Scientific American*, May 2001: 35-43.
2. Buchler, J. (1955), *Philosophical Writings of Peirce*. Dover Publ., New York.
3. Carlile, P. R. (2002), A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development, *Organization Science,* 13(4): 442-455.
4. Corbett, D. (2004), Interoperability of Ontologies Using Conceptual Graph Theory. In *Proc. of the 12th Intl. Conference on Conceptual Structures (ICCS 2004), Huntsville, AL, USA, July 2004*, LNAI 3127. Springer, Berlin, pp. 375-387.
5. Delugach, H. S. (2003), Towards Building Active Knowledge Systems With Conceptual Graphs, in *Proc. of* the *11th Intl. Conf. on Conceptual Structures (ICCS 2003), Dresden, Germany, July 2003*, LNAI 2746. Springer, Berlin, pp. 296-308
6. de Moor, A. (2004), Improving the Testbed Development Process in Collaboratories. In *Proc. of the 12th Intl. Conference on Conceptual Structures (ICCS 2004), Huntsville, AL, USA, July 2004*, LNAI 3127. Springer, Berlin, pp. 261-274.
7. de Moor, A., Keeler, M. and Richmond, G. (2002), Towards a Pragmatic Web. In *Proc. of the 10th Intl. Conference On Conceptual Structures (ICCS 2002), Borovets, Bulgaria, July 2002*, LNAI 2393. Springer, Berlin, pp. 235-249.

8. de Moor, A. and Weigand, H. (2004), Business Negotiation Support: Theory and Practice, *International Negotiation*, 9(1):31-57.
9. Fillies, C., Wood-Albrecht, G. and Weichhardt, F. (2003), Pragmatic Applications of the Semantic Web Using SemTalk, *Computer Networks,* 42: 599-615.
10. Gruber, T. (1994), Towards Principles for the Design of Ontologies Used for Knowledge Sharing. In N. Guarino and R. Poli (eds.) *Formal Ontology in Conceptual Analysis and Knowledge Representation*. Kluwer.
11. Habermas, J. (1981) *Theorie des kommunikativen Handelns* (2 vols.). Suhrkamp, Frankfurt.
12. Harper, L.W., and Delugach, H.S. (2004), Using Conceptual Graphs to Represent Agent Semantic Constituents. In *Proc. Of the 12th Intl. Conference on Conceptual Structures (ICCS 2004), Huntsville, AL, USA, July 2004,* LNAI 3127. Springer, Berlin pp. 325-338.
13. Kim, H. and Dong, A. (2002), Pragmatics of the Semantic Web. In *Semantic Web Workshop 2002, Hawaii*.
14. Kreuz, R. J. and Roberts, R. M. (1993), When Collaboration Fails: Consequences of Pragmatic Errors in Conversation, *Journal of Pragmatics,* 19: 239-252.
15. Kurtz, C. F. and Snowden, D. J. (2003), The New Dynamics of Strategy: Sense-Making in a Complex and Complicated World, *IBM Systems Journal,* 42(3): 462-483.
16. McCarthy, J. (1996), Elephant 2000: A Programming Language Based on Speech Acts, Technical Report, Stanford University.
17. McLaughlin, W. S. (2003), The Use of the Internet for Political Action by Non-State Dissident Actors in the Middle East, *First Monday,* 8(11).
18. Mey, J. L. (2003), Context and (Dis)ambiguity: a Pragmatic View, *Journal of Pragmatics,* 35: 331-347.
19. Meyer, J. J.-C. and Wieringa, R., eds. (1993), *Deontic Logic in Computer Science: Normative System Specification*. John Wiley & Sons, Chichester.
20. Mineau, G.W. and Gerbe, O. (1997). Contexts: A Formal Definition of Worlds of Assertions. In *Proc. of the 5th Intl. Conference on Conceptual Structures (ICCS 1997), Seattle, Washington, USA, August 1997*, LNCS 1257. Springer, Berlin, pp.80-94.
21. Nonaka, I., Reinmoeller, P. and Senoo, D. (1998), The 'ART' of Knowledge: Systems to Capitalize on Market Knowledge, *European Management Journal,* 16(6): 673-684.
22. Repenning, A. and Sullivan, J. (2003), The Pragmatic Web: Agent-Based Multimodal Web Interaction with no Browser in Sight. In *Human-Computer Interaction - INTERACT'03*. IOS Press, IFIP, pp. 212-219.
23. Schuler, D. (2002), A Pattern Language for Living Communication. In *Participatory Design Conference (PDC'02), Malmo, Sweden, June 2002*.
24. Shanks, G., Tansley, E. and Weber, R. (2003), Using Ontology to Validate Conceptual Models, *Communications of the ACM,* 46(10): 85-89.
25. Singh, M. P. (2002a), The Pragmatic Web, *IEEE Internet Computing,* May/June: 4-5.
26. Singh, M. P. (2002b), The Pragmatic Web: Preliminary Thoughts. In *Proc. of the NSF-EU Workshop on Database and Information Systems Research for Semantic Web and Enterprises, April 3-5, Amicalolo Falls and State Park, Georgia*.
27. Skagestad, P. (1993), Thinking with Machines: Intelligence Automation, Evolutionary Epistemology, and Semiotic, *Journal of Social and Evolutionary Systems,* 16(2): 157-180.
28. Sowa, J. (2002), Architectures for Intelligent Systems, *IBM Syst. Journal*, 41(3):331-349.
29. Spyns, P. and Meersman, R. A. (2003), From Knowledge to Interaction: from the Semantic to the Pragmatic Web. Technical Report STAR-2003-05, STARLab , Brussels.
30. Spyns, P., Meersman, R. A. and Jarrar, M. (2002), Data Modelling versus Ontology Engineering, *ACM SIGMOD Record,* 31(4): 12-17.
31. Weigand, H. and de Moor, A. (2003), Workflow Analysis with Communication Norms, *Data & Knowledge Engineering*, 47(3):349-369.