Is Your Ontology a Burden or a Gem? – Towards Xtreme Ontology Engineering

Olga Tatarintseva[1], Vadim Ermolayev[1], Anna Fensel[2, 3]

[1] Department of IT, Zaporozhye National University,
66 Zhukovskogo st., 69600 Zaporozhye, Ukraine
tatarintseva@znu.edu.ua, vadim@ermolayev.com
[2] FTW Forschungszentrum Telekommunikation Wien GmbH,
Donau-City-Straße 1, A-1220 Vienna, Austria
[3] STI Innsbruck, Technikerstr. 21a, A-6020 Innsbruck, Austria
anna.fensel@sti2.at

Abstract. One of the commonly acknowledged shortcomings of Semantic Technologies that prevents their wide adoption in industry is the lack of the commitment by the intended domain experts and users. This shortcoming becomes even more influential in the domains that change sporadically and require appropriate changes in the respective knowledge representations. This discussion paper argues that a more active involvement of the intended user community, comprising subject experts in the domain, may substantially ease gaining the required commitment of the critical mass of the domain users to the developed domain ontology. As a possible approach for building an instrumental platform for that, the paper suggests the use of the Semantic MediaWiki based collaboration infrastructure for maintaining and discussing ontology descriptions by the community of its intended users and developers. We also report how a prototypical ontology documentation wiki has been used for gaining the commitment of ontology users in the ACTIVE European project.

Keywords: ontology engineering methodology, stakeholder commitment, OntoDocWiki, xtreme ontology engineering.

1 Introduction

Building and refining practically useful knowledge representations in different domains is one of the major challenges in making semantic technologies publicly accepted today. The problem is not only in creating the proper encodings of the tacit knowledge of subject experts, or user behavior observations [1] but also in gaining a commitment to the developed ontologies by the users who are supposed to exploit these modules of formalized and explicit knowledge – either directly in their daily activities or by empowering their software tools. Evidently creating ontologies is not a routine task. It requires substantial intellectual effort.

Moreover, refining ontologies, making them better covering the intended requirements of the user community, is even more challenging and effort consuming. As well known from knowledge elicitation practice, *five* subject experts will most definitely have *seven*[1] different opinions. So, the commitment of those individuals can only be reached if a knowledge representation is aligned and harmonized alone their subjective and tacit interpretations of the domain knowledge.

Gaining a commitment to the ontology by a wider group or a community of intended users is even more complex. One reason is that the majority of users have to adapt themselves not only to the suggested knowledge representation but also to the knowledge carried by this formal representation – which could both be novel to them. In our opinion the difficulty of gaining users' commitment is the major obstacle for a broader adoption of the semantic technologies in industries and the reason[2] for the criticism expressed to those technologies. The problem becomes even more challenging for the knowledge about the domains that change frequently. Ontologies describing those domains have to be changed accordingly. The changes in the knowledge representations have to be accepted by the subject experts and the users.

Hence, offering support for facilitating a better and less effort consuming comprehension, alignment and harmonization of knowledge representations by a user community may become a substantial step forward in reusing domain knowledge by knowledge workers and their software systems. Contemporary ontology engineering methodologies put insufficient emphasis on offering ways to gaining such a commitment. The analysis of this shortcoming is given in Section 3.

We believe that a more active involvement of the intended ontology users in the processes of ontology development and refinement is required for lowering their comprehension barriers. A software tool facilitating this active involvement will inevitably be a collaboration platform that allows discussing knowledge representations and expressing opinions and arguments by any community member.

Developing tools for collaborative knowledge engineering and knowledge reuse is one of the mainstreams in the semantic technologies community. However, the vast majority of the tools available today are tailored to the use by knowledge engineers, but not by domain experts or users. One interesting exception is the development of ontology games and collaborative (social) semantic mark-up tools for Web 2.0. Yet, these approaches yield too lightweight models – insufficiently expressive for the majority of industrial applications. The analysis of the state of the art in collaborative platforms for ontology engineering, onto-gaming and semantic mark-up on the Social Web is given in Section 4.

One possible solution for the outlined problem is making these divergent courses meet. Tool support for ontology engineering would benefit from adopting "croud-sourcing" features of collaborative knowledge representation development by Social Web users. A meeting point that will allow for the proper comprehension of knowledge structures is a collaborative platform for presenting and discussing the documentation of the ontologies by the subject experts and intended users alone the development process. There are several obstacles on this way. One is the lack of a proper incentive mechanism motivating subjects to take their active part. Another one lies in the nature of the work to be done – it is out of the scope of the core professional competence of intended audience. Yet more obstacles are caused by the lack of the tool support for: (i) the development and versioning of the ontology

[1] The numbers are indicative.
[2] … apart of the incurred computational overhead.

documentation in line with the evolution of the ontology; (ii) the discussion of ontology documentation as a representation of knowledge that is more easily comprehendible by the users than the code of the ontology. Last but not least is the need for a mechanism of reaching and spreading consensus among the participants. The requirements to the envisioned collaborative platform are presented in Section 5.

Our experience in developing and experimenting with a prototype collaboration platform for involving domain specialists in the active discussion of domain knowledge representations is presented in Section 6. The prototype platform implements some of the outlined requirements. It is based on the Semantic MediaWiki [2] with an extension for moderated discussions. The prototype platform has been used in ACTIVE Project (http://active-project.eu/) for representing and discussing the PSI Suite of Ontologies (http://isrg.kit.znu.edu.ua/ontodocwiki/) describing projects and processes in microelectronic engineering design.

2 Why is an Ontology Often a Burden?

Given the effort and intellect invested into the development of ontologies as consensual descriptive formal models of domains of discourse [3] in the last two decades, it could have been expected that ontologies had already become the core enablers for the ICT infrastructures and services in many industrial branches. However this is not entirely the case. “Unfortunately, the number and quality of actual, “non-toy” ontologies available on the Web today is remarkably low” [4]. Several technical limitations and practical challenges preventing easy adoption remain unsolved. These barriers for technology and methodology uptake *as perceived by industries*[3] (c.f. [4]) are as follows:

(i) **Unjustified benefits**. Industrial users and policy makers tend to think that using ontologies in their industrial setting is an artificial requirement to a large extent. They (sometimes wrongly) assume that the required information could be presented and processed in a more easy way using lighter-weight mark-up languages (like XML) and corresponding parsers.

(ii) **Considerable effort expectations**. Industries consider that even if the need for an ontology is justified, the effort required for building it is too large to be acceptable for the incurred benefits of use.

(iii) **Insufficient expressivity and comprehension gap**. Industrial technical specialists fear that the expressivity of ontology specification languages is insufficient for fully and adequately describing their subject domain. Consequently, it can not be granted that the intended users of the ontology easily grasp the meaning of the ontology elements as intended by the knowledge engineers who created the code.

(iv) **Computational overhead and poor scaling**. Software developers in industry estimate that ontology based software solutions are too heavy-weight. The software spends too much computation power for ontology processing and

[3] This information has been acquired from a several year experience of industrial partnerships in knowledge intensive research and development projects.

reasoning because of the complexity of the problems that are solved. As a result ontology-based applications scale quite poorly to be acceptable in industrial settings. More lightweight solutions are demanded.

(v) **Insufficient maturity of ontology engineering process.** If it is not their core competence, industrial engineers will seek for and accept a methodology that is well defined and based on the use of standardized working patterns. However, the leading experts in the field state that the development of ontologies is still much more a craft work or a non-trivial mental exercise than a rigorous and standardized engineering process. One of the particular shortcomings is that ontology engineering cycles are too long. The result is that we can not build ontologies that adequately follow the changes in quickly evolving domains.

Last but not least, and as a consequence of the combination of the outlined barriers, people in industries hardly believe that ontologies will solve their practical problems and help effectively in the development of their applications. Therefore it is often difficult to obtain their commitment to the ontologies offered to or developed for them.

It may be also noticed that the requirements the industries implicitly put up-front in (i–v) are not properly balanced and are sometimes clashing. For example a desire to have easily comprehendible lightweight knowledge representations is contrary to the demand of more expressive power for a more adequate representation of a domain. One good approach to resolve those clashes is to offer a layered representational structure with a more coarse and easy to grasp descriptions on the top down to fully detailed and formally coded knowledge representation modules in the bottom. As suggested in [5] those layers may be offered as different representation facets to different categories of specialists and at different development phases.

We do not intend to resolve all the fears of industrial experts in this paper. Our objective is to evaluate the existing ontology engineering approaches, development praxis, and methodologies by looking at how they facilitate better and broader commitment to the developed knowledge representation artifacts and, by that, relax some of the existing barriers.

3 Shortcomings of Ontology Engineering Methodologies

Knowledge Engineering as a subfield of Artificial Intelligence or broader – of Computer Science is a vibrant research discipline for already more than two decades. It involves integrating knowledge into computer systems. Knowledge has to be therefore represented in a way that a computer system is able to process. These representations are often elaborated as ontologies using the instruments provided by Ontology Engineering. This discipline develops knowledge representation frameworks to ensure adequate rigor for making the outputs tractable and processible by machines. Formal knowledge representation languages (e.g. OWL 2.0 [6]) are developed and standardized for that purpose.

Ontology Engineering is also concerned about the development of the methods and methodologies for building ontologies to fulfill the requirements of the intended user audience. The results are sought to cover the user interpretations of the common sense or a target subject domain so much and completely as the expressivity of the formal representation allows. Several ontology engineering methodologies have been developed up to date. In particular METHONTOLOGY [7], DILIGENT [8], On-To-Knowledge [9], Uschold&King [10], Delphi [11], Compendium [12], HCOME [13], CommonKADS [14], NeOn [15] are the methodologies that are mentioned in the literature most frequently. A reader may be pointed to [16] for a more comprehensive list and analytical survey. Recently the methodologies taking into account the economical aspects of ontology engineering appear – e.g. OntoCOM [17].

Among those we are particularly interested in the methodologies which explicitly support: (i) collaborative ontology engineering; and (ii) ontology refinement process with evolving requirements. These methodologies are METHONTOLOGY, DILIGENT, On-To-Knowledge, Compendium and NeOn. In terms of ontology engineering lifecycle all the five methodologies suggest a variation of a process schematically pictured in Fig. 1, where the iterative parts are described in terms of either spiral or iterative waterfall process models.

The differences are:

(i) METHONTOLOGY distinguishes support and development activities and focuses on knowledge elicitation and result evaluation routine
(ii) DILIGENT focuses on distributed deployment and local changes and provides an argumentation framework for harmonization
(iii) On-to-Knowledge distinguishes ontology refinement as a separate important stage of the development process
(iv) NeOn offers a flexible scenario-based decision procedure for choosing the most appropriate lifecycle model and puts significant emphasis on the re-use of ontology patterns and available distributed ontologies
(v) Compendium explicitly concentrates on the ways of organizing collaborative work at knowledge elicitation and evolving prototyping phases

With respect to the relaxation of the barriers for gaining ontological commitment the contribution of all the mentioned methodologies is limited. NeOn suggests re-using good ontology engineering practices in the form of ontology patterns. Implicitly it suggests that using these good practices results in making the ontologies more correct and reliable – thus the commitment to these results is expected to be higher. Only OntoCOM elaborates the incentives for individuals and organizations for introducing ontologies. However it does not suggest mechanisms for gaining the commitment when the ontology is being developed. Compendium offers an approach to collaborative development based on moderated discussions and accounting for the evolution of the requirements. Unfortunately it does not mention the incorporation of the experts who carry the tacit knowledge about the subject domain that has to be elicited. None of the reviewed methodologies pays attention to the presentation of the developed knowledge representation in the form that is easily comprehendible by the intended users. The ontology documentation activity is considered only a support activity to the development process. None of the methodologies, except DILIGENT and NeOn to some extent provide the means for reaching consensus in ontology design decisions. DILIGENT does that by offering a harmonization framework. NeOn suggests consensual seeds in the form of reusable design patterns.

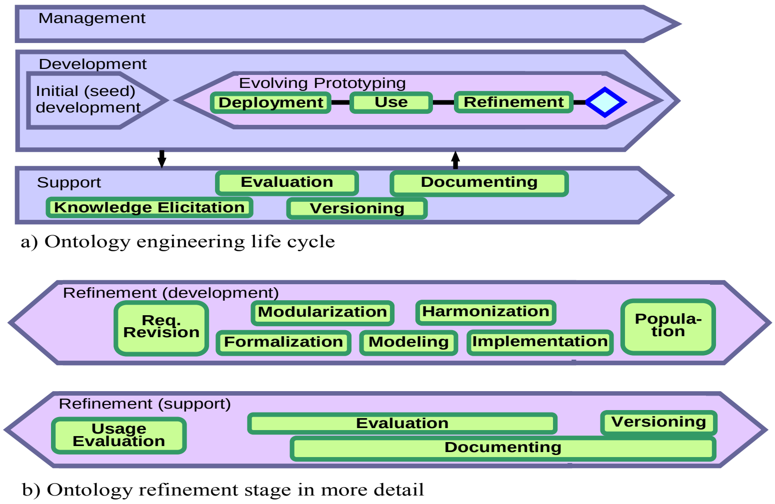

Fig. 1. Ontology engineering lifecycle specified in ISO/IEC 24744 graphical notation for describing methodologies [18].

4 Social Tagging and Games with a Purpose

One of the possible ways to check if the conceptualization of the domain is correct and complete is to evaluate the model against the outcome of users' grasp of the meaning of the content of the representative set of relevant documents. From the other hand, the labels or comments the users put on these documents or resources may be effectively used to infer the conceptualization. Such labels are often denoted as tags or annotations.

If web resources of different modalities are thought of as the representative set of data we find ourselves in the exploding field of collaborative or social tagging and annotation – a substantially characteristic part of the Web 2.0 phenomenon. A good survey of the field of social tagging is offered by Gupta et al in [19]. Tags created by the community of online users are exploited for different purposes. Taxonomy generation is one of the applications particularly relevant to our subject. The shortcomings of social tagging on Web 2.0 as analyzed in [19] as follows: (i) tags are simple bags of words without any more expressive semantics; (ii) tags are often not correct, especially if generated by spammers; (iii) tags are often ambiguous because different users apply terms to documents in different ways; tags are often sparse and do not cover the elements of the resource uniformly. Those shortcomings effect in low quality of tags. The reason is that the taggers do not use the terms of a consensual domain model in their activity.

Semantic annotation and tagging approaches further refine social tagging techniques by offering the collections of terms that are taken from such knowledge representations in the forms of taxonomies, folksonomies, thesauri. Please refer to

[20] for a comprehensive analyses of the requirements and results in the field. However, the backbone knowledge representations have to be obtained before semantic annotation may be undertaken. This remains the work for ontology engineers and is not regarded as a task for non-specialist users.

Hybrid approaches for collaborative tagging and annotation aimed at enrichment of the seed knowledge representations by the activity of the user community doing tags or annotations are also reported. For example [21] reports about the implementation of such a hybrid approach in Digital Libraries by tagging data through a combination of a standardized model, a harvesting protocol and a metadata mapping. It is concluded that both the custodians and users of digital repositories enabled with this collaborative annotation system benefit from the potential value of collaborative tagging without a need for a prior knowledge of the backend annotation systems. The weaknesses of traditional social tagging systems are attempted to be overcome by combining the best features of the Social and Semantic Webs. Unfortunately the problem of motivating users actively taking part in annotating resources remains open even in the reported advanced collaborative semantic tagging systems.

As already mentioned above, good ontologies should match consensual interpretations of domain knowledge by the representative set of domain users as closely and completely as possible. Therefore, ontology engineering in any form is by its very nature a process that has to involve as many domain experts and intended users as affordable. Increasing numbers of people are willing to spend their time using Web 2.0 applications, for example in adding tags and sharing their mark-ups in a group or community. Remarkably, it is totally on the contrary for ontology development – nobody reported about involving big user groups in creating knowledge representations so far. As recognized by many experts in Semantic Technologies one of the possible reasons is that traditional ontology engineering methodologies detach the effort from the benefits (c.f. [22]) hence de-motivating the involvement of those people whose interpretations of the domain are critically required.

One of the promising approaches for motivating more people take part in creating or refining ontologies is offering social software or, alternatively, a game with a purpose to a group of intended users. Several results in this direction are reported in the literature. Ontology creation can be implicitly embedded in social software, namely, social networking portals where users would be creating, evolving and confirming ontology items implicitly in the background, while simply providing information to a the portal for sharing content and communication with other users [23]. Games with a purpose is another approach that has been used mainly for collaborative tagging of resources having different modalities: images [24], music [25] – to mention just few. Gaming approach has also been tried for inferring human intentions from their recorded actions (Common Consensus game, [26]) and for evaluating how well commonsense facts fit to the interpretations of random users (FACTory Game by Cycorp, http://game.cyc.com/). For involving users in creating domain knowledge representations several game scenarios have been developed [22] for ontology building and refinement, ontology matching, annotating content using lightweight ontologies. Even though ontology backed up social networking portals and games with a purpose scenarios differ in relation to the users' motivation to contribute, they are in compliance to the OntoElect approach proposed in the next

Section. Both approaches offer possibilities to identify whenever users start to agree on and share certain ontological items

The positive features that are common to all these gems of related work are: (i) the pattern of user involvement adopted from Web 2.0 is used to motivate people taking part; (ii) all games with a purpose hide their purpose under the gaming scenario – offering fun in reward for providing useful results; (iii) the scenarios are designed in a way that assists in structuring the pool of players by their reputation.

There is also the shortcoming which is inherent to this approach – the knowledge representations or mark-ups that are crafted by non-specialist players can only be lightweight. Otherwise the overhead for ramping-up the players would consume all the offered incentives. Though ontology fragments are aligned with consensual user interpretations, it would be hard to ensure that the quality of those fragments is sufficient for, say, industrial use. Therefore, a joint motivated involvement of ontology engineering professionals, subject experts, and a sufficiently big group of intended users with domain knowledge and expertise is required.

5 OntoElect Approach for Xtreme Ontology Engineering

Charles Petrie in his editorial article [27] argued that the correctness of "… semantics, is evident in its use …". Ontologies are often denoted as descriptive theories that specify domain semantics – so they may only be validated by the users in their daily work in that domain. Martin Hepp in [4] backs up this view by stating that "… commitment can be achieved only by successful joint action – that is, successful usage of the ontology". Emphasizing the role of user commitment he observes that the perceived utility of the ontology grows with the number of users who commit to it. Hence, an effective ontology engineering methodology has to offer a mechanism for gaining the commitment by the intended group of users – the sooner – the better and as broadly as possible. We believe that a correct way to go is to involve the subject experts and the intended users in the development of the ontology at the earliest phase possible.

Our proposal of a possibly effective approach for attracting subject experts to play a more active role in the development and ownership of ontologies is inspired by the observations of social and political life – in *public election campaigns*. Indeed, the desired outcome of an election campaign for every candidate is to gain as much commitment of the electorate as possible. Such a commitment is measured adequately by the number of votes. The candidates compete for the votes by presenting their programmes, making coalitions, taking part in public debates – proving that his or her programme is the best match to the expectations of the majority of the electorate.

In the case of ontology engineering alternative ontology seeds for the same subject domain could be treated as election candidates. Each candidate ontology offered in a, so to say, *ontology election campaign* may be evaluated compared to the other candidates by the ability to answer the competency questions of the electorate. The more competence in answering the requirements of the electorate is demonstrated by the ontology, the more commitment it is potentially able to gain. The members of the electorate are the intended users in the domain. Their commitment could be measured in simple votes or using a more sophisticated scoring mechanism. A good example of such a mechanism is evaluating submissions in a peer review process. Candidate ontologies may be presented by their election committees composed of the knowledge engineers and subject experts who took part in the initial development of the artifact. The presentations could be compared by the electorate like it happens in politics to the political programs. The candidate ontologies may be invited and take part in the public debates. Their competences will thus be cross-checked by the members of the competing election committees. The results of the debates will provide more evidence to the electorate for making their informed votes. It may happen that none of the candidate ontologies receives the majority in the campaign. Such an outcome in politics may be treated as an event leading to a one more round of elections. In ontology engineering a new election round may also be used constructively for the refinement of the candidates. The development teams may make use of the election results by concentrating on answering the competency questions they failed to answer by the previous revision. Coalitions may also be fruitful if it turns out that merging some candidate ontologies will substantially improve their joint competency. The election committees of the merged ontologies of course have to reach an agreement on how their "societal influence" will be redistributed after their victory in elections. Following the outlined procedure for ontology election, if a particular ontology gets the majority then it can be expected that the electorate's commitment to this ontology is strong enough to ensure its smooth uptake. Hence, a high level research hypothesis in our proposal is:

Introducing "democracy" in ontology engineering by incorporating a competitive and transparent procedure of ontology elections makes the process effective, development cycles shorter, and enables better results transfer to industry. The effectiveness is ensured by the fact that the developed artifacts will be appropriately refined following the intended requirements of the users in iterations (election rounds) with active participation of these intended users. The incentive for the active involvement of the domain specialists and knowledge engineers is reciprocal as they share common objectives. Moreover, both parties are naturally motivated by the competitive nature of the process. The iterations become shorter and pursue better defined and more focused objectives that reflect the desires of the domain specialists adequately. Hence, a better transfer is ensured by the fact that the commitment of the intended users to the winner is the strongest among the alternative candidates.

There are several research questions that have to be answered in more detail and rigor for proving this research hypothesis:

(i) Why will industrial domain specialists be willing to join election committees and vote in election rounds?

This research question is very similar in its nature to the question about the proper motivation for people to join political groups and vote in elections. The answer may be sought by devising appropriate reciprocal incentive schemes motivating industrial domain specialists and knowledge engineers. For that looking at the results in several European projects may be useful. FP7 ACTIVE project derives the recommendation on possible incentive schemes, in particular for ontology development, by looking at the teams of knowledge workers as social structures. The high level objective of the

FP7 INSEMTIVES project (http://www.insemtives.org/) is to bridge the gap between human and computational intelligence and providing incentives for users to contribute to the massive creation of semantic content. A general framework for organizing campaigns may be adapted from the EU infrastructure project SEALS (http://www.seals-project.eu/) which develops a reference infrastructure, the SEALS Platform, to facilitate the formal evaluation of semantic technologies.

(ii) What are the proper ontology representation notation and collaboration platform for transparent election debates?

The first part of this research question is about the proper balance between the expressivity of the notation for representing ontologies and the ease of comprehension of this notation for the users that are not knowledge engineers. The expressivity has to be equal to the tractable subset of the chosen ontology specification language (for example OWL 2.0 – a de-facto standard ontology representation language to date). For ensuring the ease of comprehension we have to take into account that the representatives of the electorate are industrial knowledge workers – the engineers who develop, adapt, or adopt IT solutions in their businesses. A UML-based language is one of the commonly used notations for these professionals. Therefore, a visualized ontology specification in UML or its extension (for example OntoUML [28], or other UML variants with appropriate expressiveness [29]) accompanied by the textual description of ontology elements in a natural language would be appropriate – please refer to Section 6 for more details.

The second part of the question is about the collaboration platform that enables efficient debates. We consider that Semantic MediaWiki (SMW) with extensions is very appropriate as such a platform. An argument in favor of a Wiki-based infrastructure is that it is a Web 2.0 platform that intrinsically supports the exchange of opinions and has been extensively used for collaborative content development, "crowd-sourcing", user community development, etc.

(iii) How to ensure the swift convergence of the series of election rounds to the appearance of a single good scorer ontology?

The answer to this research question has to be sought by looking for the proper set of heuristic rules and policies for ontology elections. It could be rightfully argued that such a statement is a way too succinct to answer the question. However we do not have a more detailed recommendation at the moment and leave this very important issue for future research and experiments.

6 Xtreme Documenting and Ontology Discussions On-Line

One of the hypotheses we pointed out is: the ontology developed with active involvement of the intended users has to be presented in a form that is easily comprehendible by these users. A rich self explanatory notation with a user interface that is native for the target user group has to be exploited for that. We believe that a proper way of presenting the ontology to the subject experts is the documentation of the ontology as it combines the formal definitions of the ontology elements in textual and graphical representation with the informal descriptions of the semantics of those elements. Of course the documentation has to be developed in line with the ontology design to be available in proper time for discussing the design decisions. Ideally, the documentation describing a concept, a property, an ontological module has to appear at the same time with the appearance of the design of this ontology element. Apart of that, the documentation has to contain the information about the ontological context of the element and the information about its evolution.

As a proof of concept we have developed the electronic documentation site for ontologies – OntoDocWiki (http://isrg.kit.znu.edu.ua/ontodocwiki/). This resource is based on the SMW with LiquidTreads extension (www.mediawiki.org/wiki/Extension:LiquidThreads) as a basic collaboration infrastructure. Currently the resource contains the documentation of the PSI Suite of Ontologies v.2.3. This Suite of Ontologies has been developed in the Performance Simulation Initiative (PSI) project[4] and further refined for the needs of the ACTIVE Project. The wiki articles represent documentation for the Suite of Ontologies (Fig. 2), each individual ontology module (Fig. 3), each individual concept (Fig. 4) have been semi-automatically generated [30] based on the reference specification [31].

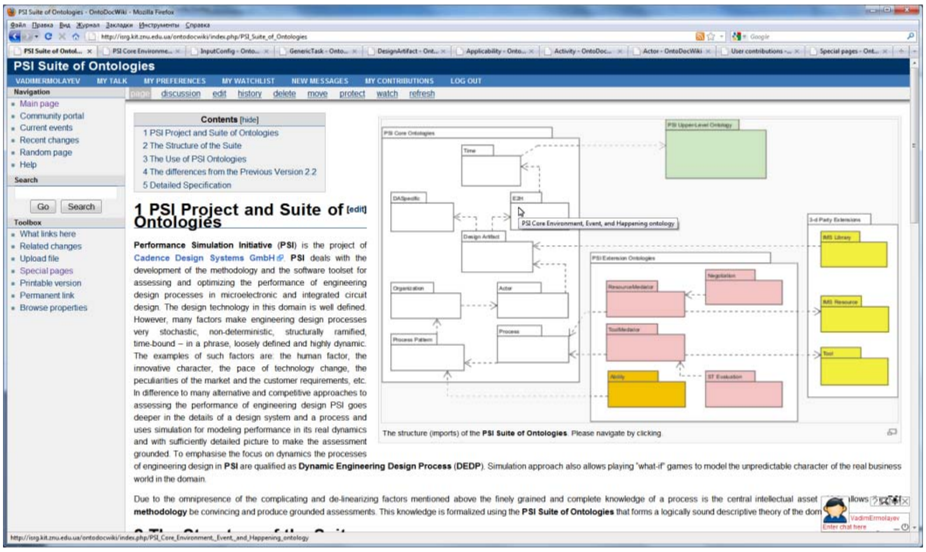

Fig. 2. OntoDocWiki article describing the PSI Suite of Ontologies v.2.3.

As pictured in Fig. 2–4, the articles comply with all the outlined requirements for describing ontology elements in an easily comprehendible and non-professional user friendly, yet informative manner. They combine textual descriptions with the graphical (UML) representation of the ontological contexts. These graphical representations are implemented as image maps and allow easy navigation in the pictured ontological contexts. For example, as pictured in Fig. 2, clicking the package representing the E2H ontology in the package structure diagram opens the article describing the E2H ontology (Fig. 3). Similarly, the class diagram of the particular ontology allows navigation to individual concept articles (e.g. the concept of a

[4] PSI is the accomplished R&D project of Cadence Design Systems GmbH.

Happening, Fig. 4) or related ontological modules represented in the ontological context.

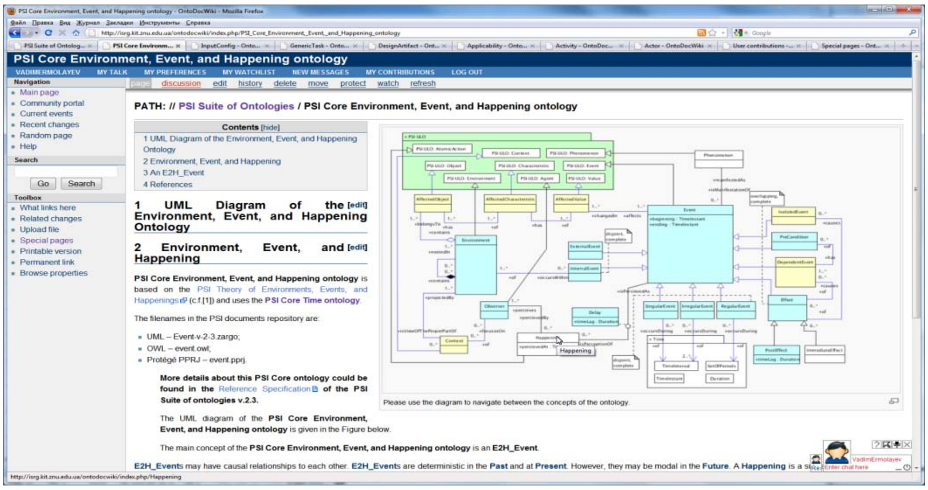

Fig. 3. OntoDocWiki article describing the PSI Environment, Event, and Happening (E2H) ontology v.2.3.

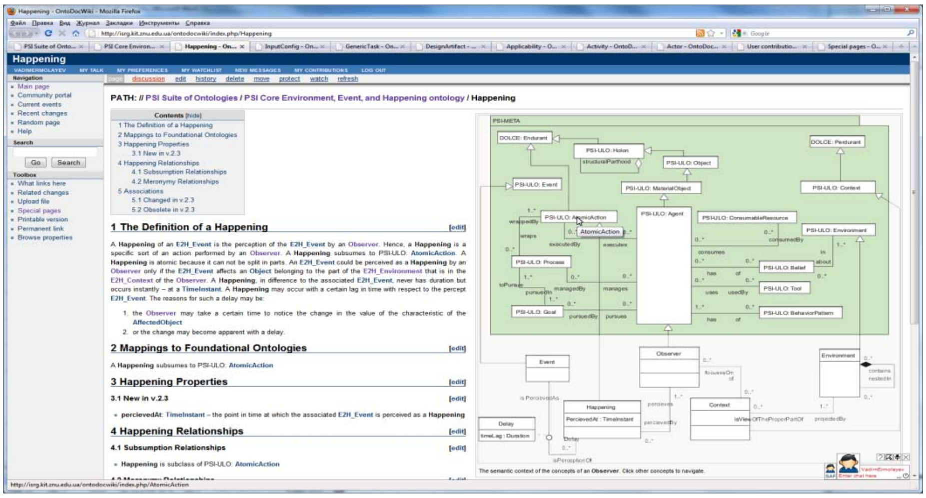

Fig. 4. OntoDocWiki article describing the concept of a Happening of the PSI E2H ontology v.2.3.

At a concept level the documentation is informative enough for allowing a non-professional user to evaluate the semantics of the concept and the surrounding ontological context [29]. As shown in Fig. 4, the article documenting a concept contains the information about:

- The relationship of the concept to the higher-level or foundational ontologies. For example, the concept of a Happening is described as a subclass of an AtomicAction – the concept of the PSI Upper-Level ontology.

– The explanation of the semantics of the datatype properties.
– The object properties grouped by the type of relationship: subsumption (if a concept is a subclass of another concept), part-whole relationships (represented as aggregations or compositions in the UML class diagram), and associations.

The descriptions of individual properties are structured in a way to present the evolution of those properties. For the properties that have been changed in the current version the information about the change is given. For those properties that have been introduced or became obsolete the rationale for this design decision is described. For the properties that may be used differently depending on the deployment of the Suite the variants of use are explained. For example, in some applications, like ProjectNavigator [32] the Suite of PSI Ontologies may be used without its Upper-Level ontology.

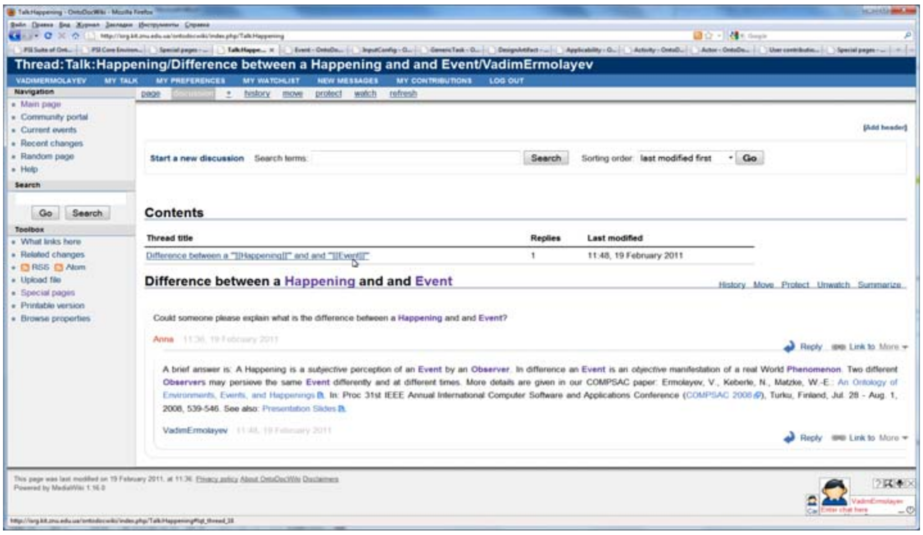

Fig. 5. OntoDocWiki discussion page for the concept of a Happening.

In ACTIVE project the PSI Suite of Ontologies has been used as the background knowledge for describing the subject domain of one of its case studies – knowledge processes in microelectronic engineering design. The fully functional prototype of the ACTIVE Design Project Visualizer [33] has been developed and validated based on this backbone knowledge representation. Prototype development has been done by several project partners. Only one of the partners was the owner of the PSI Suite – the rest were not familiar with these ontologies. Therefore we had to develop an ontology documentation and discussion resource for ramping-up the collegues. It turned out that OntoDocWiki became very helpful in both: explaining the PSI ontologies to the software developers, some with a marginal background in knowledge engineering; and collaboratively refining the ontologies in response to the project requirements. In fact the v.2.3 release of the Suite has been developed in this collaboration based on the extensive use of the OntoDocWiki platform.

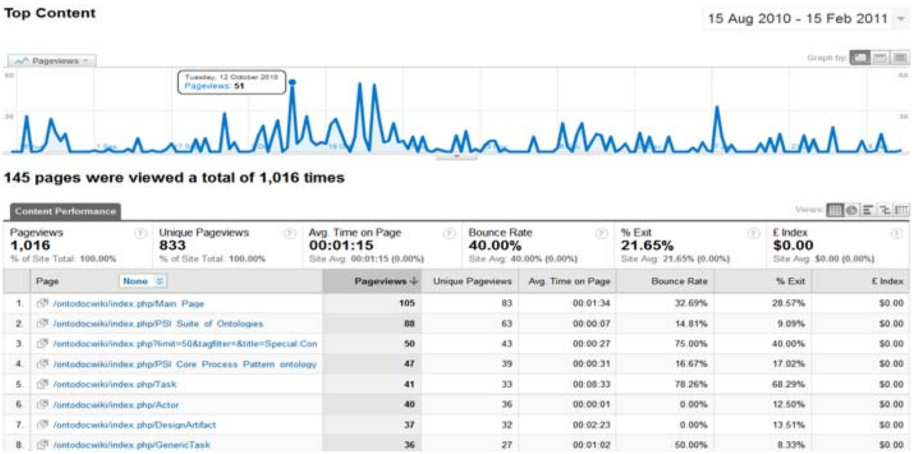

	Page	Pageviews	Unique Pageviews	Avg. Time on Page	Bounce Rate	% Exit	£ Index
1.	/ontodocwiki/index.php/Main Page	105	83	00:01:34	32.69%	28.57%	$0.00
2.	/ontodocwiki/index.php/PSI Suite of Ontologies	88	63	00:00:07	14.81%	9.09%	$0.00
3.	/ontodocwiki/index.php?limit=50&tagfilter=&title=Special:Con	50	43	00:00:27	75.00%	40.00%	$0.00
4.	/ontodocwiki/index.php/PSI Core Process Pattern ontology	47	39	00:00:31	16.67%	17.02%	$0.00
5.	/ontodocwiki/index.php/Task	41	33	00:08:33	78.26%	68.29%	$0.00
6.	/ontodocwiki/index.php/Actor	40	36	00:00:01	0.00%	12.50%	$0.00
7.	/ontodocwiki/index.php/DesignArtifact	37	32	00:02:23	0.00%	13.51%	$0.00
8.	/ontodocwiki/index.php/GenericTask	36	27	00:01:02	50.00%	8.33%	$0.00

Fig. 6. The statistics of use of the OntoDocWiki resource provided by Google analytics.

One of the features that proved to be useful and effective in ramping-up the users was discussions about ontology elements. Fig. 5 pictures an example of such a discussion about the semantics of the concept of a Happening. The utility of a discussion is straightforward – anyone may pose a question or present an argument about the semantics of an element in the ontological context of the related wiki article; anyone else may offer an answer or a counter-argument. The outcome of the discussion is often a better and deeper comprehension of the semantics. It also turned out some times that the discussions led to the changes in the ontology.

Apart of the more active involvement of the users in the ontology development process, the use of the Wiki as a platform allowed making the arguments or the statements in discussions more grounded by linking existing articles to the parts of the discussion statements. For example, the links to the items describing the concepts that are related for the discussion thread were inserted in the headings and the discussion statements using MediaWiki markups for hyperlinks. Those were either internal OntoDocWiki articles (e.g. Event) or the pointers to the external resources (for example [34]) offered as back-up information in support of the discussion statements.

Fig. 6 shows some of the usage statistics for the OntoDocWiki. It is topical to notice that the peaks in access hits (one was on the 12-th of October 2010) depict the activity of users exactly in the periods of ontology discussions that occurred in different phases of the prototype development. It is also interesting that the most frequently visited OntoDocWiki pages were the ones describing the ontology concepts that were the most complex in semantics compared to the others. So, it took people more time and more visits to comprehend these ontology elements. Such information on the intensiveness of the use of different parts of ontology documentation turns out to be extremely important for the knowledge engineering team. Indeed, it objectively measures the complexity of the comprehension of the ontological concepts that could become a problem for gaining the ontological commitment by the community of the intended ontology users. These measures may be valuable in adjusting the tactics and foci in the ontology election campaigns.

7 Concluding Remarks

Ontology engineering as a field has been in intensive research and development for a substantial period of time. As we have hopefully shown in the paper, ontology engineering technology has passed the peak of inflated expectations on the Gartner's hype cycle curve (see e.g. http://en.wikipedia.org/wiki/Hype_cycle). Currently the position is in the proximity of the through of disillusionment, probably a little bit shifted to the slope of enlightenment. Most evidently, semantic technologies in broad and ontology engineering technologies in particular, will be better accepted by industries if and when they reach Gartner's plateau of productivity. Our intension while writing this discussion paper was to analyze the reasons for the disillusionment and, perhaps, to shed the light on the possible way up the slope to the region of industrial maturity. We believe that the right way is marked by the increase in the commitment of the intended industrial users.

We have outlined our views on a possible methodological framework for ontology engineering – OntoElect. We believe that our approach may be capable to relax several barriers on the way of gaining better and broader commitment to the use of ontologies in industrial applications. These beliefs are backed-up by our experience in implementing and using some of the elements of this framework in our research and development work in several European projects. We had positive experience in testing parts of the approach. For example, our work on PRODUKTIV+[5] ontologies happened to be the informal competition of the two separate groups which views further converged to a single (merged) ontology suite – in debates. Another example that highlights the necessity of involving industrial subject experts and users as early as possible was our work on the PSI Suite of Ontologies together with the experts and users from Cadence Design Systems GmbH. The use case for testing the user friendly way for representing ontologies on a collaborative platform was the one in the ACTIVE project – presented in Section 6.

Acknowledgements

PSI Suite of Ontologies has been developed in the PSI project funded by Cadence Design Systems GmbH. OntoDocWiki development has been partially supported by the ACTIVE project funded in part by the European Commission Framework Programme 7. The authors would also like to thank the anonymous reviewers for their thorough comments that helped substantially to improve the paper.

References

1. Gruber, T. R.: Toward Principles for the Design of Ontologies Used for Knowledge Sharing. Int. J Human-Computer Studies 43, 907--928

[5] PRODUKTIV+ is the accomplished R&D project funded by the German Bundesministerium für Bildung und Forschung (BMBF).

2. Krötzsch, M., Vrandecic, D., Völkel, M., Haller, H., Studer, R.: Semantic Wikipedia. J of Web Semantics 5(4), 251--261 (2007)
3. Uschold, M., Grüninger, M.: Ontologies: Principles, Methods, and Applications. Knowledge Eng. Rev. 11(2), 93--155 (1996)
4. Hepp, M.: Possible Ontologies: How Reality Constrains the Development of Relevant Ontologies. IEEE Internet Computing 11(1), 90--96 (2007)
5. Guizzardi, G.: Theoretical foundations and engineering tools for building ontologies as reference conceptual models. Semantic Web 1-2(1), 3--10 (2010)
6. OWL 2 Web Ontology Language Primer, http://www.w3.org/TR/owl2-primer
7. Gómez-Pérez, A., Fernández-López, M., Corcho, O.: Ontological Engineering. Springer, London (2004)
8. Pinto, H.S., Tempich, C., Staab, S., Sure, Y.: DILIGENT: Towards a fine-grained methodology for distributed, loosely-controlled and evolving engineering of ontologies. In: de Mántaras, R.L., Saitta, L. (eds.) 16th European Conf. on Artificial Intelligence. ECAI, pp. 393--397, IOS Press (2004)
9. Sure, Y., Studer, R.: On-To-Knowledge methodology: On-To-Knowledge: Semantic Web enabled Knowledge Management. In: Davies, J. (eds.), J. Wiley and Sons (2002)
10. Uschold, M., King, M.: Towards a Methodology for Building Ontologies. In: Skuce, D. (eds.) IJCAI'95 Workshop on Basic Ontological Issues in Knowledge Sharing, pp.6.1--6.10. Montreal, Canada (1995)
11. Holsapple, C.W., Joshi, K.D.: A collaborative approach to ontology design. Comm. ACM, vol. 45, pp. 42--47. ACM Press, New York (2002)
12. Buckingham-Shum, S., Motta, E., Domingue, J.: Augmenting design deliberation with compendium: The case of collaborative ontology design. In: HypACoM 2002: Facilitating Hypertext-Augmented Collaborative Modeling. ACM Hypertext'02 Workshop, Maryland University (2002)
13. Kotis, K., Vouros, G.A., Alonso, J.P.: HCOME: tool-supported methodology for collaboratively devising living ontologies. In: SWDB'04: Second International Workshop on Semantic Web and Databases, Co-located with VLDB. Springer-Verlag (2004)
14. Schreiber, G. et al.: Knowledge Engineering and Management. The CommonKADS Methodology. MIT Press, Cambridge, USA (1999)
15. Suárez-Figueroa, M. C. et al. D5.4.1: NeOn Methodology for Building Contextualized Ontology Networks. Technical report, The NeOn Project (2008)
16. Pâslaru-Bontaş, E.: A Contextual Approach to Ontology Reuse: Methodology, Methods and Tools for the Semantic Web. PhD Thesis, Freie Universität Berlin, Berlin (2007)
17. Simperl, E., Tempich, C., Sure, Y.: A Cost Estimation Model for Ontology Engineering. In: Cruz, I. F., et al. (eds.) ISWC 2006. LNCS, vol. 4273, pp. 625--639. Springer, Berlin, Heidelberg (2006)
18. Henderson-Sellers, B., Gonzalez-Perez, C.: Standardizing Methodology Metamodelling and Notation: An ISO Exemplar. In: Kaschek, R., Kop, C., Steinberger, C., Fliedl, G. (eds.) UNISCON 2008. LNBIP, vol. 5, pp. 1--12. Springer, Berlin, Heidelberg (2008)
19. Gupta, M., Li, R., Yin, Z., Han J.: Survey on Social Tagging Techniques. SIGKDD Explorations 12(1), 58--72 (2010)
20. Uren, V., Cimiano, P., Iria, J., Handschuh, S., Vargas-Vera, M., Motta, E., Ciravegna, F.: Semantic annotation for knowledge management: Requirements and a survey of the state of the art. Science. Services and Agents on the World Wide Web 4(1), 14--28 (2006)
21. Hunter, J., Khan, I., Gewrber, A.: HarvANA – Harvesting Community Tags to Enrich Collection Metadata. In: Paepcke A, Borbiha J, Naaman M (eds.) 8th ACM/IEEE-CS Joint Conference on Digital Libraries, pp. 147--156. ACM New York, New York (2008)

22. Siorpaes, K., Hepp, M.: Games with a Purpose for the Semantic Web. IEEE Intelligent Systems 23(3), 50--60 (2008)
23. Zhdanova, A.V.: Community-driven Ontology Construction in Social Networking Portals , Web Intelli. and Agent Sys. 6, 1, 93--121 (2008)
24. Von Ahn, L.: Games with a Purpose. Computer 29(6), 92–94 (2006)
25. Law, E. et al.: Tagatune. In: Proc. Int'l Conf. Music Information Retrieval, pp. 361--364, Austrian Computer Soc., Ismir (2007)
26. Lieberman, H., Smith, D., Teeters, A.: Common Consensus: A Web-Based Game or Collecting Commonsense Goals. In: Proc. Workshop Common Sense for Intelligent Interfaces, ACM Conf. Intelligent User Interfaces. ACM Press (2007)
27. Petrie, C.: Pragmatic Semantic Unification. IEEE Internet Computing, vol. 9, no. 5, pp. 95--96 (2005)
28. Guizzardi, G.: Ontological Foundations for Structural Conceptual Models. Telematica Instituut Fundamental Research Series No. 15, Netherlands (2005)
29. Ermolayev, V., Copylov, A., Keberle, N., Jentzsch, E., Matzke, W-E.: Using Contexts in Ontology Structural Change Analysis. In: Ermolayev, V., Gomez-Perez, J.-M., Haase, P., Warren, P. (eds.) CIAO, CEUR-WS/vol. 626 (2010)
30. Ermolayev, V., Tatarintseva, O.: Applied Research and Development in Cooperation with Industry. Information Technologies in Education 5, 16--26 (2010)
31. Ermolayev, V., Jentzsch, E., Keberle, N.: Performance Simulation Initiative. The Suite of Ontologies v.2.3. Reference Specification. Technical Report, PSI-ONTO-TR-2010-1, VCAD EMEA, Cadence Design Systems, GmbH (2010)
32. Sohnius, R., Jentzsch, E., Matzke, W.: Holonic simulation of a design system for performance analysis. In: HoloMAS '07: Proc. of the 3rd international conference on Industrial Applications of Holonic and Multi-Agent Systems, pp. 447--454. Springer-Verlag, Berlin, Heidelberg (2007)
33. Ermolayev, V., Dengler, F., Fortuna, C., Stainer, T., Bosser, T., Melchior, E-M.: Increasing predictability and sharing tacit knowledge in electronic design. In: Warren, P., Simperl, E., Davies, J. (eds.) Context and semantics in knowledge management, Springer Verlag, Berlin, Heidelberg (2011)
34. Ermolayev, V., Keberle, N., Matzke, W.-E.: An Ontology of Environments, Events, and Happenings. In: Proc 31st IEEE Annual International Computer Software and Applications Conference, pp. 539--546. Turku, Finland (2008)