The CKC Challenge: Exploring Tools for Collaborative Knowledge Construction

Natalya F. Noy, Abhita Chugh, Harith Alani

Abstract

The great success of Web 2.0 is mainly fuelled by an infrastructure that allows web users to create, share, tag, and connect content and knowledge easily. The tools for developing structured knowledge in this manner have started to appear as well. However, there are few, if any, user studies that are aimed at understanding what users expect from such tools, what works and what doesn't. We organized the Collaborative Knowledge Construction (CKC) Challenge to assess the state of the art for the tools that support collaborative processes for creation of various forms of structured knowledge. The goal of the Challenge was to get users to try out different tools and to learn what users expect from such tools—features that users need, features that they like or dislike. The Challenge task was to construct structured knowledge for a portal that would provide information about research. The Challenge design contained several incentives for users to participate. Forty-nine users registered for the Challenge; thirty-three of them participated actively by using the tools. We collected extensive feedback from the users where they discussed their thoughts on all the tools that they tried. In this paper, we present the results of the Challenge, discuss the features that users expect from tools for collaborative knowledge constructions, the features on which Challenge participants disagreed, and the lessons that we learned.

1   Why the Challenge?

The great success of Web 2.0 is mainly fuelled by an infrastructure that allows web users to create, share, tag, and connect content and knowledge easily. In general, the knowledge created in today's applications for collaborative tagging, folksonomies, wikis, and so on is mostly *unstructured*: tags or wiki pages do not have any semantic links between them and usually are not related to one another in any structured form. By contrast, ontologies, database schemas, taxonomies are examples of *structured* knowledge: these artifacts usually contain explicit definitions of, and links between, their components, often with well-defined semantics.

It is natural that both researchers and practitioners are now beginning to explore how the power of creating web content in a social environment can be used to acquire, formalize and structure knowledge. This integration of Web 2.0 and Semantic Web approaches can come from both directions: through applications that add explicit semantics and structure to the sets of artifacts produced by today's Web 2.0 applications; and through using the social and collaborative approaches to defining the knowledge structures and models themselves.

The new generation of tools that support these paradigms ranges from extensions of wikis that support semantic links between wiki pages (e.g., Semantic MediaWiki [9], BOWiki [2], Platypus Wiki [7], etc.), to tools to tag bookmarks collaboratively and to organize the tags in some semantic

structure (e.g., BibSonomy [1], SOBOLEO [10], etc.), to full-fledged ontology editors that support distributed and collaborative development of ontologies (e.g., Collaborative Protégé [8], OntoWiki and pOWL [6], etc.). Commercial tools such as Freebase [4] are also entering the field. Most of these tools are in fairly early stages of development and the whole field of constructing structured knowledge collaboratively is nascent. There are few, if any, user studies that are aimed at understanding what users expect from such tools, what works and what doesn't.

We organized the Collaborative Knowledge Construction (CKC) Challenge to assess the state of the art for the tools that support collaborative processes for creation of various forms of structured knowledge. We ran the Challenge under the auspices of Workshop on Social and Collaborative Construction of Structured Knowledge held at the 16th International World Wide Web Conference at Banff, Canada. [1] Our main goal was to get users to try out different tools for collaborative construction of structured knowledge, and to discuss requirements for such tools, features that users need, features that they like or dislike. We use the term "structured knowledge" to describe anything from a simple hierarchy of tags to a fully axiomatized OWL ontology.

The Challenge ran for two weeks prior to the workshop and everyone was invited to participate. One did not have to attend the CKC workshop in order to take part in the Challenge. For the six tools that participated (Section 2), we invited users to enter information on their research domains: their institutions, papers, various events, such as conferences and workshops, research topics, and anything else they find relevant. Because the participating tools differed in their expressive power and the type of information they supported, the exact information that users could enter in each of the tools was different as well.

At the end of the Challenge, we invited users to provide feedback on each of the tools that they tried and to reflect on their experience in general, leading up to the question of what their ideal tools for collaborative knowledge construction would be. During the Challenge, we ran a competition for the most active user and the most insightful feedback (promising mystery prizes to the winners) to encourage users to participate.

The goal of the Challenge was not to compare the tools themselves: the field is still in its infancy and the tools are too different from one another in terms of their functionality to have a competition between the tools. Rather, our goal was to get some direction on the path forward in the development of such tools. In fact, the developers of the Challenge tools have all actively supported the Challenge and were very keen to observe user experience with their tools and to get the detailed feedback. In this paper, we present the results of the Challenge, discuss the features that users expect from tools for collaborative knowledge constructions, the features on which Challenge participants disagreed, and the lessons that we learned.

2 The Challenge Tools

We solicited tool demonstrations for participation in the workshop. From the accepted demos, any tool that was domain independent, dealt with structured knowledge, and allowed users to work collaboratively became part of the Challenge. Table 2 lists the six tools that participated. The tools range from folksonomy-style tools that enable users to add some structure to their tags to full-fledged ontology editors. We now take a slightly more detailed look at the participating tools. Papers describing each of the tools are available in the CKC workshop proceedings [15].

---

[1] http://km.aifb.uni-karlsruhe.de/ws/ckc2007

BibSonomy	A web-based social resource-sharing system allows users to organize, tag, and share bookmarks and bibliographic entries	University of Kassel, Germany
Collaborative Protégé	An ontology and instance editor with additional support for collaborative development (history, annotations, discussions, etc.)	Stanford University, US
DBin	An application that enables collaborative editing of knowledge bases. The CKC brainlet contains an ontology and instance editor and discussion facilities.	Università Politecnica delle Marche, Italy
Hozo	An ontology editor that enables asynchronous collaborative development	Osaka University, Japan
OntoWiki	An ontology and instance editor that provides collaborative capabilities, such as history and ratings	University of Leipzig, Germany
SOBOLEO	A web-based system that supports annotation of web resources and creation of taxonomies	FZI Research Center for Information Technology, Karlsruhe, Germany

Table 1: The Challenge tools and their descriptions. Detailed descriptions of all the Challenge tools are available in the CKC workshop proceedings [15]

2.1 BibSonomy

BibSonomy [14, 1] is a web-based social resource sharing system that allows users to organize, share, and tag bookmarks and publications collaboratively (Figure 1). Users can tag resources with terms that describe topics, names, places, events. Users can then search for resources by specifying the associated tags. Additionally, BibSonomy allows users to create relations between tags. For instance, a user can create a relation such as `Programming` ← `Java`, indicating that `Programming` is a more general tag than `Java` and the system should include resources tagged with `Java` when the user searches for `Programming`.

Users share their tags and tagged resources with other BibSonomy users. Relationships between tags however are kept in the user's private space.

2.2 Collaborative Protégé

Collaborative Protégé [8] is an ontology editor that supports a number of collaboration features (Figure 2). In addition to the common ontology-editing operations, it enables annotation of both ontology components and ontology changes (i.e., users can attach notes to either of these). Annotations may be of different types, such as comments, questions, examples, and so on. For instance, a user might attach an annotation to a class suggesting a new property be added to the class; a user might annotate a change that he made with an explanation for the change. Collaborative Protégé supports searching and filtering of user annotations based on different criteria. Users may also create discussion threads, create proposals for changes, and vote on these proposals; there is

also support for live chat within the editor. Multiple users access and edit the same ontology at the same time. All the changes made by one user are immediately visible to all the others.

2.3 DBin

DBin [17, 3] is a Peer-to-peer application that enables groups of users to create knowledge bases collaboratively. DBin allows domain experts to create domain- or task-specific user interfaces—*brainlets*—from a collection of components, such as ontology or instance editors, discussion forums, specific views, and other plugins. Specifically, the CKC brainlet—the brainlet we used for the Challenge—includes an ontology-editing and instance-acquisition components (Figure 3). Users edit an ontology directly on the server and can see changes made by other users. The system keeps and makes readily available provenance information such as who created a particular class or property. DBin runs as a local application that connects to a peer-to-peer network.

2.4 Hozo

Hozo [16, 5] supports asynchronous construction of ontologies (Figure 4). Ontologies are usually subdivided into multiple inter-connected modules. A developer checks out and locks a specific module, so that others cannot edit it, edits it locally, and then checks it back in. At the time of check-in, Hozo uses the dependencies that were declared between the modules to identify conflicts. Users can view the list of changes and accept or reject the changes made by other users. Hozo supports definitions of classes and properties. Although the standard version of Hozo supports instance editing, this capability was disabled in the version that participated in the Challenge for technical reasons.

2.5 OntoWiki

OntoWiki [11, 6] is a web-based application that supports collaborative building of ontologies and creation of instances (Figure 5). OntoWiki provides different views on instance data. For example, users can select instances that contain geographical information and view them on a map. Instances having property values of datatype *xsd:date* can be viewed on a calendar. OntoWiki enables users to rate content which helps in measuring the popularity of content. The system tracks all changes applied to the knowledge base and users can review this information. If users want to edit classes or properties, OntoWiki redirects them to pOWL, a web-based ontology editor.

2.6 SOBOLEO

SOBOLEO [10] is a web-based tool for annotation of web resources (Figure 6). It enables users to assign tags to web resources and to organize tags in a hierarchy. This hierarchy and tags are developed collaboratively and users immediately see changes performed by other users. Users can search through the annotated web resources using concepts from the taxonomy.

2.7 Summary of the Challenge Tools

As mentioned earlier, the goal of the challenge was not to compare the tools themselves. The functionality and focus of all six tools is different, thus giving the Challenge participants the opportunity to try out a variety of features. Table 2 summarizes some key features supported by the

tools in the Challenge. Note that some of these tools exist in other versions that support more of the features in this table. However, our description here is based entirely on the versions of the tools that their developers entered in the Challenge (i.e. what the Challenge participants had access to).

All six tools support creation of hierarchical information. Collaborative Protégé, DBin, and OntoWiki focus on editing ontologies. Hozo's main focus is building and integrating ontology modules, and detecting any resulting conflicts. BibSonomy and SOBOLEO are primarily designed for tagging web resources and creation of taxonomies by specifying relationships between tags. SOBOLEO and Collaborative Protégé allow users to chat and create discussion threads. Collaborative Protégé and OntoWiki enable users to rate ontology components and to add comments to ontology components. BibSonomy separates a user's personal space from the shared space of all users; in the other tools all information is meant to be shared. Collaborative Protégé and OntoWiki maintain the log of changes made to the ontology by different users and make it available during editing. Hozo supports asynchronous construction of knowledge and enables users to lock ontology components before editing them.

From the descriptions above we can see that the tools that participated in the Challenge target different scenarios for collaborative knowledge construction. In Collaborative Protégé and Hozo, for example, the final product of the collaborative effort are the ontologies themselves. In BibSonomy and SOBOLEO the main goal is annotation of web resources with tags. The tag hierarchies in BibSonomy and SOBOLEO are *byproducts* of the collaborative annotation effort, and are mainly used for searching and browsing the annotated resources. DBin and OntoWiki focus on knowledge acquisition, were the community of users can create and share knowledge bases, and enter new instance data using customized interfaces.

3   Preparing for the Challenge

The subject area that we selected for the Challenge was the representation of research domains; a subject that all participants would be familiar with. We hoped that having a domain where each of the users can say something about themselves would be a subtle factor that would encourage participation: Challenge participants would want to make sure that the knowledge bases contained at least the information about them and their institutions and papers—something close to their heart. More specifically, we suggested that users construct structured knowledge for a hypothetical portal that will access information about research. This knowledge would include information about research topics, conferences, workshops, seminars, talks, researchers and research groups, roles of researchers, publications (in these conferences, as well as journals, books, etc.), and relevant web page URLs (e.g. for papers, conferences, demos, topics). We suggested that users start by capturing the information about their own research topics, research groups, publications and conferences or other events that they were involved in. Users were free to expand the knowledge base in any direction.

Because the tools in the Challenge were quite different from one another and could represent different types of information, we did not expect users to enter the same information in each tool.

Prior to the Challenge we (the Challenge organizers) used each of the tools to enter some bootstrapping information about the CKC workshop itself (papers and first authors). By doing so, we have effectively tested each of the tools and its suitability to the task. Furthermore, for the tools that provided facilities for discussion, annotation or rating facilities were provided, we used

	BibSonomy	Collaborative Protégé	DBin	Hozo	OntoWiki	SOBOLEO
Hierarchy of concepts	√	√	√	√	√	√
Properties		√	√	√	√	
Instances of concepts in the hierarchy (includes tags)		√	√	†	√	√
Comments on ontology components		√	√		†	
Ratings		√			√	
Asynchronous editing of ontology modules				√		
Personal space	√			√		
History of changes		√	†	√ (module level)	√	
Discussion		√	†			√
Instant Chat		√				√
Creation of content through an integrated browser button	√					√
Web-browser interface	√				√	√

Table 2: Features in the Challenge tools. This table shows the main features of each tool based on the tool version that their developers made available for the Challenge. Features marked with † are available in some versions of the tools at the time of the Challenge, but were not available in the versions used for the Challenge.

these features to demonstrate their usage.

From our experience, this bootstrapping period was absolutely essential for the success of the Challenge. This phase helped us to make sure that all the tools work in the Challenge setting and are set up correctly. We identified tool features that were difficult to use or find and suggested that the tool developers add clarifications to a how-to page for using their tool in the Challenge. At this stage, we actively collaborated with the developers of the tools to ensure that the Challenge can realize the full potential of their respective tools.

4 Running the Challenge

The Challenge ran for a two-week period and we tried to advertise it as widely as possible. Anyone could participate as a user in the Challenge. We asked users to use the same username for all the

tools so that we could aggregate their activity.

One of our main challenges was to encourage as active participation as possible. In order to do that, we announced a competition for the most active user. The idea of a competition among the users was inspired by several earlier reports indicating that the game or competition component works as an excellent incentive for users [13, 18]. Each day, we collected the activity information from each of the tools and aggregated it on the Challenge Web site. Thus, the users could track their activity with respect to the other users throughout the Challenge (not unlike in a computer game). We should note that the "number of facts entered" was computed differently by different tools and adding information about a paper in one tool would count as a different number of operations than in another tool. However, because the competition was rather informal (and prizes symbolic), this discrepancy was not a major concern.

Our scoring system explicitly encouraged users to try more than one tool: making, say, 30 changes in a single tool counted less than using three tools and making 10 changes in each. Each edit in the second and subsequent tools (in terms of the user's own activity) was more valuable than the edits in the tool that he used the most.

Finally, we asked users to fill out a feedback form for each of the tools that they tried and a general feedback form reflecting on their experience. For the tool-specific feedback, we asked users which features of the tool they liked, what they thought needed to be improved, which type of content was easy to enter, which type of content was not supported, and so on. After trying several tools, users could provide general feedback with their thoughts on collaborative knowledge construction. We asked users how they thought an ideal tool for collaborative construction of structured knowledge should look like.

5   The Challenge Results

During the two-week Challenge, the CKC Challenge website had visitors from over 60 countries. User activity was at its peak during the last 3 days of the Challenge. Forty-nine users registered for the Challenge; thirty-three of them participated actively. Over 50% of the 33 active participants tried out one tool only (Figure 7). We have received 36 feedback forms: 31 tool-specific feedbacks, and 5 general feedbacks.

Among the 6 participating tools, SOBOLEO had the highest user activity: 65% of active users tried out SOBOLEO (Figure 8).

Table 5 shows how much information and what type of information was added to each tool. BibSonomy clearly had the largest amount of information entered; some of it though was uploaded from BibTex files, but the tags were entered manually. Recall that Hozo uses the check-in/ check-out mechanism and the users who used the tool made the changes in their local versions but forgot to check them back in. Hence, the Hozo server recorded no activity.

The tools in the Challenge represented quite a wide range of functionality. SOBOLEO and BibSonomy focused on annotating resources, such as web pages and bibliography entries; the other tools focus on building structured data and ontologies. We observed user preference for one type of system over the other when we aggregated results from all six tools. For instance, users who tried SOBOLEO, also tried BibSonomy; users who tried Collaborative Protégé, also tried OntoWiki. Those who used more than two tools had their top two tools (in terms of their activity) of the same type.

From the usage statistics (Figure 8), it is clear that download was a big impediment for users.

Tool	Structures entered	Popular features
BibSonomy	1006 bookmarks (846 unique) 307 BibTex entries (187 unique) 35 relations 1118 tags	`postBookmark` and `postPublication` buttons in a browser Upload to EndNote
Collaborative Protégé	32 classes 212 instances	Discussion, voting, chat and their integration with the rest of the tool Robustness of the tool
DBin	4 classes 17 instances	Customizable user interface
Hozo	No data available	Visualization features
OntoWiki	3 classes 80 instances Over 100 ratings	Maps Ratings
SOBOLEO	202 concepts 155 documents Average of 3 topics per document 393 concept relations	Ease of use

Table 3: The summary of the data entered by users in the tools during the two week evaluation period. These numbers do not include the initial seed data that the Challenge organizers entered in the tools. The third column lists the most popular features that users liked in each tool

DBin and Hozo required download and installation and got used the least. This intuition was confirmed by the feedback from users. It is worth noting, however, that the participants in the Challenge used the tools only for the sake of this Challenge rather than to do any of their usual work. Perhaps if there is a more pressing need for a particular tool then download would not be such an impediment. However, several users pointed out explicitly that a Web interface (and not, say, an applet, as Collaborative Protégé) was ideal for such tools.

We noted that there was very little discussion, even in the tools that enabled such a feature (Collaborative Protégé, SOBOLEO). Perhaps, this lack of discussion was due to the domain being too straightforward: there was simply not much to discuss or to argue about. Another possible explanation for the lack of discussion is that users had no real use for the resulting knowledge and thus were not too bothered with some imperfections.

6   What We Learned

Our main goal for the Challenge was to assess the state of the art for the tools for collaborative construction of structured knowledge, but, more important, to understand what users ultimately want from such tools in terms of features and functionalities. Our discussion in this section is based on the information in the feedback forms and the discussion of the Challenge at the CKC workshop at WWW 2007. Even given the large perceived diversity of the tools, Challenge participants agreed on a number of issues that will be critical in the tools for collaborative construction of structured

knowledge.

The first thing that we discovered was that "collaboration" means different things to different people. To many, a tool supports collaborative development if it simply provides distributed users access to the same ontology or knowledge base. To others, a tool supports collaborative development only if it includes support for annotations from different users, provenance, ways to reach consensus, and so on.

However one interprets the term, the tools that support collaborative construction of structured knowledge are still, for the most part, in their infancy. We, as a community, still don't have a good idea of what these tools should do or how they should work.

One of the questions we asked the participants was *"If you were to design your perfect tool for collaborative construction of structured knowledge, which features from which tools would you incorporate?"* Figure 9 presents some of the answers. In general, however, everyone agreed that no such perfect tool exists yet. Different scenarios require different **workflows** and hence different tool support. For instance, a group of users working on developing an ontology in the context of a specific project will have different requirements compared to an open community developing a lightweight taxonomy that anyone can edit. The Challenge participants expressed different expectations from the tools for adding structure to folksonomies than for the tools for collaborative development of full-fledged ontologies. Everyone seemed to agree that there should be several of such tools to satisfy the varying requirements and settings of different collaborative development scenarios. However, it will be crucial for the various tools to interoperate by enabling content to be imported to and exported from one another. The users praised the Challenge tools that had such functionalities.

The spectrum of **collaboration scenarios** that the tools must eventually support is quite broad. In some cases, tools should support specific protocols for making changes, where some users can propose changes, others can discuss and vote on them, and only users with special status can actually perform the changes. At the other end of the spectrum are settings where anyone can make any changes immediately. Thus, tools need to support different **mechanisms for building consensus**, depending on whether environment is more open or, on the contrary, more controlled. While none of the tools in the Challenge had specific support for any consensus-building protocol, several of the tools provided some technical means to support such a process. These mechanisms included rating, voting, and discussion functionalities. However, we noticed that these features were hardly used in the course of the Challenge. One of the possible reasons was that the nature of the selected domain was fairly straightforward and not controversial; furthermore, the setting of the Challenge did not require reaching consensus before committing a change.

Ratings of concepts and individuals was one of the features provided by the tools (Collaborative Protégé and OntoWiki). Users have noted however that ratings without comments and explanations attached to the ratings were not very useful by themselves. Other users suggested using conventional chat software (e.g. Yahoo, MSN) for the discussions instead of implementing tool-specific chat services. There were some general concerns that discussions and voting techniques in general could slow down the editing process if tools demand a consensus on every editing action before it can take effect.

Different tasks and settings also require different level of **expressive power** in the tool. For instance, a simple taxonomy may be sufficient for organizing tags for shared annotation of Web resources; users developing a medical terminology collaboratively may need a more expressive ontology. Thus, not all communities and settings will require sophisticated features provided by some of the tools. Even though a number of tools in the Challenge supported ontology development,

none provided capabilities for using a very **expressive ontology language**, such as OWL. [2] This lack of expressive power did not bother any of the participants and some suggested that perhaps it may be difficult to edit very expressive ontologies in a collaborative setting.

Naturally, more open environments, where anyone can join the editing process will need more advanced support for determining **trust and credibility** of various users. Given the complex nature of the task such as ontology editing, poor entries may result not only from malicious intent but also from simple incompetence. One possible approach is to measure users reputations based on how the community votes or rates their edits and change proposals. Similarly, support for and easy access to **provenance information** is critical in distributed collaborative projects. Users must be able to see who made the changes and when, to read a comment by the author of this change, to understand what was the state of the knowledge base when the change was made, to access concept and change histories, and so on.

Facilitating **ontology comprehension**—something at which ontology-development tools never seem to succeed fully—becomes an even bigger issue in the collaborative setting [12]. Users need to understand and modify ontology elements developed by others and having visual aids and other features to increase ontology comprehension is critical. In fact, difficulty of understanding both the ontologies themselves and the features of the more sophisticated tools may be one of the bigger challenges for this field. The users liked some of the visual aids that some of the tools provided, such as the interactive graphs of Hozo, and the map and calendar widgets of OntoWiki.

Several participants noted that a **good user interface** is critical. Users mentioned that having a lot of features is useless if they are hard to find or use. In addition, many of the tools in the Challenge ran very slow and such **performance** is unacceptable if these tools are to be used for real work. In general, participants agreed that having a Web interface to the tool becomes increasingly important and new technologies, such as AJAX, enable tools to support a wide range of sophisticated user interface features in this setting. Some users suggested building browser plugins to speed up the process of inserting and accessing knowledge from some of the lightweight tools (e.g. BibSonomy, SOBOLEO).

It is interesting that simple features like **autocomplete** for property values and names of new objects can also make all the difference in such tools. Users found autocomplete very useful and several of them mentioned it in their feedback. Furthermore, good autocomplete helps reduce unnecessary duplication of data: For example, the Freebase designers noted[3] that once they made autocomplete fast, the number of duplicates in the system was reduced dramatically. It turned out that the users were creating their own objects without waiting for autocomplete to finish.

One of the topics that appeared to be more controversial and that elicited emotional response from proponents on both sides was whether the tools should support **personal spaces** for user's own structures and data (e.g., BibSonomy) or whether everything should be immediately **shared** with the community (e.g., Collaborative Protégé, OntoWiki, SOBOLEO). In the third (intermediate) case, the data is ultimately shared, but a user can develop it asynchronously, in his own personal space before posting it to the world (e.g., Hozo). Ultimately, there was no agreement on which model works best, and, again, different setting will probably necessitate different boundaries between personal and shared space. Personal **views**, which none of the tools implement currently, may address some aspects of this issue: Users may often want to restrict the view to reflect their

---

[2]In fact, Collaborative Protégé does support editing of OWL ontologies, but the Challenge ontology was set up in the Protégé frames format.

[3]`https://km.aifb.uni-karlsruhe.de/ws/ckc2007/StoneSoup-www2007.pdf`

personal decisions, or their trust network, for example.

Finally, it remains to be seen whether "complex" ontologies are even amenable to collaborative editing. Some users felt that the space is sufficiently complex to be amenable to this type of work. As one of the participants pointed out *"It's hard enough to agree on ontological decisions when people are sitting at the same table; I doubt that any nontrivial ontology can be built using only asynchronous communication such as the discussion tool."* At the same time, it is possible that the tool-supported environment can actually help by providing a record of changes and comments, recording the discussion, and directly linking it to the components in the ontology, and helping frame the context of the discussion.

While our goal was to gain some insight in what users want from tools for collaborative knowledge constructions, as a side effect, we also learned some lessons on running the challenges of this sort. First, the initial period where the Challenge organizers tried each tool as it was supposed to be used in the Challenge was critical. In fact, every single tool had some problems (either bugs or wrong setup) and these problems were corrected during this period.

Second, the user-competition component seems to have indeed encouraged participation. There were several users who were clearly racing to be the winners and entered large amount of information particularly as the end of the Challenge drew near. In the end, the two most active uses entered more than 2,000 assertions each over the course of the two weeks.

Finally, it appeared that users largely followed the representation patterns that the Challenge organizers established when entering the initial data. This phenomenon, known as cumulative advantage [19], is probably a necessary evil in this setting. Not having any initial data in the tools from which to model their input would have probably intimidated many Challenge participants and discouraged them from entering anything. Yet, entering this data seems to have constrained what the users ended up entering.

7 Conclusions

Although the tools aim at different scenarios for collaborative knowledge construction, it seems that there is a common list of features that users wish to see in all such tools, such as having private and public spaces, being able to import and export in various formats, editing of hierarchies and properties, provenance and change tracking, visualizations and nice and simple interfaces.

However, demand and interest in features for discussions and ratings were much higher for tools for editing ontologies than for tools for editing simple tag hierarchies and annotations. Users expect ontologies to be agreed and verified by the community, but the requirements of consensus are less stringent for less formal knowledge structures, such as tag graphs and instance data.

One of the main outputs of this Challenge was the realization of how the different tools can learn from each other and exchange some of their features. For example several users suggested adding the feature of web resource annotation used by BibSonomy and SOBOLEO to core ontology editing tools such as Collaborative Protégé and Hozo to be able to annotate concepts with external web resources. The authors of BibSonomy and SOBOLEO are now working on integrating these two tools together to bridge the gap between the tag-based approach of BibSonomy with the more ontology-based approach of SOBOLEO.

Acknowledgments

References

[1] BibSonomy. http://www.bibsonomy.org/.

[2] BOWiki. http://onto.eva.mpg.de/bowiki/.

[3] Dbin. http://dbin.org/.

[4] Freebase. http://freebase.com.

[5] Hozo. http://www.hozo.jp/.

[6] OntoWiki. http://3ba.se.

[7] Platypus Wiki. http://platypuswiki.sourceforge.net/.

[8] The Protégé project. http://protege.stanford.edu.

[9] Semantic MediaWiki. http://meta.wikimedia.org/wiki/SemanticMediaWiki.

[10] SOBOLEO. http://soboleo.fzi.de:8080/webPortal/.

[11] S. Auer, S. Dietzold, and T. Riechert. OntoWiki–a tool for social, semantic collaboration. In *Fifth International Semantic Web Conference, ISWC*, volume LNCS 4273, Athens, GA, 2006. Springer.

[12] A. Gibson, K. Wolstencroft, and R. Stevens. Promotion of ontological comprehension: Exposing terms and metadata with web 2.0. In *Workshop on Social and Collaborative Construction of Structured Knowledge at WWW 2007*, 2007.

[13] B. Good, E. Tranfield, P.Tan, M. Shehata, G. Singhera, J. Gosselink, E. Okon, and M. Wilkinson. Fast, cheap and out of control: A zero curation model for ontology development. In *Pacific Symposium on Biocomputing*, 2006.

[14] A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme. BibSonomy: A social bookmark and publication sharing system. In *Proceedings of the Conceptual Structures Tool Interoperability Workshop at the 14th International Conference on Conceptual Structures*, pages 87–102, 2006.

[15] N. F. Noy, H. Alani, G. Stumme, P. Mika, Y. Sure, and D. Vrandecic, editors. *Workshop on Social and Collaborative Construction of Structured Knowledge (CKC 2007) at WWW 2007*. Banff, Canada, 2007. Available on-line at http://ceur-ws.org/Vol-273.

[16] E. Sunagawa, K. Kozaki, Y. Kitamura, and R. Mizoguchi. An environment for distributed ontology development based on dependency management. In *Second International SemanticWeb Conference (ISWC2003)*, Sanibel Island, FL, USA, 2003.

[17] G. Tummarello, C. Morbidoni, and M. Nucci. Enabling semantic web communities with DBin: an overview. In *Fifth International Semantic Web Conference, ISWC*, volume LNCS 4273, Athens, GA, 2006. Springer.

[18] L. von Ahn and L. Dabbish. Labeling images with a computer game. In *SIGCHI conference on Human factors in computing systems*, pages 319–326. ACM Press New York, NY, USA, 2004.

[19] D. Watts. *Six Degrees: The New Science of Networks*. Vintage, 2004.
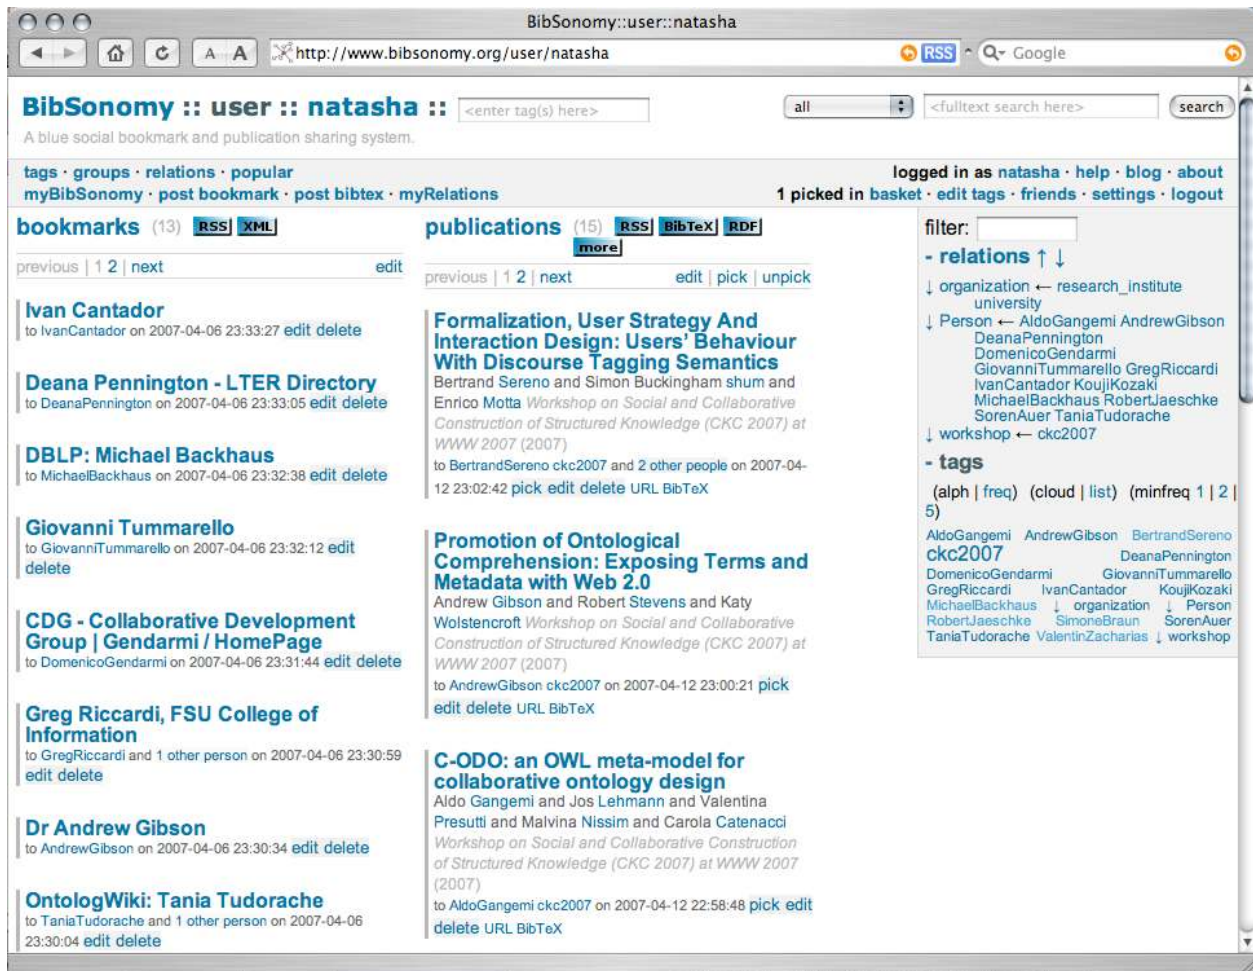
Figure 1: BibSonomy: The personal space (*myBibSonomy*) of a user. The tagged bookmarks and publications are on the left; the relations between tags and the user's tags are on the right.
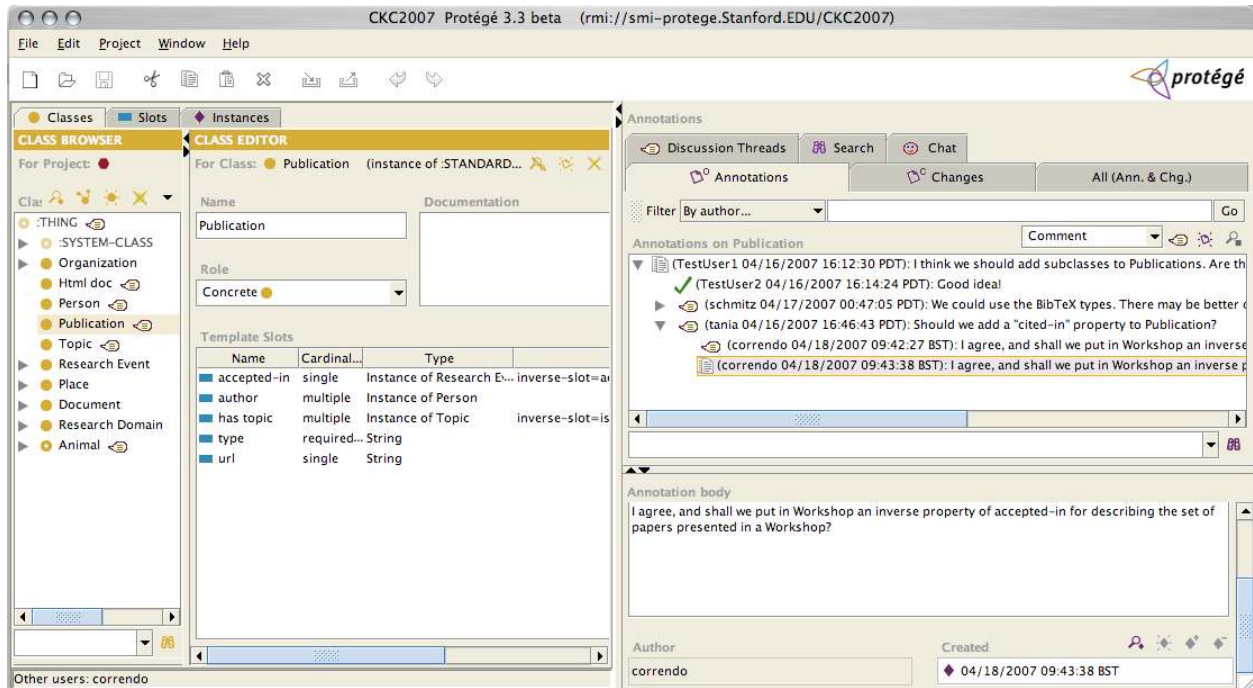
Figure 2: Collaborative Protégé: The Collaborative Protégé editor window has two main panels: the classes, properties and instances view on the left and the collaborative panel on the right. The figure shows an example of annotations attached to the `Publication` class. Users can view discussion threads, annotations, and changes gpt a particular class by selecting the class in the class tree on the left and the respective tab in the collaborative panel.
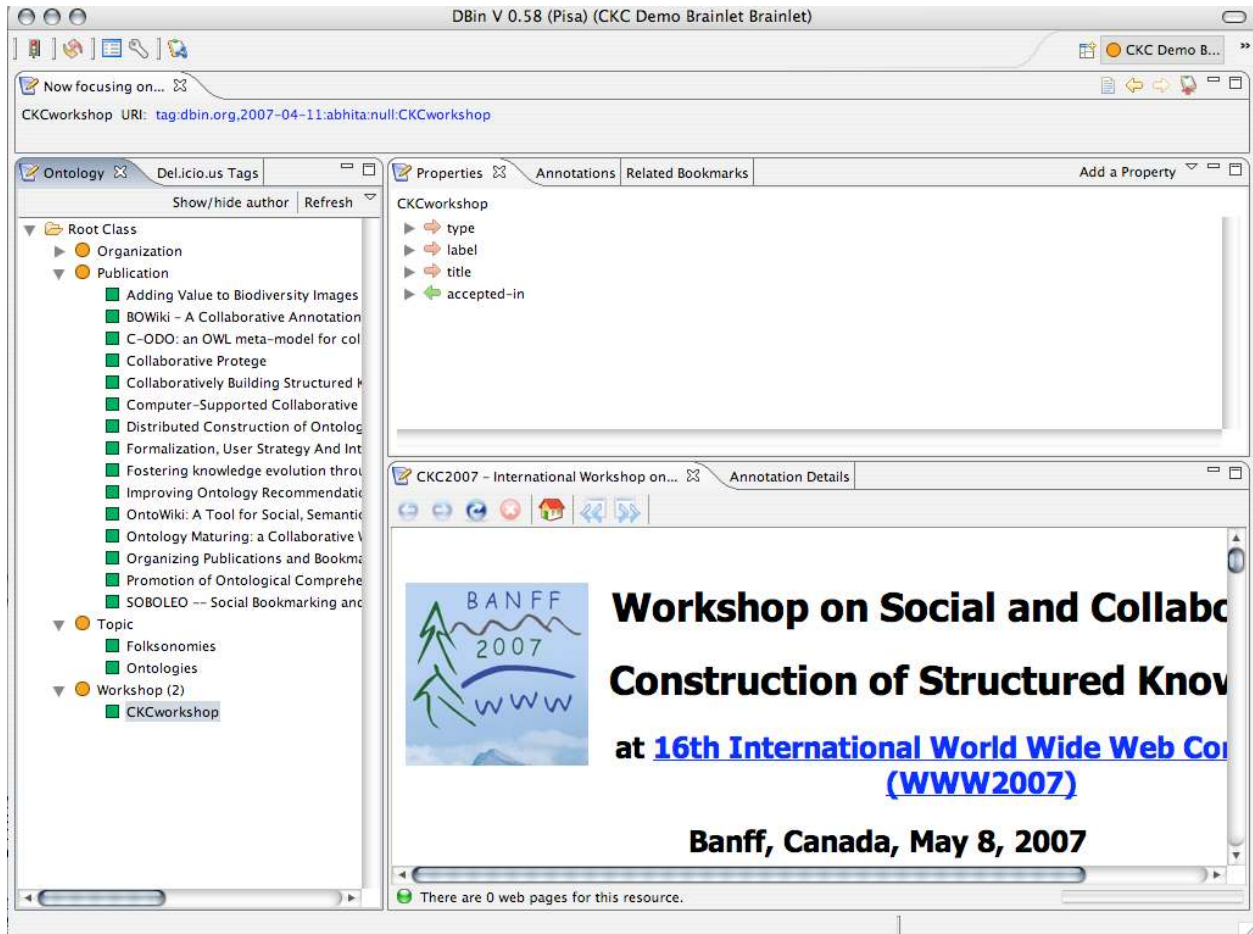
Figure 3: DBin: The class and instance hierarchy is on the left and the panel on the right enables users to annotate resources. DBin allows users to create domain specific user interfaces and this figure shows one such possible user interface
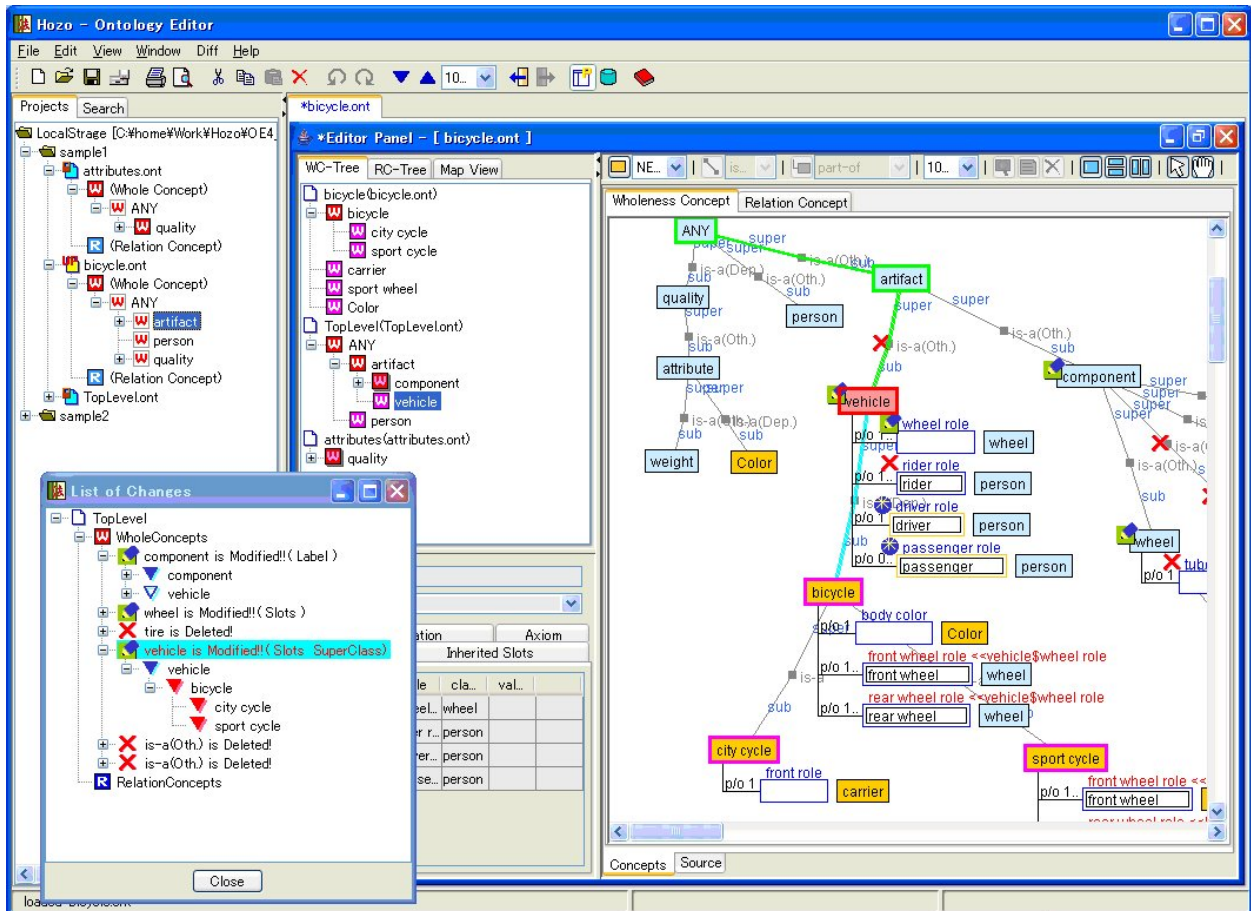
Figure 4: Hozo: The panel on the left shows the class hierarchy and the list of changes. The panel on the right shows the dependencies between concepts.

17
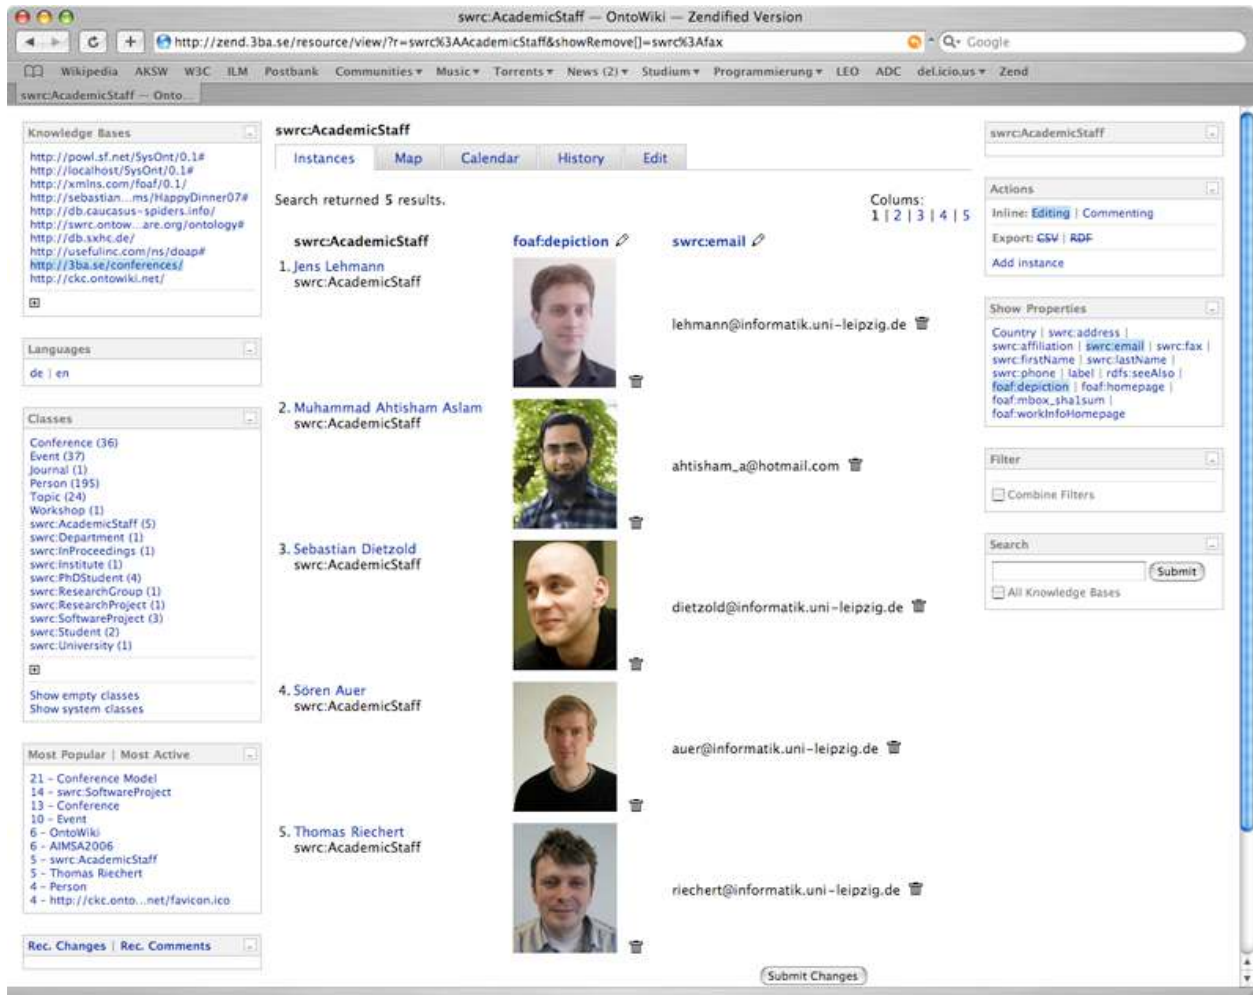
Figure 5: OntoWiki: The user interface is divided into 3 parts: the left sidebar, a main content section and a right sidebar. The class browser is displayed in the left sidebar. Once a user selects a class, instances of the class appear in the main content section which link to views for individual instances. The right sidebar offers filtering and editing tools and information specific to the selected content.

Figure 6: SOBOLEO: The interface for editing the taxonomy. The left column shows the hierarchy of tags; the editing panel is in the middle and allows user to define concept names and relationships between concepts. In the figure above, Organization is defined as a broader concept of FZI: every document about FZI is also about Organization. The messages panel on the right side enables users to chat with other users editing the taxonomy at the same time.
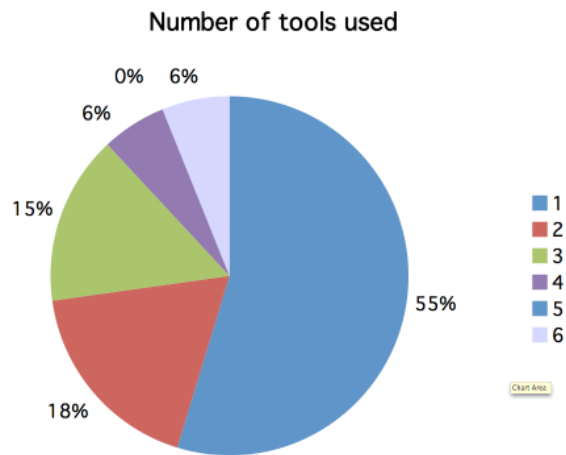
Figure 7: 55% of active users tried out 1 tool; 6% tried out all 6 tools.
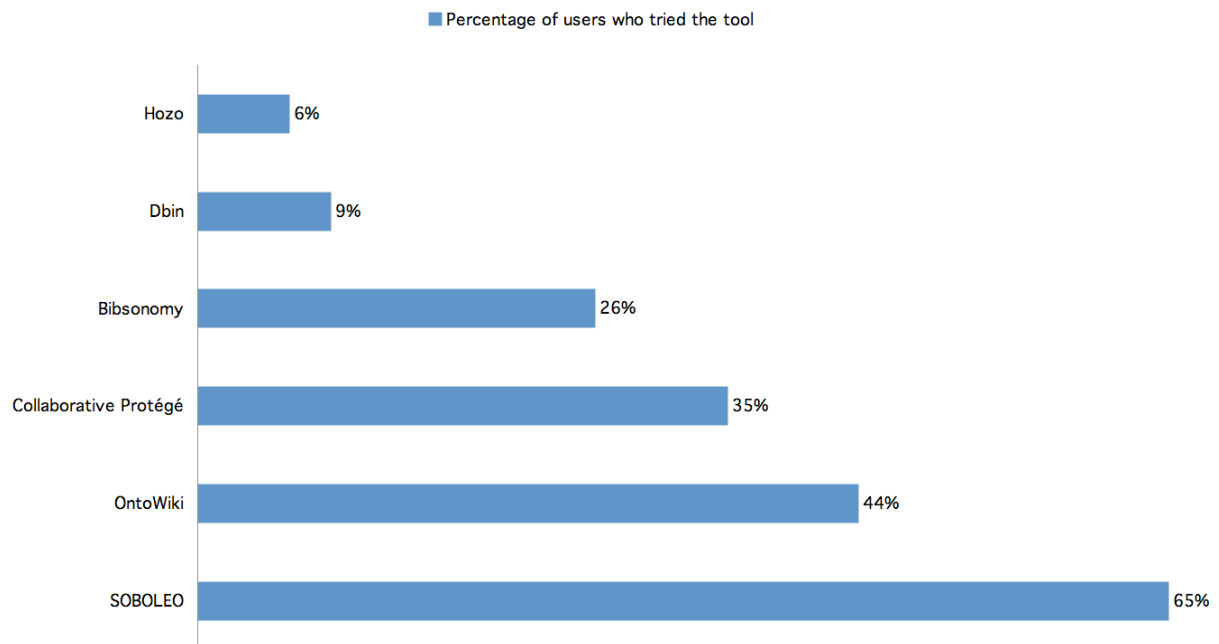

Figure 8: 65% of the 33 active participants tried out SOBOLEO; 44% used OntoWiki and 35% used Collaborative Protégé

Which features from the Challenge tools would you incorporate into your perfect tool?

I would use Hozo and Soboleo for creating my GUI for ontology. Importing and exporting from BibSonomy. Connections between users from Protégé, and history from OntoWiki

Bibsonomy's method of posting bookmarks and publications, Collaborative Protégé's method of creating instances of classes. OntoWiki's map feature.

Figure 9: The features from the "perfect tool" for collaborative construction of structured knowledge, according to some Challenge participants.