FEMwiki: Crowdsourcing Semantic Taxonomy and Wiki Input To domain Experts While Keeping Editorial Control: Mission Possible!

Patty Kostkova
University College London
London
UK

p.kostkova@ucl.ac.uk

Vladimir Prikazsky
European Centre for Disease
Prevention and Control (ECDC)
Sweden

Vladimir.prikazsky@ecdc.europa.eu

Arnold Bosman
European Centre for Disease
Prevention and Control (ECDC)
Sweden

Arnold.Bosman@ecdc.europa.eu

ABSTRACT

Highly specialized professional communities of practice (CoP) inevitably need to operate across geographically dispersed area - members frequently need to interact and share professional content. Crowdsourcing using wiki platforms provides a novel way for a professional community to share ideas and collaborate on content creation, curation, maintenance and sharing. This is the aim of the Field Epidemiological Manual wiki (FEMwiki) project enabling online collaborative content sharing and interaction for field epidemiologists around a growing training wiki resource.

However, while user contributions are the driving force for content creation, any medical information resource needs to keep editorial control and quality assurance. This requirement is typically in conflict with community-driven Web 2.0 content creation. However, to maximize the opportunities for the network of epidemiologists actively editing the wiki content while keeping quality and editorial control, a novel structure was developed to encourage crowdsourcing – a support for dual versioning for each wiki page enabling maintenance of *expert-reviewed* pages in parallel with *user-updated* versions, and a clear navigation between the related versions.

Secondly, the training wiki content needs to be organized in a semantically-enhanced taxonomical navigation structure enabling domain experts to find information on a growing site easily. This also provides an ideal opportunity for crowdsourcing. We developed a user-editable collaborative interface crowdsourcing the taxonomy live maintenance to the community of field epidemiologists by embedding the taxonomy in a training wiki platform and generating the semantic navigation hierarchy on the fly. Launched in 2010, FEMwiki is a real world service supporting field epidemiologists in Europe and worldwide. The crowdsourcing success was evaluated by assessing the number and type of changes made by the professional network of epidemiologists over several months and demonstrated that crowdsourcing encourages user to edit existing and create new content and also leads to expansion of the domain taxonomy.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

DH'15, May 18–20, 2015, Florence, Italy.

ACM 978-1-4503-3492-1/15/05.

<http://dx.doi.org/10.1145/2750511.2750523>

Categories and Subject Descriptors

algorithms, semantic web, real world assessment, semantic navigation

General Terms

crowdsourcing, semantic web, evaluation, field epidemiology, taxonomy

Keywords

FEMwiki, social and Semantic Web, user engagement, evaluation

1 INTRODUCTION

In many modern disciplines professional experts are often widely geographically dispersed. It has become essential to maintain a single repository of knowledge about the domain online to avoid error-prone practices such as sending information from person to person by email. Experts increasingly desire to be able to contribute to a shared repository using web 2.0 technology such as wikis and crowdsourcing to the community what traditionally was developed by experts committees. However, crowdsourcing shared professional content development to real world users requires easy-to-use tools for domain experts to make their contributions. While the collaborative aspect of developing the repository is important, it is also vital to ensure that the quality of the repository is maintained. This required is unquestionably of paramount importance in the medical domain. However, editorial controls should not stifle the pace of contribution to the portal. Therefore, it is important to provide user friendly Web 2.0 tools for experts to collaboratively maintain a knowledge repository online, while at the same time, provide an editorial control system that maintains quality, but does not interfere excessively with the process of updating the resource. Further, the crowdsourced wiki content to potentially hundreds of users needs to be organized in a semantically-enhanced taxonomical navigation structure enabling domain experts to find information on a growing site easily, one that is easy to maintain by domain experts as the project grows. Both features, content creation and taxonomy maintenance require active cooperation from the domain experts.

In this paper, we present the Field Epidemiology Manual Wiki (FEMwiki) Framework crowdsourcing model enabling collaborative editing of the actual content as well as navigation taxonomy. FEMwiki (www.femwiki.com),

funded by the ECDC (European Centre for Disease Prevention and Control), is used by field epidemiologists to maintain a repository of knowledge used for training purposes. FEMwiki has its origins in a training manual for the EPIET training course (European Program for Intervention Epidemiology Training), and was converted into an wiki-style repository using crowdsourcing.

FEMwiki framework is structured using a domain taxonomy editable by users in the same way as the actual content. The taxonomy browser on the front page of the wiki allows users to immediately see and navigate the organisation of the repository (Figure 1).

Section 2 provides a background to the project and sets the scene for section 3 where we present the FEMwiki crowdsourcing framework for both wiki editing and semantic taxonomy development. In section 4, we discuss the evolution of the project and evaluation results with real world field epidemiologists. Section 5 brings discussion while section 6 concludes.

2 BACKGROUND AND RELATED WORK

Crowdsourcing owns its growing popularity to the simple fact that a large number of users can make a small effort on a shared task enabling a large scale collaborative work performed easily [1]. Crowdsourcing has many forms and could be implemented over a number of platforms. Typically, collaborative Web 2.0 technologies enable users to create and modify content in a shared repository instead of merely being passive consumers. In addition to sharing the work, the risk of bottlenecks is reduced. The most well known example of crowdsourcing must be Wikipedia¹ with over 4.7 million articles, being increased every day with over 800 new articles as of March 2015. However, user contributions remain sparse. Wikipedia has also been studied as a cultural phenomena reaching trusted level of information through crowdsourcing [2].

Large wikis such as Wikipedia can be difficult to navigate, as they are large at repositories and there is no native support for organising the content. However, for domain-specific wikis, this problem could be overcome by organizing the pages according to semantic taxonomy representing the domain entities (also called nodes) forming the basis for content navigation using the parent-child relationship (entity becomes a wiki page). Semantic ontologies and taxonomies are used in a wide variety of disciplines. Perhaps their most notable successes are in the life sciences and medicine (for example, the Open Biological and Biomedical Ontologies (OBO) [3], MeSH², ICD-11³ and SNOMED-CT⁴). In such domains, the ontologies are usually highly formal but require a considerable amount of expertise, time and effort to build. Stevens and Jupp [4-5] argue that many other medical ontologies are rather taxonomies as they do not follow first order logic relationships between entities but better describe the complex medical domain.

Figure 1. The taxonomy browser - FEMwiki front page

The notion of using crowdsourcing to develop domain taxonomies is attractive, as a large number of experts in the domain can each make contributions (small or large), rather than a small number of experts expending a large amount of effort. The resulting taxonomy would (in theory) respect a consensus view on the domain, rather than the view of a smaller number of experts. However, despite the experts' efforts, the resulting web 2.0 taxonomy often contains gaps and errors. Also the taxonomy is never final or static - there is a need to fill in gaps and modify the taxonomy structure and maintain it as the domain develops. A good example of an ontology generated by crowdsourcing is DBpedia⁵ [6], where structured data is extracted from Wikipedia, and made available on the Web. In this case, the contributors do not create the ontology directly, but make edits to Wikipedia, and are probably not aware that some of their contributions are being used to create an ontology. Semantic wikis are more direct attempts to combine semantics with crowdsourcing. The users are generally aware that they are also contributing structured data as well as human-readable text. In semantic wikis, each page (also called an article) corresponds to an entity in the resulting ontology (either a class, individual, or property). The most widely used semantic wiki is perhaps Semantic MediaWiki (SMW)⁶, built on the MediaWiki platform that is used for Wikipedia. However, attempts have been made to provide easier-to-use tools usable by domain experts. Project Halo [7] is an extension to SMW that provides a semantic toolbar to allow page editors to annotate pages. However, the arrangement of pages into a class hierarchy is performed by placing semantic annotations directly into the wikitext, which is probably not ideal for most IT non-specialist users. OpenDrugWiki [8] is a semantic wiki system that holds drug interaction data. The wiki pages can be edited by any registered user, but to maintain quality the peer review mechanism is performed by a set of editors manually checking all edits. Only approved pages are made available for querying posing a clear issue with scalability. HJ Jung proposes quality assurance in crowdsourcing using matrix factorisation [9]. Further, the NeLI⁷ project provides another example of domain experts-led development on infection taxonomy using SKOS/owl

¹ http://en.wikipedia.org/wiki/Main_Page (English version)

² <http://www.nlm.nih.gov/mesh>

³ <http://www.who.int/classifications/icd/en/>

⁴ <http://www.ihtsdo.org/snomed-ct/>

⁵ <http://dbpedia.org> (the ontology is accessible at <http://wiki.dbpedia.org/Ontology>)

⁶ <http://semantic-mediawiki.org/>

⁷ <http://www.neli.org.uk>

[10], however, in this case, assistance of the computer science researchers was required to facilitate the process. SweetWiki [11] is a system with a WYSIWYG ontology editor, based on semantic tagging using a wiki object model. However, unlike SweetWiki, we use the actual wiki structure for developing the semantic taxonomy. The importance of user-friendly interface to semantic technologies has been highlighted by Madle et al [12] and Oliver [13]. Crowdsourcing intelligence for semantic has been also argued for by Auer and Kontokostas [14].

Therefore, designing a collaborative Web 2.0 wiki utilizing crowdsourcing while keeping editorial control over quality remains an issue. Further, engaging users in semantic navigation taxonomy maintenance for domain wikis remains attractive but a suitable user-friendly interface is essential for unsupervised collaborative input from domain experts.

In this paper we propose a solution to these two problems: the FEMwiki framework and conduct the initial evaluation. User engagement has been a growing discipline with developed models for assessment [15] but methods for encouragement remain an open problem.

3 THE FEMWIKI CROWDSOURCING FRAMEWORK

FEMwiki consists of a wiki-based repository, user forums for discussion of wiki pages, and user personal profiles where users can give more information about themselves.

The schematic organisation of the wiki part of FEMwiki is shown in Figure 2. Wiki pages - representing a term from the domain of field epidemiology - may contain texts and graphics, and are organised into a hierarchical structure.

Figure 2. A schematic diagram of the FEMwiki framework content organisation.

In addition to navigating by following links to other wiki pages, the main organisational feature is a navigation hierarchy of semantically connected parent-child wiki pages derived from semantic taxonomy of the epidemiology domain. Although there is no strict meaning given to the parent-child relationship between wiki pages in the actually platform, semantically connected parent-child pages create a naturally formed navigation taxonomy representing the semantics of the domain knowledge. Each node in the taxonomy is a wiki page, which can have text and graphics, as well as child pages. Additional features, such as tagging, may assign more terms to a single wiki page flexibly. Pages can also contain cross reference (untyped)

hyperlinks to other pages anywhere in the wiki. In the framework it is not required that there is a single root node (there can be multiple disjoint trees), but to avoid confusion we will assume in the following that there is such a root node.

Figure 3. The position of a page in the hierarchy altered during editing.

The taxonomy browser is immediately visible on the front page of the wiki giving users the opportunity to visualize the domain taxonomical structure. This browser page is instantly regenerated when changes are made allowing users to see updates instantly. More importantly, users can enjoy a seamless experience as taxonomy changes are done through the same wiki interface as updates to normal wiki content. The taxonomy can be altered by specifying a parent page while editing any wiki page (Figure 3). In particular, the position of a page in the hierarchy can be altered during editing, by specifying a new parent page (if no parent is specified, the page is placed at the topmost level).

The taxonomy and all wiki pages are viewable by the public. In order to edit a page, and to post in the forums, users must create an account. As in other wiki systems, page histories are stored. While versioning is a typical feature of wikis, it has a special importance for the editorial control system, see in the following section.

3.1 FEMwiki Dual Versioning Editorial Model

While the system aims to encourage any registered user can add to or change a wiki page, it is important to ensure that the content can be trusted. In the FEMwiki framework, the mechanism ensuring this is assigning an editorial role to senior domain experts to approve specific pages and keep dual version system in parallel. Editors are senior epidemiologists assigned by FEMwiki management at ECDC to ensure approved content is of high quality, up-to-date and strictly evidence based.

The expert-reviewed version of a page (Figure 4), is displayed with a colour-coded green bar at the top of the page containing a link to the latest unapproved version (if one exists). The editor and other contributors to the page content are listed on the right hand side. The latest unapproved version of a page (Fig 4) has a yellow bar at the

top, with a link to the expert-reviewed version (if one exists). The editor of the page can approve this version by clicking a button (not visible to any other readers).

The page history is not affected by the dual editorial support. Only the latest version of a page can be edited. There is only one expert version permitted at a given time (or none). Thus, if an expert-reviewed version of a page exists, together with one or more later versions, only one of these later versions could be later approved as the "new" expert-reviewed. Thus, the page history that is stored is a sequence of pages, together with the version number of the latest expert-reviewed page (Figure 5).

Figure 4. The expert-reviewed and latest versions of a wiki page.

Figure 5. The page history at FEMwiki framework, illustrating expert-reviewed and unapproved pages.

The content on the right hand side of the wiki pages illustrates the users involved in the editorial process, thus adding an extra layer of transparency and trust, unusual in other wiki projects.

3.2 FEMwiki Semantic Navigation Model

One of the main challenges of implementation of semantic technologies is the cost of development and maintenance of domain ontologies and taxonomies. This is of particular importance in life critical domains where the need to keep the ontology up to date is paramount. User-friendliness of ontology editors is another challenge. In the FEMwiki framework, we utilized the wiki user interface users have been using for collaborating editing of the medical content for entirely different purpose: the wiki page also serves as a user-friendly taxonomy editor, thus, offering a seamless experience to users.

Therefore, in order to elicit more edits from users the entire field epidemiology taxonomy is displayed on the navigation page (rather than just pages with existing content). A colour coding is used to draw user attending to empty pages ("stubs") and to distinguish between various types content, see Figure 6. The taxonomy editor supports colour-coding for the dual versioning of pages: (A) YELLOW: link to the latest version of the page (B) GREEN: link to the expert-reviewed (and approved) page, clicking on the text "approved version" will lead to the reviewed version. Further, (C) QUESTION MARK: pages that do not have an expert-reviewed version (indicated by the question mark icon), the link will lead to the latest version, and finally, (D) GREEN ONLY: indicates pages where the latest version is also the expert-reviewed version, the link leads to this common version. Any edits to the page will cause a new latest version to be created. Finally, (E) RED illustrates (and visually draws attention to) to pages tagged as "stubs" where content has not been developed yet.

By simply looking at the colour-coded taxonomy browser, user can see which pages have expert-reviewed versions, and can either choose to see that version, or a later unapproved version if one exists (Figure 6). The user can also see "stub" pages marked in red - this feature is specifically designed to highlight parts of the wiki content that need to be filled in, and to encourage users to start this process.

Figure 6. The FEMwiki taxonomy browser
 (A) latest version
 (B) expert-reviewed page
 (C) no expert version exists, only latest version
 (D) latest version is also the expert-reviewed version
 (E) empty “stubs” marked in red.

4 THE FEMWIKI EVOLUTION AND EVALUATION RESULTS

As outlined in the Introduction section, the Field Epidemiology Manual Wiki (FEMwiki) [16], funded by the ECDC (European Centre for Disease Prevention and Control), is used by field epidemiologists to maintain a repository of knowledge used for training purposes. It was developed and hosted by City ehealth Research Centre until January 2012, and was subsequently migrated to ECDC. FEMwiki was developed using Telligent Community software⁸ (which provides typical wiki functions, such as editing, and conflict resolution). The FEMwiki serves primarily field epidemiologists in Europe - this community of practice was investigated by Fowler et al [17].

4.1 The evolution of the FEMwiki content

The basis of the FEMwiki was the a training manual developed by a training programme run by ECDC, EPIET, which was organised into 17 chapters. Each chapter was originally written by a trainer/lecturer in the EPIET programme - the manual was intended to be studied and also taught like a textbook, from start to finish. During the process of converting the manual to FEMwiki, an editorial board was appointed to oversee the process of reviewing each chapter and converting it into a wiki page(s) [16]. Where possible, these were the original lecturers, otherwise, new senior experts were appointed.

The first version of FEMwiki retained the chapter sequence of the EPIET manual, with a home page for each chapter. The semantic nature of the FEMwiki platform (i.e. the taxonomy browser) was not utilised at this initial stage (see Figure 7-1). Utilising the FEMwiki framework potential for semantic taxonomic representation and navigation provided an opportunity to develop a taxonomy of public health and field epidemiology not covered by existing medical taxonomies. In order to organise the content, a taxonomy was developed in consultation with

domain experts in Stockholm and London which attempted to cover the knowledge required to train field epidemiologists (Figure 7-2).

The navigation taxonomy has undergone a major development during the project to enhance the simplicity and actively engage users. A noticeable feature of the resulting taxonomy is that it is still somewhat linear, with features of a course to be followed rather than a classification of a domain. Another issue is that the names of some nodes were not self-contained. For example, the name “Evaluation” was used for two different nodes in different places (under both *Establishing a surveillance system* and *Screening*) in the hierarchy. These were not intended to be the same node, but rather should have been named in a context-independent way (for example *Evaluation of a surveillance system*).

A mapping exercise was then carried out, linking the initial FEMwiki pages to the taxonomy terms, which led to renaming of these context-dependent terms to make them self-explanatory and standalone in the taxonomical hierarchy (Figure 7-3). Finally, the FEMwiki was edited collaboratively, with pages and the taxonomy able to be edited online (Figure 7-4).

Figure 7. FEMwiki evolution

4.2 Evaluation study design

How successful the FEMwiki crowdsourcing framework actually is with real world users?

Our evaluation aimed to find the number and types of changes that have been made by the domain experts in editing FEMwiki. The FEMwiki project was launched in November 2010 with the original EPIET chapter-driven structure. The new taxonomical navigation interface was launched in September 2011 which was the starting point for evaluating the evolution with real domain users (Figure 7).

⁸ <http://telligent.com>

For evaluation purposes, we discuss two distinct versions of the FEMwiki content and structure:

- (1) the taxonomy derived from consultation with experts, completed in September 2011 and used as the basis for the first version of FEMwiki (Figure 7-2);
- (2) a snapshot of the FEMwiki content in early January 2012
- (3) a snapshot of the FEMwiki content hosted by ECDC, taken from May 2012.

It is important to realise that the taxonomy in version (1) was not used directly in the FEMwiki implementation, but was used as part of the process of reorganising the content in September 2011. We performed an initial exploration of the changes between the taxonomy version (1) and later versions. The changes that we found were of several types. Firstly several nodes were renamed and many were simple alterations, for example *Classifying / measuring risks* became *Classifying and measuring risk*. Others were less trivial changes, for example, *Evaluation* changed to *Evaluation of surveillance systems*.

This latter example is a case of renaming to make a node context-independent. Some nodes were moved to a different location in the hierarchy between versions. There was also a "flattening" of the hierarchy between the two versions. For example, the node *Analysis by person characteristics* is in the fourth level of the hierarchy in version (1), but in the third level in version (2). Some superfluous intermediate nodes were removed in order to simplify the structure of the hierarchy. We mainly concentrated subsequently on comparing versions (2) and (3), where the changes between the versions were entirely made online.

In our evaluation, we specifically considered two important aspects of change to the FEMwiki content:

- (a) evolution of the semantic taxonomy: which terms were created, deleted, moved, and renamed, and also the shape of the taxonomy tree, and
- (b) evolution of page content: changes to existing pages; new content added to "stubs"

We make comparisons between taxonomies snapshots taken in January 2012 and May 2012. Table 1 shows a summary of the changes between versions. The measure "*inheritance richness*" [18] is the average number of subclasses for each non leaf node in the hierarchy. A higher number indicates a wider tree, and a lower number indicates a taller tree (for trees with the same number of nodes). In our case, the differences in inheritance richness do not seem significant.

4.3 FEMwiki taxonomy: Results

In order to make comparisons between versions of the FEMwiki taxonomies, we extracted OWL class hierarchies from the SQL databases used in the FEMwiki framework to store information about wiki pages. The class hierarchies were then compared by counting nodes, and by using the PROMPT [19] plug-in to Protege to find changes between the successive versions.

Between the two versions at the study period (January 2012 and May 2012), 12 new pages were created, including 8 pages on competency requirements, pages on EU legislation, and mathematics (*Probability*). 17 pages were deleted, although it seems likely that most of the content of these pages was transferred into other pages. Inheritance richness decreased from 3.77 to 3.65 illustrating the taxonomy tree was getting "taller".

Table 1. Comparison of FEMwiki versions

	(2) FEMwiki (Jan 2012)	(3) FEMwiki (May 2012)
Total number of nodes	283	278
Number of stub nodes	90	75
Number of new nodes	n/a	12
Number of deleted nodes	n/a	17
Number of renamed nodes	n/a	0
Inheritance richness	3.77	3.65

There were several rearrangements of the taxonomy. Two new top level categories were created, *Public health law* and *Public health informatics*. A page under the category *Uncategorised* was moved to the new *Public health law* category. More complex changes were also performed, for example: the page *Case definitions* together with its three child pages was moved from one branch of the taxonomy to another.

4.4 FEMwiki content: Results

As it is possible to reconstruct the history for page contents (which was not possible for the taxonomy), in general we can give results as graphs over time, instead of just comparing snapshot versions. However, the number of stub nodes cannot be easily reconstructed in this way, so we have to look at the two snapshots. From Table 1 we can see that the number of "stub" nodes has reduced from 90 to 75 between January and May 2012, indicating that the content is gradually being filled in. It is likely that the high visibility of the stub nodes in the taxonomy browser acts as a prompt for users to add content, thus produces active engagement with the site. Fully understanding this phenomena would require further research and user feedback.

In the FEMwiki portal, there are over 1000 registered users in 2015 (at the end of the evaluation period the total number was 814). However only a small core of these are actively involved in making changes to the wiki (this is not usual for online communities, as discussed for example by Preece et al. [20]). the overall aim for the development of the wiki resource is to encourage more of the inactive users to start making contributions. There have been wiki page edits made by a total of 32 different registered users. On the forums for discussing pages, 37 distinct users have made comments. The overlap between these groups (those who edit and post on the forum) is 20 users. Figure 8 shows the monthly number of page edits and forum posts on FEMwiki. Following the launch in 2010, there was a steady increase in user registration. After the new system with the editable taxonomy was launched in November 2011 there was another sharp increase in user registrations, and a large amount of activity while changes to the taxonomy and wiki pages were made.

Figure 8. FEMwiki monthly numbers of page edits and forum posts

5 DISCUSSION AND FUTURE WORK

Crowdsourcing has numerous forms and models. In this study, we demonstrated novel framework, FEMwiki, enabling collaborative development of a domain taxonomy and wiki content through a single user-friendly interface and successfully evaluated the FEMwiki portal evolution with real-world users.

However, there are some interesting lessons learned. The transition from the original EPIET training manual to the current FEMwiki knowledge repository and the subsequent online web 2.0 evolution shows a clear pattern of moving away from a linear course structure towards a collection of articles based on the structure of the domain. Also the number of levels in the hierarchy was reduced, which seems to indicate a movement towards simplifying the organisation of the content (which is perhaps not so important for students following a course structure). It is important to realise that the process is still ongoing, as pages may be added, stub pages may be modified, and possibly complex rearrangements will take place to the taxonomy.

Future work, jointly with ECDC, includes investigation of specific incentives increasing users' interest in contribution, either to the taxonomy or to the wiki itself and further evaluation of content and taxonomy growth over the years. This is currently an ongoing project aiming to identify and understand the key characteristics of online community development and long-term sustainability.

6 CONCLUSIONS

Crowdsourcing using wikis platforms provides a unique way for a geographically dispersed professional communities of practice to collaborate on content creation, curation, maintenance and sharing. User-friendly navigation structure allowing a quick and easy access to growing collection of resources is required to support the content growth.

Launched in 2010, FEMwiki is a real world social combining collaborative wiki service supporting field epidemiologists in Europe and worldwide. With over 100 000 page views, over 1000 users signed up, ongoing user-driven semantic taxonomy updates and maintenance by crowdsourcing to the professional

network, the FEMwiki is a unique example of a successful crowdsourcing research with a real-world impact.

The portal features two novel models: dual versioning enabling to keep user-generated version of pages next to expert-reviewed, and user-editable semantic taxonomy. We described the effect of the crowdsourcing strategy in terms of enhancements to the wiki and the taxonomy in the initial phase: semantic taxonomy with the same interface to modify the parent-child tree structure as the pages themselves further engages users in contribution and resulted in 12 new nodes in taxonomy (not possible to add without editable taxonomy) and 15 filled stubs which otherwise would not be created. 4% of registered users contribute new content and those who contribute to editing pages are also active in discussion forums. Overall, the evaluation with real users demonstrated that the FEMwiki platform provides an interface likely encouraging domain experts to contribute to the wiki content and the taxonomy through a single interface, however, further research is required to better understand incentive strategies for long-term engagement and retention.

7 ACKNOWLEDGMENTS

We acknowledge ECDC for funding this project, Julius Weinberg and Mike Catchpole for their input into the FEMwiki taxonomy. Many thanks also go to former CeRC researchers Martin Szomszor and Simon Hammond for the implementation of the current version of the FEMwiki and David Fowler for his input into previous versions of this paper and the research study.

8 REFERENCES

- [1] J Howe. The Rise of Crowdsourcing. Wired, June 2006
- [2] S. Lindgren. Crowdsourcing Knowledge. Culture Unbound, Volume 6, 2014, pp. 609–627, 2014
- [3] B Smith, M Ashburner, C Rosse, J Bard, W Bug, W Ceusters, LJ Goldberg, K Eilbeck, A Ireland, CJ Mungall, N Leontis, P Rocca-Serra, A Ruttenberg, S-Assunta Sansone, RH Scheuermann, N Shah, PL Whetzel, and S Lewis. The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. Nature biotechnology, 25(11):1251, 5, November 2007.
- [4] S Jupp, R Stevens, S Bechhofer, P Kostkova, and Y Yesilada. Semantic navigation: Ontology or Knowledge Organisation System? In NETTAB2007 A Semantic Web for Bioinformatics Goals Tools Systems Applications, 2007.
- [5] S Jupp, R Stevens, S Bechhofer, Y Yesilada, and P Kostkova. Knowledge Representation for Web Navigation. SWAT4LS, 2008.
- [6] C Bizer, J Lehmann, G Kobilarov, SSoren Auer, C Becker, R Cyganiak, and S Hellmann. DBpedia - A crystallization point for the Web of Data. Web Semantics Science Services and Agents on the World Wide Web, 7(3):154{165, 2009.
- [7] N S. Friedland, PG. Allen, G Matthews, M Witbrock, D Baxter, J Curtis, B Shepard, P Miraglia, J Angele, S Staab, E Moench, H Oppermann, D Wenke, D Israel, V Chaudhri, B Porter, K Barker, J Fan, SY Chaw, P Yeh, D Tecuci, and P Clark. Project halo: Towards a digital

-
- aristotle. *Artificial Intelligence Magazine*, 25(4):29-47, 2004.
- [8] A Kostlbacher, J Maurus, R Hammwöhner, A Haas, E Haen, and C Hiemke. OpenDrugWiki - Using a Semantic Wiki for Consolidating, Editing and Reviewing of Existing Heterogeneous Drug Data Why Use a Semantic Wiki? In Ch Lange, J Reutelshoefer, S Schaert, and H Skaf-Molli, editors, *SemWiki-2010: 5th Workshop on Semantic Wikis*, Hersionissos, Heraklion, Crete, Greece, 2010.
- [9] HJ Jung. Quality Assurance in Crowdsourcing via Matrix Factorization based Task Routing. *WWW'14 Companion*, April 7-11, 2014, Seoul, Korea. pp. 3-7
- [10] G Diallo, P Kostkova, G Jawaheer, S Jupp, R Stevens. Process of building a vocabulary for the infection domain. 21st IEEE International Symposium on Computer-Based Medical Systems, 2008. *CBMS'08*, pp 308-313, 2008
- [11] M Bua, F Gandon, G Ereteo, P Sander, C Faron. SweetWiki: A semantic wiki. *Web Semantics Science Services and Agents on the World Wide Web*, 6(1):84-97, 2008.
- [12] G Madle, P Kostkova, J Mani-Saada, A Roy. Lessons learned from evaluation of the use of the National electronic Library of Infection. *Health informatics journal*, Vol 12, No 2, pp. 137-151, 2006
- [13] H Oliver, G Diallo, E De Quincey, D Alexopoulou, B Habermann, P Kostkova, M Schroeder, S Jupp, K Khelif, R Stevens, G Jawaheer, G Madle. A user-centred evaluation framework for the Sealife semantic web browsers. *BMC bioinformatics*, Vol 10, S14, 2009
- [14] S Auer, D Kontokostas. Towards Web Intelligence through the Crowdsourcing of Semantics. *WWW'14 Companion*, April 7-11, 2014, Seoul, Korea. pp. 991.
- [15] J. Lehmann, M. Lalmas, E. Yom-Tov and G. Dupret. Models of User Engagement, 20th conference on User Modeling, Adaptation, and Personalization (UMAP 2012), Montreal, 16-20 July 2012
- [16] P Kostkova and M Szomszor. The FEMwiki project: a conversion of a training resource for field epidemiologists into a collaborative Web 2.0 portal. In 3rd International ICST Conference on Electronic Healthcare for the 21st century, 2010, Springer LNICST, pp. 119-126, 2012
- [17] D Fowler, M Szomszor, S Hammond, J Lawrenson, P Kostkova. Engagement in Online Medical Communities of Practice in Healthcare: Analysis of Messages and Social Networks. In 3rd International ICST Conference on Electronic Healthcare for the 21st century, 2010, Springer LNICST, pp. 154-157, 2012
- [18] S Tartir, I. Budak Arpinar, and AP. Sheth. Ontological Evaluation and Validation. In Roberto Poli, Michael Healy, and Achilles Kameas, editors, *Theory and Applications of Ontology: Computer Applications*, chapter 5, pages 115-130. Springer, 2010.
- [19] NF. Noy, H Kunnatur, M Klein, and MA. Musen. Tracking changes during ontology evolution. In *Proceeding of the 3rd International Semantic Web Conference (ISWC2004)*, pages 259-273. Springer Verlag, 2004.
- [20] J Preece, B Nonnecke, and D Andrews. The top twelve reasons for lurking: improving community experiences for everyone. *Computers in Human Behaviour*, 20(2):201-223, 2004.