

Massive Ontology Interface

Matt Stannett
Computer Science Department
University of Waikato
New Zealand
matt.stannett@gmail.com

Catherine Legg
Philosophy Program
University of Waikato
New Zealand
clegg@waikato.ac.nz

Samuel Sarjant
Computer Science Department
University of Waikato
New Zealand
sarjant@waikato.ac.nz

Abstract

This paper describes the Massive Ontology Interface (MOI), a web portal which facilitates interaction with a large ontology (over 200,000 concepts and 1.6M assertions) that is built automatically using OpenCyc as a backbone. The aim of the interface is to simplify interaction with the massive amounts of information and guide the user towards understanding the ontology's data. Using either a text or graph-based representation, users can discuss and edit the ontology. Social elements utilizing gamification techniques are included to encourage users to create and collaborate on stored knowledge as part of a web community.

An evaluation by 30 users comparing MOI with OpenCyc's original interface showed significant improvements in user understanding of the ontology, although full testing of the interface's social elements lies in the future.

Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation]: User Interfaces – Evaluation/methodology, Natural language, User-centered design

I.2.4 [Knowledge Representation Formalisms and Methods]: Representations (procedural and rule-based)

General Terms

Design, Human Factors.

Keywords

Ontology, web interface, graphic visualization, gamification, crowd-sourcing, OpenCyc.

1. INTRODUCTION

An ontology is a highly structured knowledge base. Such machine-readable representations of concepts are often held up as the 'silver bullet' (Fensel, 2001) for solving problems in information extraction, machine translation, database integration (Suchanek et al, 2008, Medelyan et al, 2009), and the Semantic Web (Berners-Lee 2001, 2002). A great deal of academic work has been done in this research field, producing many ontology languages and tools, and specific ontologies of which the

Permission to make digital or hard copies of all or part 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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '10, Month 1–2, 2010, City, State, Country.

Copyright 2010 ACM 1-58113-000-0/00/0010 ...\$15.00.

'ontology search engine' Swoogle¹ claims to search over 10 000. But examination reveals that many of these are lamentably small and sketchy. There is an increasing realization that in order to get past this proof-of-concept stage, ontology hand-coding is insufficient and some form of crowd-sourcing is required. It is specifically to this task of user-contribution to ontologies that our new interface is addressed.

Ontology interfaces are a remarkably underexplored area in HCI research, with some notable exceptions (e.g. Domingue, 1998). The MOI lays the groundwork for a significant advancement for a number of reasons: i) It displays as much information as possible in plain English, so users can understand and contribute to the ontology without prior experience with formal languages, ii) It graphically visualises ontology concepts and relations, the notoriously rich interlinkages between which are at least part of what makes ontological engineering difficult for many users, iii) It enables users to discuss aspects of the ontology publicly and work collaboratively, keeping logs of all comments and design decisions and finally, it is hoped that, iv) its gamification elements will offer an engaging platform so that users will get involved and contribute.

Design ideas for the new interface were taken from Wikipedia's editing framework, highly successful websites such as Google, and the needs of the Knowledge Miner project of which this interface is a part (Section 2, below). The Bootstrap framework formed the basis of the interface's design, as it provided built-in default styling for HTML elements and simple Javascript functionality which allowed development to focus on the content and how best to display it, rather than the appearance and styling of individual elements.

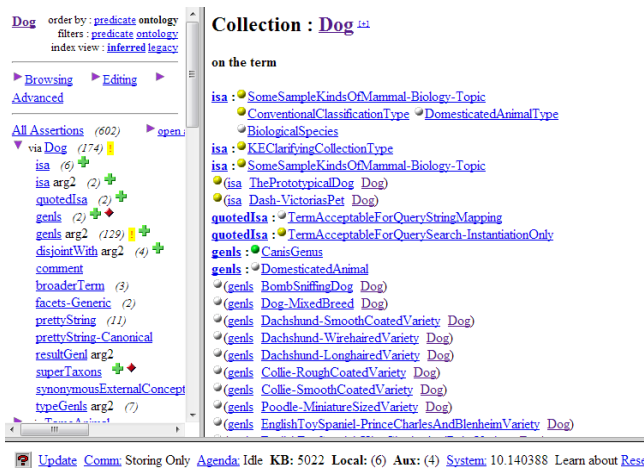
The interface has received preliminary testing via a user survey with 30 subjects performing specified tasks on both it and the OpenCyc interface (Section 5).

2. BACKGROUND

OpenCyc. The OpenCyc ontology is a publicly available spin-off of the large-scale, long-running Cyc project,² whose aim is to codify common sense knowledge for use in strong Artificial Intelligence (AI) applications (Lenat, 1995). A simple example of Cyc's common sense knowledge is, "Bill Gates is not a parking meter," a statement that is obvious to a human, and can be logically proven true by Cyc, though never explicitly asserted there (Legg and Sarjant, 2012). Thus Cyc uses a custom inference engine to answer queries and justify accepting or rejecting assertions with a series of logical deductions from existing

¹ <http://swoogle.umbc.edu/>

² <http://www.cyc.com/platform/overview>



knowledge. The information structures within Cyc closely follow formal logic, consisting of *concepts* which combine to form **Figure 1. An example of the native OpenCyc concept page.**

assertions. Each assertion is a tuple consisting of a predicate followed by a set of arguments in order to make a statement presented as true. Every concept in an assertion represents either a predicate, a collection, or an instance of a collection. For example, (*#\$isa #BillGates #MaleHuman*) states that *#BillGates* is an (*#\$isa*) instance of the *#MaleHuman* collection.

Various drawbacks of Cyc have been discussed in the research literature, such as that its largely in-house manual creation is impractical, and its knowledge overly US-centric (Sowa, 2004). The solution to these problems, as so often in this *Web 2.0* era, would seem to lie in getting the world involved in building and maintaining its ontology, and efforts have been made in this

direction. However, an impediment to this is an arguable further drawback of Cyc which has received much less attention, namely its *user interface*. This interface consists of a HTML-based browser, whose design is rather dated (Figure 1). It uses frames, is almost exclusively text-based, presents concepts and assertions in the CycL formal language, and its documentation is minimal. In the context of the Cyc project these features are understandable since the interface was designed for in-house use by experienced ontology encoders. But if we want to the world to get involved in ontology-building we need to provide something more attractive

Wikipedia is today's single biggest source of free information, and a great example of what can be achieved when the web community collaborates online. Its popularity can be attributed to its support of multiple languages, ease of use, and that the web community actively contributes to it and ensures its accuracy. Mining data from Wikipedia's "goldmine of information" has become very popular (Medelyan, et al., 2009). There are now a number of projects using Wikipedia specifically for automated ontology building, such as DBpedia (Auer et al. 2007, Bizer et al, 2009) and YAGO (Suchanek et al, 2008, Hoffart et al, 2013). Our own effort, Knowledge Miner, for which MOI is the interface, combines Wikipedia and OpenCyc, and will now be described.

Knowledge Miner. This project builds on earlier work with the OpenCyc ontology, which began by mapping 52K Cyc terms to Wikipedia articles with 93% precision (Medelyan, Legg, 2008). Several mining heuristics were then integrated into the mapping, increasing accuracy from 83% to 91%, as well as extending the Cyc ontology by 35,000 concepts mined from Wikipedia (Sarjant et al., 2009). Current work is focused on increasing the quality and scope of the information that can be extracted, mapped, and inserted into the ontology (Legg, Sarjant, 2013). As noted, as part of this work we seek to get the public involved in adding knowledge and correcting its errors, hence the importance of the

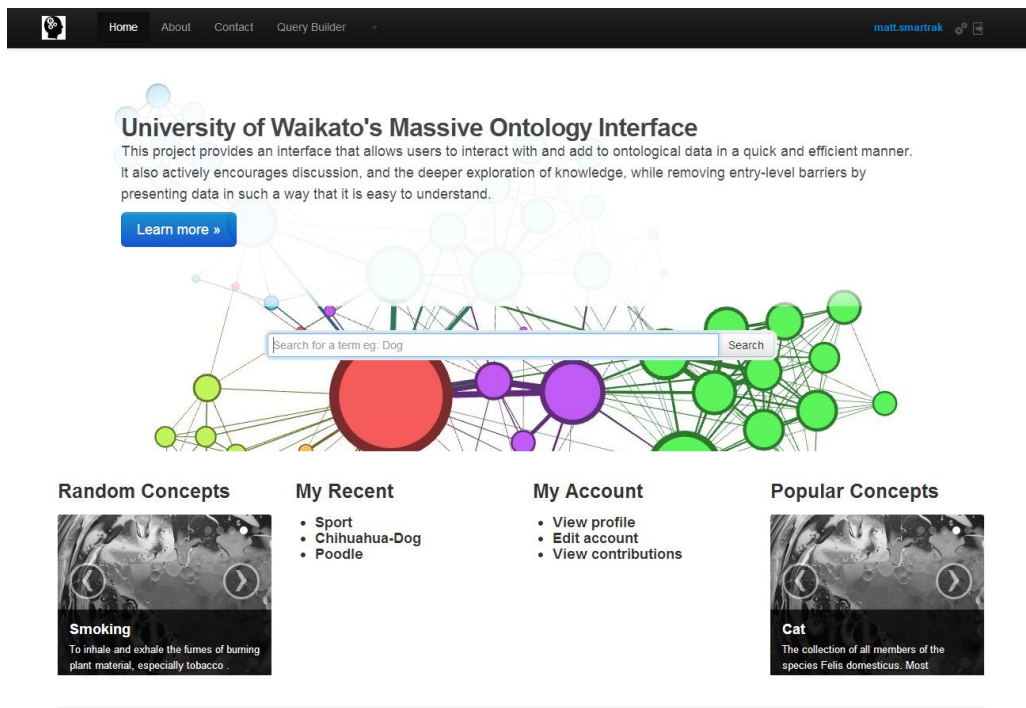


Figure 2. The MOI homepage, containing links to various other key interface elements.

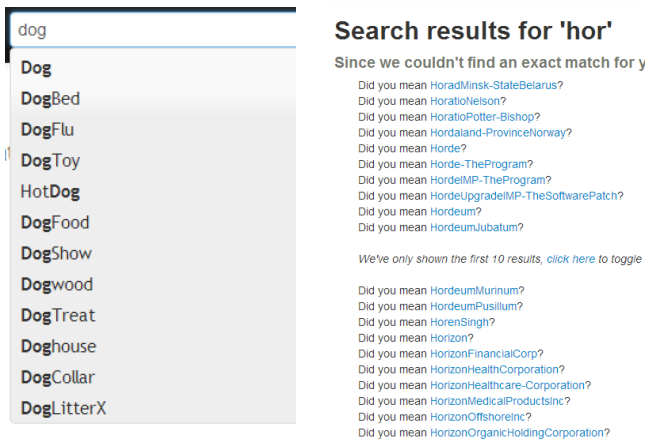


Figure 3. (Left) FAYT Search Bar searching for concepts beginning with "Dog". (Right) Concept listing page.

MOI interface’s usability.

3. KEY INTERFACE ELEMENTS

This section outlines the key structural features of the interface, considered as a website. The user experience will be analysed more specifically in Section 4.

The MOI homepage (Figure 2) is a user’s first entry point to the interface. It contains a link to each key interface element, and a short blurb outlining the website’s intent, with a large button linking to more information. In the page’s centre is the *Find-As-You-Type Search Bar* (Section 3.1), which is also replicated in the top of every other page. This search bar will take the user to the *Concept Page* (Section 3.2) which presents the ontological information for a single concept. Users may also view and

navigate the ontology in a graphical format using the *Visual Taxonomy* (Section 3.3). Users can also navigate to concepts from the homepage using the *Random*, *Recent* and *Popular* links (Section 3.4). The user account details (Section 4) can be accessed either from the *My Account* link at the bottom of the homepage or the links at the top right of the window.

3.1 Find-As-You-Type Search Bar

A Find-As-You-Type (FAYT) Search Bar is a dynamic search control that intermeshes the typing and searching processes together to provide a rapid set of results to the user, allowing them to quickly find concepts in the ontology without needing to type full concept names. When a user has entered three or more characters in a FAYT search bar, the ontology is queried to find a subset of concepts with case-insensitive names that either begin with the search string, or have a synonym beginning with the search string. Every search bar in MOI uses FAYT to streamline searching.

For example, as shown in Figure 3 (left), searching on “dog” produces auto-complete suggestions such as: ‘Dog’, ‘HotDog’, ‘DogPile-SearchEngine’ and ‘Dogma-Movie.’ With this tool, users can easily find ontological concepts by only searching for objects in vague terms and selecting a concept suggestion.

When the search is submitted, the user is either redirected to the *Concept Page* for the named concept, or if the search string did not match any concept, they are redirected to a listing of concepts containing the search string in their name or synonym (Figure 3, right). By contrast, the native OpenCyc interface search functionality does implement a FAYT search, though results are instead ordered alphabetically, which can result in complicated concept names being suggested before shorter-named concepts. When a search is submitted in the OpenCyc interface, the user is either redirected to the concept page, or shown a case-sensitive

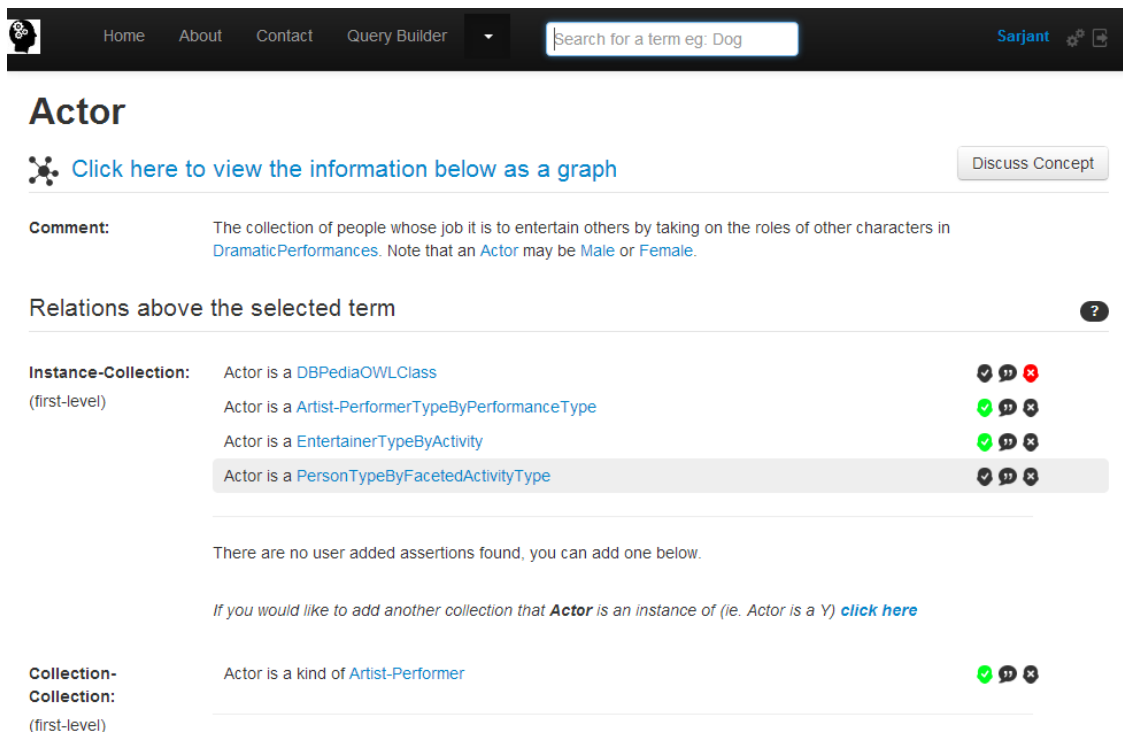


Figure 4. The concept page, listing all assertions about a concept.

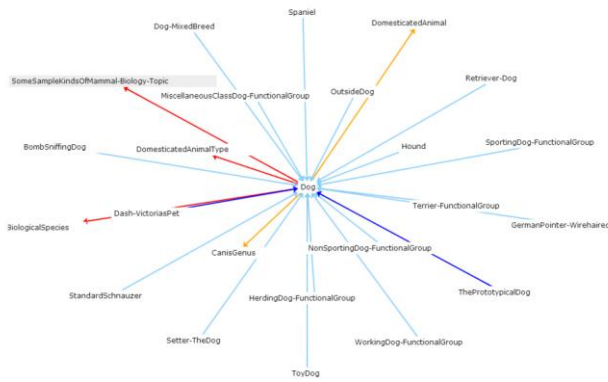


Figure 5. The graphical visualisation of a concept and its upward and downward assertions.

alphabetical listing of results.

3.2 Concept Page

The *Concept Page* (Figure 4) is where all pertinent information about a given ontological concept is displayed. The page is dynamically populated with all assertions that include the concept, laid out vertically down the page including: its comment/description, taxonomically ‘upward’ and ‘downward’ assertions, and then all other assertions. ‘Upward’ and ‘downward’ define broader or more specific concepts respectively, represented in Cyc using the *isa* and *genls* predicates (e.g. (*isa Actor EntertainerTypeByActivity*) and (*genls Actor Artist*) are upward assertions, and (*isa GeorgeClooney Actor*) and (*genls ComedicActor Actor*) are downward assertions). ‘All other’ assertions include assertions using predicates such as *birthDate*, and *numberOfInhabitants*.

To break up the page visually and avoid the cluttered look present in the original interface we used: spacious design, bold text and titles, horizontal rules, icons to replace text, coloured elements and hover-over help text for various elements.

Each assertion provides links to other concepts, such that users can freely navigate between them. To the right side of each assertion are three buttons: *agree*, *discuss*, and *disagree*. Agree and disagree allow a user to directly provide feedback on the ontological data. The aggregated user feedback will be used to alter KnowledgeMiner’s extraction heuristics and remove heavily disagreeable assertions. The discussion button leads to the Discussion Page (Section 4.2), where a user can provide textual feedback regarding the assertion. There is also a discussion page for each *concept*, accessible at the top right of a concept page.

Users may also contribute to the ontology by adding assertions directly. Currently, a user may add additional upwards or downwards relations to concept, using FAYT controls to guide the user in selecting the assertion’s arguments. If the added information is valid (i.e. is ontologically consistent with existing concept information), it is queued for moderation and later incorporated into the ontology. Assertions contributed by users who are moderators or administrators are immediately accepted (see Section 4.3 for more on moderation).

3.3 Visual Taxonomy Page

A user may opt to view a concept’s assertions in visual format by selecting the graph visualisation option on the concept page. On

Kind relations

- [CanisGenus](#)
- ▼ [DomesticatedAnimal](#)
- ▼ [Dog](#)
 - [HerdingDog-FunctionalGroup](#)
 - [NonSportingDog-FunctionalGroup](#)
 - [WorkingDog-FunctionalGroup](#)
 - [MiscellaneousClassDog-FunctionalGroup](#)
 - [Dog-MixedBreed](#)
 - [BombSniffingDog](#)
 - [StandardSchnauzer](#)
 - [GermanPointer-Wirehaired](#)
 - [Setter-TheDog](#)
 - [Terrier-FunctionalGroup](#)
 - [Retriever-Dog](#)
 - [Spaniel](#)
 - [Hound](#)
 - [ToyDog](#)
 - [OutsideDog](#)
 - [SportingDog-FunctionalGroup](#)

the *Visual Taxonomy* page there are two features to the page: a spring graph, and a tree-based listing.

The graph is a force-directed Springy graph (Figure 5) which displays upward and downward assertions, denoted by different coloured lines connecting the central concept to the more/less broad concept. Users can click on each concept to navigate through the graph, proceeding in an upward/downward direction from the central concept. The graphical representation of the concept presents an alternative view that should help users understand the general structure of the ontology.

The second part of the Visual Taxonomy page are two tree-based listings of the upward and downward relations, one listing for *isa* (instance of) assertions, and one for *genls* (kind of) assertions (Figure 6). These listings present a textually hierarchical view of the concept’s upward and downward assertions as an alternative to the Javascript-based graph.

3.4 Popular, Recent and Random Concepts

Figure 6. Tree-based listings of upward and downward *genls* (is a kind of) assertions.

The Popular, Recent, and Random Concept pages, accessible from the homepage or from the top menu bar, all share a similar design and function (Figure 7). On each page there are up to 24 concepts shown, each in its own boxed element. Inside each box element is the concept name (linking to the concept page), comment, and the aggregated number of agree/discuss/disagree interactions carried out on it and its assertions. As on the *Concept Page*, hover-over tooltips explain the interaction icons’ meaning.

These pages present a visual index of different concepts, organized in a variety of ways. The *Popular* page displays the concepts that have the most interactions from the entire user base, so concepts currently attracting the most attention will feature prominently here and entice further interactions. The *Recent* page is the only page wholly influenced by an individual user, providing a quick and easy method for the user to view their recent work, as well as how other users have been interacting with the same concepts. Each user can only view their own *Recent* page but any user can view another user’s interaction data by viewing their profile. The *Random* page displays a random sample of the ontology’s concepts to encourage users to interact with a broader scope of concepts.

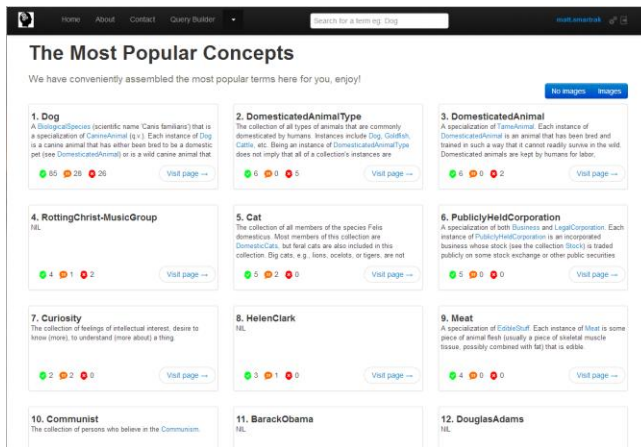


Figure 7. An example of the Popular Concept page.

4. USER EXPERIENCE

This section deals with how user interaction has been incorporated into the MOI. This includes how users access the ontology’s knowledge, how they communicate with it through the interface, how moderation features work, and the extensive database logging of events.

4.1 User Profiles

Each user is assigned a profile, with standard profile actions such as *change email*, *username*, *password*, etc. The View Profile Page (**Error! Reference source not found.**) is for users to view their own profile and that of other users. It is designed to enhance sociality and facilitate interaction between users of the Massive Ontology Interface by allowing each other to see where other users have recently agreed/discussed/disagreed on concepts and assertions. The ten of the user’s most recent interactions are shown, with an option to show more. Each interaction uses colour to distinguish its type, and can be filtered into specific interaction types.

Also included on this page is a summary of a user’s contributions: the number of assertions they have added that have been accepted, and the number of interactions they have carried out – agreed, commented, and disagreed. This data is used to calculate a user’s Action Points (AP); a simple form of gamification designed to entice users to interact with the ontology. Gamification is the “use of video game elements in non-gaming systems to improve user experience and user engagement,” which is a means to encourage and sustain mass-collaboration (Deterding, et al., 2011). We have incorporated gamification elements to gain more social and emotional buy-in from the user base. Each new level of interactivity or way that the user can link the MOI to themselves creates a deeper level of personal investment and a greater level of interaction.

A user’s AP is calculated as:

$$10A + 0.25I$$

where A is the number of accepted added assertions and I is the number of (dis)agreement interactions with assertions. Effectively a contribution is worth 10 AP and an interaction is worth 0.25 AP. The imbalance between contributions and interactions is deliberately calculated to place more importance on users contributing to the ontology than simply interacting with it.

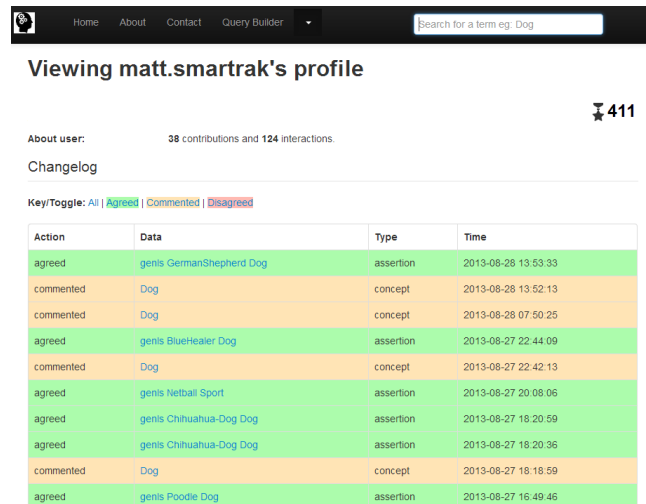


Figure 8. User profile page, listing recent activity and 'Action Points.'

Future development will include other gamification elements such as reward and reputation systems with points, badges, levels and leader boards, which are also known to improve user engagement (Deterding, et al., 2011). AP is based upon the achievement system present in two large online communities that allow user contribution and discussion: Stack Overflow³ and Reddit.⁴ Currently AP is roughly analogous to Stack Overflow’s ‘reputation’ and Reddit’s ‘karma,’ except not influenced by other user’s ratings of contributions.

4.2 Discussion Page

The *Discussion Page*, accessed from the *Concept Page*, provides a framework to foster discussion of an ontological concept or assertion, such as debating the truth of an assertion. Each comment a user submits is added to the discussion and tagged with the user’s username. A comment must contain at least 30 characters. As well as providing a forum-like interface for discussion, the page also includes aggregate totals of user interactions for the information, the author of the information (either user-added, mined, or pre-existing information from OpenCyc), and the internal ontological representation of the data.

User-based moderation functions are also present in the *Discussion Page*. Every user has access to a report function which allows them to flag a comment as offensive in the moderation backend, where a moderator can later make a decision on whether to remove it. In all reporting cases, information about the commenter, reporter, and moderator is logged to prevent future abuse of these functions.

4.3 Administration and Moderation

To prevent exploitation by malicious users or programs, moderation tools are used to protect both the ontological and the social elements of our application. Users exist at one of three levels: ‘normal’ user, moderator, and administrator.

Moderators are normal users that have shown themselves invested in the ontology by the quantity and quality of their contributions and are trusted to safeguard the integrity of the ontology to some

³ <http://meta.stackoverflow.com/help/badges>

⁴ <http://www.reddit.com/wiki/awards>




Figure 9. Responses for Questions 1, 3, and 4.

degree. They have the power to approve or reject a user’s assertions, or send the user a message regarding their addition. Moderators can also adjudicate on whether a comment that has been reported as offensive should be removed. Administrators have the same power as moderators, but may also view and change user levels, and view the log of moderation events.

The area that will likely require the most stringent moderation is the *Discussion Page* where users can freely post comments on assertions and concepts. Here we allow comments to be freely posted without moderation, but incorporate comment reporting functionality.

5. USER STUDY

A user study was conducted to test whether the MOI is better than the native OpenCyc interface for interacting with the ontology. The study was carried out as a digital questionnaire, consisting of 11 questions, 9 of which concerned both interfaces, though 5 of the 9 were ‘general feedback’ questions (text response only). In the other 6 questions, users could provide feedback using a Likert scale and enter text. Mean responses for Likert questions involving both interfaces can be found in Table 1, with confidence values for each mean.

At the beginning users were asked to identify themselves using their gender, age group and a user id randomly generated by the MOI. The user study ran for 24 days and involved 31 participants from 18-60 years old. The level of ontological experience ranged from self-reported novice to expert.

Due to the relatively low numbers of participants, the study focused on the user’s ease in accessing information in the ontology, and submitting feedback to it. We were not able to test such aspects of the social experience as users interacting with each other, and moderating each other’s contributions. We hope that in future the system will achieve sufficient uptake that we will be able to remedy this.

5.1 Questions and Results

Question 1. In ontologies, every "thing" is a "concept." Search for the "Dog" concept using the search box. How easy was it to find the concept?

This question threw off many users as many indicated they did not know what ‘a

	OpenCyc	MOI
Q1	3.74 (0.36)	4.48 (0.31)
Q4	4.06 (0.40)	4.13 (0.42)
Q6	1.77 (0.38)	3.58 (0.44)
Q9	2.55 (0.45)	3.74 (0.36)

Table 1. Mean responses for Questions 1, 4, 6 and 9 with bracketed confidence values.

concept’ represented. Nonetheless, as shown in Figure 9 (left), more users found MOI easier to use for finding the “Dog” concept. However feedback from performing this task in the new interface from multiple users was that they preferred being directed to the “first topic searched” or “most likely result” rather than being shown a list of options and having to choose.

Question 2: Study the layout of the information on the page describing Dog. This is called a Concept page. List what you like and don’t like about how the information is presented.

Feedback about OpenCyc’s native interface expresses that users found the page “quite overwhelming” and disliked that information “isn’t broken into clear sections [...] little help in differentiation or categorisation of the text on the page.” There were also comments regarding the page being “all hyperlinks [...] basically just jargon everywhere.” However, several users did express positive feelings towards the split-frame view where the (hierarchically structured) “related concepts are on the left in a separate part”.

The feedback received from performing the same task in the new interface was much more positive. It included: “I feel like I actually understand what there is to convey,” “Comment up the top is nice,” “information is split into section then it is not information over load,” and “the language used was more accessible and easy to understand.” Once again there were complaints about the presence of jargon within the page, and some users missed the hierarchy tree on the left as it displayed “how concept is derived.” This feedback has been taken into account for future work.

Question 3: As a user, you can give feedback to the system to help improve the quality of assertions (MOI only). Try to agree/disagree with an assertion. How obvious was it that you could do this?

This question only applied to MOI as the functionality is not present in the original interface. Overall 70% of users found that the purpose of the icons was obvious/very obvious (Figure 9, centre). User-provided suggestions included: “simple text links would be better,” “put an ‘Assertions’ title up the top of the columns,” and “a pop up feature for first time users highlighting this ability.” This lack of recognition is a serious problem as the main purpose of MOI is for users to interact with data. The feedback led to the development of a one-time pop-up for users to summarise various aspects of the *Concepts Page*.

Question 4: Try to search for a famous character (actor, musician, superhero etc.) How easy was it to find the page about the character?

For this question, both interfaces had nearly identical results (Figure 9, right). This is probably because the question tests knowledge in the ontology, rather than the interface. As both




Figure 10. Responses for Questions 5, 6, and 9

interfaces feature a search bar in the top, this question, in hindsight was a poor choice for the user study.

Question 5: Concepts can also be visualised as a graph (MOI only). Is the graph less/more clear in communicating a concept’s ontological information than the textual interface?

This is another question that only applied to MOI. The graph view was designed to be supplementary to the textual interface, providing an alternative method of traversing the information within the ontology. The general consensus (20 for, 6 against; Figure 10, left) was that the graph is clearer than the textual interface, but some users did express concerns: “the help text [...] needs further clarification”, “clearer but it does not work if the window is too small,” and that “the textual interface has more information.” It was also conveyed that for such a nice-looking feature it is relatively difficult to find and should be made more prominent within the page. Many users were also not aware of the tree-based listings on the page, so a one-time popup was also added to the visual interface to describe the page’s features.

Question 6: How clear is the meaning of the information on the Assertion page?

Both Likert feedback and comments received for this question showed that the MOI representation of an assertion is much more readable than OpenCyc’s, even for most users with prior ontological experience. For example, one user stated: “looks like a mess” with OpenCyc, then for MOI stated “Clear!!!!”. Many users were confused by OpenCyc’s use of coloured dots for assertion types. However the formal ontological syntax is confusing for some users even in MOI: “Everything makes sense except for ‘isa’ and ‘microtheories’”.

Although the page greatly improved assertion comprehension, approximately 20% of users still had trouble understanding the meaning of the information, which is something for us to improve on in future work (Figure 10, centre).

Question 7: List one thing that the interface did well.

OpenCyc’s native interface was praised for the “quantity of information”, “quick and responsive” nature of the interface, and the “search box at top left easy to find.” Navigation using the tree structure in the left pane was brought up as an advantage over MOI. Users also commented that the interface appeared to be geared towards “expert users to drill down and find the detailed information they were looking for.”

The feedback gathered about MOI mainly related to how user-friendly the interface was: “very approachable for a new user” and “displayed information in a clear and simple way.” Participants also commented on how the search functionality was quicker, as

they were taken to a direct match for a page, rather than a list of matches, as in the OpenCyc interface.

Generally this feedback gave the impression that the interfaces were for two mutually exclusive user groups, experienced and inexperienced, rather than one being an improvement or iteration on another.

Question 8: List one thing that the interface did not do so well.

Feedback on the OpenCyc interface focused on poor layout and design decisions. Participants described the original interface as “looks jumbled up” and “overwhelming in text,” with possible improvements being “to have the [concept] comment at the top.” This reflects the steep learning curve, or prior experience required to effectively navigate and use the original interface.

In regard to MOI, it was once again expressed that “the icons are not immediately apparent as clickable” on the Concept page. There were also suggestions that the interaction icons should be “in colour or a bit bigger or even a key.” The one-time popup for the *Concept Page* attempts to address these issues. Other comments touched on there being “less information than [the OpenCyc interface],” which is true; MOI uses a query to filter out more redundant ontological information. (We are investigating alternative queries in this regard.) Some comments also expressed concern about the location of the *Home Page’s* search bar.

A problem with both interfaces is the restriction of Camel Case (spaces replaced by capitalisation) when searching for concepts. Future work aims to avoid this restriction where possible.

Question 9: How clear is the ontological hierarchy of concepts in this interface?

Although both interfaces are connected to the same ontology, users found that the MOI interface is able to better communicate the ontology’s concept hierarchy than OpenCyc’s (Figure 10, right). This was mainly due to the “Selective use of linking,” and that “Information is grouped into sections” in MOI, rather than the native interface’s “not enough space between the list items. The blue is harsh to read. Poor font, text needs to be bigger.”

Question 10: Each interface accesses the same ontology but visualises the information in a different manner. What would you say is the primary purpose of each interface?

This question was intended to gauge how users compared the two interfaces, in terms of their observed primary purpose. Participant opinions on the purpose of the original interface ranged from “for a machine,” to “an interface for a professional researcher” as the target audience. Other comments comparing the two interfaces said: “[OpenCyc]: Holding Information. [MOI]: Displaying Information,” and “[OpenCyc] is for more experienced users to

get information in a less flashy way whereas [MOI] is for users with less experience.” Suggestions about what the purpose of MOI was for also varied from to provide “a clear understanding for a wider range of audience” at one end, to “encourage exploration,” and a place where “people can have their say.”

It is good that users can generally determine the purpose of MOI, but we would still like to improve MOI such that it is also preferable for professional ontological researchers to use.

Question 11: Please leave any general comments you may have on the two interfaces and their usability.

These comments reiterated general preference for the MOI over the OpenCyc interface for interacting with the ontology. However, there were some final suggestions made about how to further improve the user experience, such as simplifying the *Home Page* design and rearranging the *Concept Page* for concepts that are *Individuals* (instances of collections). We have taken these suggestions into account for future work.

5.2 User Study Analysis

Even though both interfaces communicate with the same ontology, in general users found MOI easier to work with than the OpenCyc interface. Searching for a concept has nearly identical difficulty in each, but once a concept is found, MOI’s simplified design with explanatory text allowed users to better understand the ontological information. However, insofar as MOI reflects the underlying CycL syntax, this still presents a barrier to complete user understanding of the ontology.

The OpenCyc interface was frequently described as an ‘expert’ interface, whereas MOI was described as an interface for new users. While this is a good starting point, we would like MOI to be suitable for both novice and expert users.

The additional features of MOI (visual taxonomy and assertion interaction buttons) were well received, but some aspects require extra work so that users are able to take full advantage of them. The social elements of MOI were not fully tested but may not be testable until a sizable user base is engaged with the interface.

6. CONCLUSION

Ontology interface design is in its infancy. Ironically (given the relative maturity of the two fields) it is more common to try to use ontology as a tool in HCI development than to use HCI as a tool in ontology development. We have made gains in shifting a ‘classic’ ontology interface from expert-user-only to a state such that the general public may contribute to it if they wish. Our results indicate that the interface we developed does encourage users to become involved in contributing to ontological knowledge. With increased user uptake, the quality of the information contained within the ontology will improve, as well as using user feedback as an evaluation mechanism.

7. REFERENCES

- Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z. (2007). DBpedia: A Nucleus for a Web of Open Data. In Aberer, K. et al (eds.) *ISWC/ASWC 2007, LNCS 4825*. Springer, Berlin Heidelberg, 722-735.
- Berners-Lee, T. (2002). Foreword. *Spinning the Semantic Web: Bringing the World Wide Web to its Full Potential*, D. Fensel (ed.). Cambridge: MIT Press.
- Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The Semantic Web. *Scientific American* 284(5): 34-43.
- Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, C., Hellmann, S. (2009). DBpedia – A Crystallization Point for the Web of Data. *Jrnl Web Semantics* 7(3), 154-165.
- Deterding, S. Sicart, M., Nacke, L., O’Hara, K., Dixon, D., (2011). Gamification: Using Game Design Elements in Non-Gaming Contexts. *ACM CHI Conference on Human Factors in Computing Systems, CHI 2011*, May 7–12, 2011, Vancouver, BC, Canada.
- Domingue, J. (1998). Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web. In *Proceedings of the 11th Knowledge Acquisition, Modelling and Management Workshop, KAW’98, Banff, Canada*.
- Fensel, D. (2001). *Ontologies: A Silver Bullet for Knowledge Management and Electronic Management*. Berlin, Springer.
- Hoffart, J., Suchanek, F.M., Berberich, K., Weikum, G. (2013). YAGO2: A spatially and temporally enhanced knowledge base from Wikipedia, *Artificial Intelligence* 194, 28-61.
- Legg, C., Sarjant, S. (2012). Bill Gates is not a Parking Meter: Philosophical Quality Control in Automated Ontology-building, *Symposium on Computational Philosophy, AISB/IACAP World Congress 2012, Birmingham, England*.
- Legg, C., Sarjant, S., (2013). Ontological Quality Control in Large-Scale, Applied Ontology Matching. In *Proc., Workshop on Ontology Matching: OM-2013, ISWC-2013*, Sydney, Australia.
- Lenat, D.B. (1995). Cyc: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM* 38 (11).
- Medelyan, O., and Legg, C. (2008). Integrating Cyc and Wikipedia: Folksonomy meets Rigorously Defined Commonsense. *Proc., Workshop on Wikipedia and AI, AAAI08, Chicago, USA*.
- Medelyan, O., Legg, C., Milne, D., and Witten, I.H. (2009). Mining Meaning from Wikipedia. *International Journal of Human-Computer Studies*, 67(9), 716–754.
- Nielsen, J. (1993). *Usability Engineering*. Boston: Academic Press.
- Sarjant, S., Legg, C., Robinson, M., Medelyan, O. (2009). ‘All you can eat’ Ontology-building: Feeding Wikipedia to Cyc. *Proc Int Conf on Web Intelligence, WI-09, Milan, Italy*.
- Sowa, J. (2004). The Challenge of Knowledge Soup. In *Research Trends in Science, Technology and Mathematics Education*, J. Ramadas and S. Chunawala (2005), 55-90. Mumbai: Homi Bhabha Centre.
- Suchanek, F. M., Kasneci, G., and Weikum, G. (2008). Yago: A Large Ontology from Wikipedia and WordNet. *Elsevier Journal of Web Semantics* 6(3), 203-217.