The Cognitive Atlas: Employing Interaction Design Processes to Facilitate Collaborative Ontology Creation

Eric Miller, Christian Seppa
Squishymedia Inc
524 E Burnside Ste 410
Portland OR 97214
+1 503 488 5951
{eric,christian}@squishymedia.com

Aniket Kittur
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15218
+1 412 268 7505
nkittur@cs.cmu.edu

Fred W Sabb
UCLA
760 Westwood Plaza
Rm C8-854
Los Angeles CA 90095
+1 310 844 6458
fsabb@mednet.ucla.edu

Russell A. Poldrack
University of Texas
3925-B W. Braker Lane
Austin, TX 78704
+1 512 232 9504
poldrack@mail.utexas.edu

Nature Precedings : doi:10.1038/npre.2010.4532.1 : Posted 11 Jun 2010

ABSTRACT
The Cognitive Atlas is a collaborative knowledge-building project that aims to develop an ontology that characterizes the current conceptual framework among researchers in cognitive science and neuroscience. The project objectives from the beginning focused on usability, simplicity, and utility for end users. Support for Semantic Web technologies was also a priority in order to support interoperability with other neuroscience projects and knowledge bases. Current off-the-shelf semantic web or semantic wiki technologies, however, do not often lend themselves to simple user interaction designs for non-technical researchers and practitioners; the abstract nature and complexity of these systems acts as point of friction for user interaction, inhibiting usability and utility. Instead, we take an alternate interaction design approach driven by user centered design processes rather than a base set of semantic technologies. This paper reviews the initial two rounds of design and development of the Cognitive Atlas system, including interactive design decisions and their implementation as guided by current industry practices for the development of complex interactive systems.

Categories and Subject Descriptors
D.2.2 [**Software Engineering**]: Design Tools and Techniques – *Evolutionary prototyping, Software libraries, User interfaces*

H.5.2 [**Information Interfaces and Presentation**]: User Interfaces – *Graphical user interfaces, screen design, user-centered design*

H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces – *Collaborative computing, Web-based interaction*

General Terms
Design, Experimentation, Human Factors.

Keywords
Semantic Web; Ontologies; Collaboration; Cognitive Science; User Interface; User Experience Design.

WWW 2010, April 24-30, 2010, Raleigh, North Carolina.
FWCS 2010, April 26, 2010, Raleigh, North Carolina.

1. INTRODUCTION
Biomedical research has become increasingly reliant on semantic infrastructure and ontologies to guide analysis and enable new discoveries. One challenge to the development of knowledge bases in research domains arises in trying to elicit knowledge from researchers, who are not schooled in the technicalities of ontology development, and who are unlikely to use any system that requires them to learn a complex interaction design. We have developed the Cognitive Atlas project (http://www.cognitiveatlas.org) to address these issues in the domain of cognitive science and neuroscience.

The system aims to develop a cohesive ontology from researchers who are geographically and temporally distributed, and who do not share a common mental model of what the ontology should look like. Thus, in contrast to systems like GO [1] in which the building blocks of the domain are relatively well specified and contributions are carefully managed, Cognitive Atlas needs to enable end users not just to add information but also to participate in the meta-level tasks of discussion, debate, gathering evidence, and building consensus.

In turn, the Cognitive Atlas is designed to contribute to the larger informational ecosystem of neuroinformatics and scientific collaboration projects. The Cognitive Atlas is contributing content to the NeuroLex [2] and the Neuroscience Information Framework [3], and is designed and built to be consistent with the approaches used by related projects such as the Semantic Web Applications in Neuromedicine (SWAN) project [4].

The Cognitive Atlas is somewhat unique in that it is a web application concerned almost entirely with building semantic relationships but is deliberately designed to minimize user interaction with the Semantic Web technologies employed by the system. In part this is because of the nature of the subject matter to be captured: the system needed to support subjectivity, differences of opinion, overlap across domains, and ambiguity.

Ultimately, our approach is designed to address the issue of daunting interactive complexity in semantic web applications by using user-centered design processes to promote the elicitation and discussion of key semantic relations within the domain without requiring users to articulate their contributions using formal ontological structures. Our approach is to create a knowledge base containing enough structured information to support a process where we can refactor the knowledge base for expression as RDF. In effect, the system is designed to meet the users halfway; it asks them to express their knowledge in an interactive format structured around how they understand their domain knowledge, while the underlying system is designed to support the translation of the structured contributions into formally articulated Semantic Web content.

The project's focus on 'soft' design considerations and standard commercial design processes described in this paper may strike a more technically minded reader as somewhat superficial. The authors would suggest that such an interpretation misses a key point regarding the importance of a structured interaction design process. Any description of a design process risks the appearance of triviality. Or to rephrase a commonly quoted aphorism about music criticism,

"Writing about design is like dancing about architecture."

Design is a process with a tendency towards subjectivity. No single methodology, approach, or template will work for every situation. We believe that an elegant concept alone is inadequate to ensure the success of an application; it must also be successfully translated into a workable and useful embodiment that offers value to the user community and is easy to use, and a structured design process offers a way to enhance the likelihood of success.

We believe that the scientific community should take advantage of established interactive design practices and conventions in order to create effective and useful scientific collaboration applications. Otherwise, an unnecessary degree of interactive 'friction' may exist in poorly designed interactions, discouraging use and impairing the usefulness and impact of an otherwise viable scientific project. And an unusable application does not make anyone happy – the principals, the users, or the funders.

2. PROJECT OBJECTIVES

Creating an ontology of mental concepts and their relations with each other has special value for cognitive science for a number of reasons. First, as in many other domains, there are differences in terminology such that the same term could mean different things to different people, and conversely different terms could refer to the same underlying concept. Second, there are fundamental differences in opinion of what certain concepts and processes actually are: for example, some researchers conceive of the concept "working memory" as a system with different components (e.g., the phonological buffer), whereas others think of it from an attentional perspective (e.g., what is currently in the focus of attention). Both of these types of differences have cascading effects to other research that is based on cognition: for example, for a neuroscientist to claim that an area of the brain is associated with working memory, it is necessary to be explicit about what the concept of working memory refers to. Thus an ontology of cognitive processes could provide a foundation for many different areas of research related to cognition, ranging from studies of disease and mental illness to the biology that underlies them, including brain systems, signalling pathways, and even genetic markers.

In addition to the utility of the ontology itself, there are a number of benefits that come about from the process of collaboratively constructing the ontology. In contrast to the cathedral approach to knowledge building, in which a small group of individuals work according to a single plan, the Cognitive Atlas is based on the bazaar approach, in which very many people each contribute based on their own expertise and interests [5] [6]. By combining principles from low-cost distributed knowledge production systems such as Wikipedia with structural elements from more formal ontologies we aim to create a new type of knowledge building approach that harnesses the strengths of both. Some of the goals enabled by this include:

- **Group sensemaking.** One key goal for the Cognitive Atlas is for it to serve as a tool for scientists to make sense of their domain. As described above, this led us to make different design decisions than if the goal was simply the creation of an ontology. For example, coordination and consensus building have been found to be critical to the effectiveness of large scale collaboration systems such as Wikipedia and open source software [7][8]. Thus designing to support discussion and debate was critical to enable a distributed group of scientists to engage in sensemaking together.

- **Individual sensemaking**. At the individual level, we aimed to make the ontology useful for scientists or graduate students unfamiliar with a domain to understand it more easily, for example by having annotated and curated lists of relevant citations for each concept as well as the relations between them. Making the system useful for those actively engaged in using it (as opposed to the potential for being useful sometime in the future for others) has been repeatedly found to be a key factor in the success of groupware systems [9], so this aim is synergistic with the development of the ontology itself.

- **Capturing scientific discussion.** Currently, most forms of scientific discussion take place either extremely slowly (e.g., journal articles) or are limited in their reach (e.g., informal discussions at conferences, or journal clubs). While there have been some efforts to address this by allowing scientists to comment on papers online (e.g., Cell, PLoS ONE), engagement with these venues has been minimal. By providing a central place for these discussions and making them useful in the context of a larger structure (the development of an ontology of cognition) we hope to capture and share scientific discourse in a way that will benefit other researchers interested in a topic.

- **Promoting interdisciplinary research.** By having many different kinds of researchers (e.g., cognitive psychologists, neuroscientists) developing, annotating and creating links between an underlying set of concepts we aim to increase the likelihood of interdisciplinary insight. Furthermore, making it easier for scientists to make sense of an unfamiliar area can promote the chances of their finding important connections to their own areas of expertise.

3. DESIGN RATIONALES

Since the target audience for the Cognitive Atlas is made up of researchers with limited ontology creation experience, we decided to employ familiar and easy to use web application interaction design metaphors to the maximum extent possible.

Initially the project team expected to use the Semantic MediaWiki package as the foundation for the system. After the first couple of months of discussions and research, however, it became clear that the existing open-source packages were not a fit for the needs of this specific project. Ultimately, three key issues steered us towards custom development. We realized that our 'fuzzy' domain-specific semantics and knowledge structures would not easily fit into a pre-existing package; we knew that budget restrictions would limit our ability to adequately customize the existing frameworks to our needs; and we believed the success of our project depended on differentiating our application from other wiki-based knowledge elicitation initiatives by making our application visually unique and memorable. As a result of this decision, the first two releases of Cognitive Atlas have been purpose-built using the open-source LAMP development stack.

By developing a purpose-designed custom framework for the elicitation of semantic knowledge we gained additional flexibility to wrap the user interaction design around the mental models and processes of our audience. Whenever possible we wanted to avoid asking users to shift out of their practitioner-oriented conceptual framework into a technology-driven abstract and formalized framework.

4. DESIGN PRINCIPLES

- **Know the users.** Users of Cognitive Atlas will almost exclusively be researchers in the fields of cognitive psychology, cognitive science, and neuroscience. This is a relatively small and specific audience, so we can assume certain shared characteristics of the user base as a foundation for more sophisticated interaction design. We can use sophisticated domain-specific language to precisely guide interaction, for example, instead of mainstream audience oriented general language.

- **Don't make users think.** This usability principle, popularized by Steve Krug's mainstream interactive design book [10], applies even to complex systems where a sophisticated user base is assumed. Specifically, in the case of applications with a semantic web component, we were aware that many Semantic Web and Semantic Wiki tools would require end users to mentally shift into a kind of 'RDF mindset' in order to contribute to and benefit from the knowledge base. We wanted to avoid forcing users to set aside their existing conceptual frameworks that characterize their understanding of cognitive science. In addition, we assumed that any requirements to learn a new language for specifying knowledge, even something as simple as WikiText, would substantially reduce the level of participation.

- **Minimize interactive friction as much as possible.** Every time a user needs to stop and think about a particular interactive element in order to understand its functionality, a point of 'friction' has been introduced into the interaction. Even an overly heavy single pixel horizontal rule can inhibit communicative effectiveness as a user scans a page [11]. The cumulative effect of these points of friction is to discourage end users since the system requires more in effort from the user than it provides in value.

- **Ensure that the user interaction design is driven by user-centered design goals, and not driven by the technologies and data models.** We want to ensure that end users are focused on the content that they are viewing and contributing, and are not forced to think in terms of the underlying technologies, technology-driven vocabularies, or database structure.

- **Make the best use of all available design techniques.** HTML form elements are only one component of an interactive designer's toolbox. User interactions can be shaped and guided using layout, typography, negative space, relative proportion, alignment, color vocabularies, contrast, sequencing, foregrounding and backgrounding, motion, change over time, and other 'tricks' from the interactive design trade. Collectively these techniques can be orchestrated to provide subtle (or not so subtle) interaction cues and affordances for the user instead of explicitly articulating interaction requirements using the written word.

- **Engage with users throughout the design process.** Industry standard techniques including wireframing, rapid prototyping, and agile development practices help the development team to elicit valuable feedback early and often from end users.

- **Ensure that the site 'sets a hook' in users.** The project will only be a success if we get users to contribute to the system and to return on a regular basis. We need to ensure that every interaction with the system encourages users to return and to continue participating in the construction of the knowledge base. Some planned features include "knowledge management" functions that would provide researchers with added value from their contributions to the system, individual "dashboards" comprised of content relevant to the user, and personalization features inspired by social networking sites. Much of the planned development for 2010 is oriented towards this particular design principle.

- **Avoid overdesign.** We believe it is best to start by designing a simple system capable of supporting most usage scenarios but not necessarily every possible use case. A highly usable system supporting 80% of practitioner-provided information is better than a more complex system that is harder to use but accommodates a wide variety of edge cases.

- **Use industry best practices.** Many commercial design and development best practices are also applicable and valuable when developing noncommercial applications. Structured requirements gathering processes, information architecture and wireframing processes, rapid prototyping, brand and design vocabulary development, agile development techniques, structured project management, and other approaches commonly used by the private sector all apply.

5. IMPLEMENTATION

We used a rapid and iterative design, prototyping and development process to step through and refine our interactive design approaches for the system. We started with flat PDF 'wireframe' style prototypes; developed medium-fidelity HTML prototypes to explore specific design issues; then used rapid release iterations on the alpha system during development to allow quick review and refinement cycles by the project team. This Agile-like process has continued into the production cycle as practitioners begin using the tool and provide feedback to the project managers and development team.

The project team also invested in the development of an identity and design vocabulary system. Developed by a branding professional, the result was a clearly documented design vocabulary that served as a rulebook for a consistent user interface design language employed throughout Cognitive Atlas. The benefits are not merely aesthetic; usability benefits when users familiar with the system internalize the design conventions and are able to quickly understand meaning implied by subtle color and typographical cues.

Terminology was also a key to making interaction intuitive for end users. Cognitive Atlas uses the term "assertions" to describe user contributions that characterize the relationship between two elements in the knowledge base. Typical assertions supported by the system are expressions like:

- Declarative memory is the same as explicit memory
- Memory retrieval is a part of declarative memory
- PMID:12345 (a published research paper) provides empirical evidence that working memory is related to attention
- Working memory is measured by the comparison of 0-back and 1-back conditions on the item recognition task

A relatively comprehensive example of a concept definition with related assertions can be viewed on the current system:
http://cognitiveatlas.org/concept/working_memory
The task of creating assertions is made more challenging when we take into account that definitions for different concepts, tasks, or test results ("indicators") can vary from paper to paper, or from researcher to researcher. So Cognitive Atlas supports the ability to cite existing definitions or to create new definitions for any of the elements cited in an assertion.
The need to express precise semantic relationships still remains; so to facilitate usability and clarity for the end user we designed the user interaction to separate semantic linking activities from other interactions with the knowledge base.

6. CURRENT STATUS

The core functionality of the Cognitive Atlas system is functionally complete and core contributors began building the knowledge base in 2009. Phase II of the system launched in late 2009 and included full support for interoperability via Semantic Web technologies, more robust integration with third-party information sources including PubMed and PubBrain, and many refinements to the interactivity of the system.

The primary user activity as of this writing is centered on contribution and discussion of definitions for the approximately 800 terms in the knowledge base. Once a substantial number of terms have been defined, the assertions (semantic relations) component of the site will be substantially more useful for end users.

We have identified two specific challenges to be addressed in the near future. One challenge is to ensure that the system provides immediate and substantial value to the end user starting with the first interaction. Many sites and systems compete for practitioner mindshare, so we need to make sure that the system is useful from the beginning or we risk losing users and our traction within the community. Secondly, we need to refine our backend Semantic Web tools to facilitate interoperability with other cognitive science initiatives, in particular the NeuroLex [2] project.

Enhancements planned for the 2010 releases of the system include additional personalization, 'dashboard' tools, data mining tools for knowledge discovery based on a user's published abstracts, development of a version for mobile device access to the knowledge base, and visualization tools intended to maximize utility and informational value by avoiding common pitfalls [12].

7. CONCLUSIONS

We would like to be clear that the process described here was suited to this for this particular application and would not be applicable in all situations. The design abstractions we employed as well as the "lite" semantics employed in the application are not globally applicable to Semantic Web-based systems where more formal mechanisms would be more appropriate. However, we do believe that some of the processes, practices, and design techniques described here may be of benefit to the user interface design community and the Semantic Web community, particularly with regard to the elicitation of expert knowledge from a broader audience of practitioners. We also hope that some of the principles employed here may be useful for the design of systems aimed at helping scientists and practitioners to make sense of the state of knowledge in their areas for a diverse range of domains.

8. ACKNOWLEDGMENTS

The development described here was supported by grant RO1MH082795 from the National Institute of Mental Health.

9. REFERENCES

[1] Ashburner M, Ball CA et al (2000) Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. Nat Genet. 2000 May;25(1):25-9.

[2] Bug WJ, Ascoli GA et al (2008) The NIFSTD and BIRNLex vocabularies: building comprehensive ontologies for neuroscience. Neuroinformatics. 2008 Sep;6(3):175-94. Epub 2008 Oct 31.

[3] Gardner D, Akil H et al (2008). The Neuroscience Information Framework: A Data and Knowledge Environment for Neuroscience. *Neuroinformatics*. 2008 Sep;6(3):149-60. Epub 2008 Oct 23.

[4] Ciccarese P, Wu E et al. (2008) The SWAN Biomedical Discourse Ontology**.** J Biomed Inform. 2008 Oct;41(5):739-51. Epub 2008 May 4.

[5] Benkler, Y. (2002) Coase's Penguin, or, Linux and the Nature of the Firm. Yale Law Journal, 112 (3). 367-445.

[6] Raymond, E. (1999). The Cathedral and the Bazaar. Knowledge, Technology, and Policy, 12 (3). 23-49.

[7] Kittur, A., & Kraut, R. E. (2008). Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. CSCW 2008, San Diego, CA.

[8] Mockus, A., Fielding, R. T., & Herbsleb, J. D. (2002). Two case studies of open source software development: Apache and mozilla. ACM Transactions on Software Engineering and Methodology 11 (3), 309-346.

[9] Grudin, J. (1994). Eight challenges for developers. Communications of the ACM, 37, 93-105.

[10] Krug, S. (2005). Don't Make Me Think: a Common Sense Approach to the Web (2nd Edition). New Riders Publishing.

[11] Tufte, E. (1990). Envisioning Information. Graphics Press.

[12] Karger, D. (2006) The Pathetic Fallacy of RDF. SWUI06.