

The National Center for Biomedical Ontology

Mark A. Musen,¹ Natasha F. Noy,¹ Nigam H. Shah,¹ Christopher G. Chute,² Margaret-Anne Storey,³
Barry Smith,⁴ and the NCBO team

¹Center for Biomedical Informatics Research, Stanford University, Stanford, CA USA

²Division of Biomedical Statistics and Informatics, Mayo Clinic, Rochester, MN USA

³Department of Computer Science, University of Victoria, Victoria, BC Canada

⁴Department of Philosophy, University at Buffalo, Buffalo, NY USA

The National Center for Biomedical Ontology (NCBO) is now in its seventh year. The goals of this National Center for Biomedical Computing are to create and maintain a repository of biomedical ontologies and terminologies; to build tools and Web services to enable the use of ontologies and terminologies in clinical and translational research; to educate our trainees and the scientific community broadly about biomedical ontology and ontology-based technology and best practices; and to collaborate with a variety of groups who develop and use ontologies and terminologies in biomedicine. The centerpiece of the NCBO is a Web-based resource known as *BioPortal*. BioPortal makes available for research in computationally useful forms more than 270 of the world's biomedical ontologies and terminologies, and supports a wide range of Web services that enable investigators to use the ontologies to annotate and retrieve data, to generate value sets and special-purpose lexicons, and to perform advanced analytics on a wide range of biomedical data.

MISSION

Advances in computing power and new computational techniques have changed the way researchers approach biology, medicine, and indeed all of science. In biomedicine, one of the most fruitful approaches has been to use software tools and knowledge resources known as *ontologies*—machine-processable descriptions of scientific domains—that can promote the integration of disparate data sources. We have shown that such resources can enable data aggregation, improve search, and allow the detection of new associations that were previously not detectable. It is now possible to demonstrate computationally correlations among genes, diseases, treatments, and outcomes, to use these correlations to efficiently direct research into potentially fruitful areas, and to translate the insights from this research to the practice of medicine. Achieving these integrative analyses requires software systems that take advantage of the semantics of these areas and that can intelligently negotiate domains and knowledge sources, identifying commonality across systems that use different and conflicting vocabularies, while understanding apparent differences that may be concealed by the use of superficially similar terms.¹ An appropriate ontology provides the cornerstone of software for bridging systems, domains, and resources.² Ontologies are the foundation of all semantic technologies in e-science, and are a critical component of multi-disciplinary and translational research in biomedicine.³

The National Center for Biomedical Ontology (NCBO) has become a leading scientific organization for bringing semantic technology to biomedicine. With core performance sites at Stanford University, the Mayo Clinic, the University of Victoria, and the University at Buffalo, our team works to create and disseminate national infrastructure that supports the use of computer-stored knowledge in the form of ontologies.

Our overall mission comprises four main objectives:

1. We create and maintain a repository of biomedical ontologies and terminologies.
2. We build tools and Web services to enable the use of ontologies and terminologies.
3. We educate our trainees and the scientific community broadly about biomedical ontology and about NCBO technology.
4. We collaborate with a variety of groups who develop and use ontologies and terminologies in biomedicine.

OUTPUTS OF THE CENTER

The outputs of our center can be best described in terms of the overall objectives of our work.

Repository of biomedical ontologies

The NCBO's BioPortal provides access to more than 270 biomedical ontologies and controlled terminologies.^{4,5} Users come to the BioPortal Web site to browse biomedical ontologies and to search for specific ontologies that have terms that are relevant for their work. A cancer biologist may learn from BioPortal that the Gene Ontology offers the best coverage for annotating her experimental data with terms related to cell division, or that she can access more precise terms in the National Cancer Institute (NCI) Thesaurus. She may discover that the Mouse Adult Gross Anatomy Ontology can be used to describe the body parts from which her experimental specimens were obtained, or that the National Drug File–Reference Terminology (NDFRT) provides valuable information about the properties of the drugs used in her experiments.

BioPortal enables users to navigate ontologies using a standard tree browser. Users also can visualize resources in BioPortal using special tools that offer cognitive support for understanding the complexities of large ontologies (*Figure 1*).

When users need to understand the relationships between terms in two different ontologies, BioPortal provides mappings between the ontologies to enable direct comparisons. The mappings can inform the user that the term *lung* in the Mouse Adult Gross Anatomy Ontology is related to the term *lung* in the Foundational Model of (human) Anatomy or that the term *limb* in the NCI Thesaurus is related to the term *extremity* in the Mouse Adult Gross Anatomy Ontology. The mappings between ontologies in BioPortal not only allow users to compare the use of related terms in different ontologies, but also allow analysis of how whole ontologies compare with one another. They allow us to identify ontologies that cluster together⁶ and to identify the degrees of overlap among ontologies.⁷ Like the UMLS metathesaurus,⁸ the mappings in BioPortal facilitate automated translation of terms among ontologies, but entail much more content. The mappings in BioPortal form the basis for what we refer to as the NCBO *mega*-thesaurus.

BioPortal is much more than an ontology repository, however. We have created the system as the nexus of an online community of ontology developers and ontology users who use BioPortal to view, to comment on, and to discuss the content of biomedical ontologies (*Figure 2*). Registered users of BioPortal not only can upload new ontology content, but also mark up their content (or that of any other user) with highly granular comments about any ontology.⁹ Users can indicate where they believe ontologies may reflect inappropriate modeling decisions, and other users can respond to those comments in threaded discussions that the entire BioPortal community can monitor. These threaded conversations allow BioPortal to behave very much like a wiki for making annotations to ontology content, and they enable new users to locate regions of BioPortal's ontologies where modeling decisions have been particularly controversial and ontology developers to identify elements of their work that may benefit from refactoring in future versions of their ontologies. They also allow users to identify those groups of resources, such as are maintained by the Open Biological and Biomedical Ontologies (OBO) Foundry initiative,¹⁰ that have been subjected to a process of external review designed to ensure compliance with an evolving set of best practice principles.

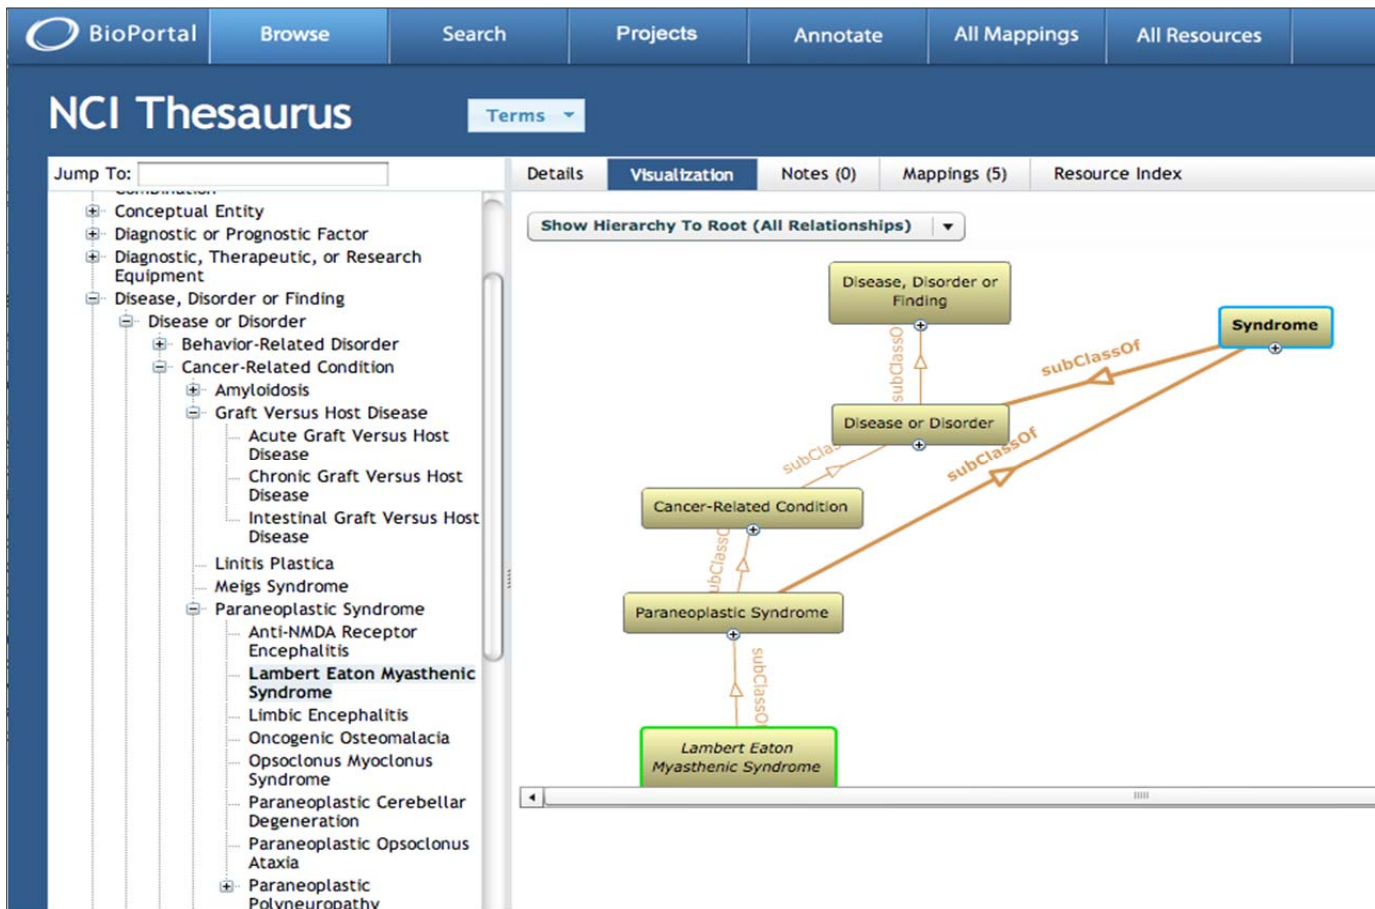


FIGURE 1: The BioPortal ontology repository. In the Figure, the user is browsing the NCI Thesaurus. A tree browser along the left-hand side of the screen allows the user to navigate the taxonomic hierarchy of the ontology. The visualization window on the right facilitates exploration of complex relationships—here, the pathway between the selected term (Lambert Eaton Myesthenic Syndrome) and its superclasses in the hierarchy. The menu bar above the visualization window allows the user to change the view to examine the details of the selected term, end-user notes regarding the term or its descendants, mappings between the term and related terms in other ontologies, or links between the selected term and the data sources referenced in the NCBO Resource Index.

BioPortal allows users themselves to post overarching reviews of the system’s ontologies—and to post online very specific proposals for changes that ontology developers might want to consider in future revisions. BioPortal thus adopts Web 2.0 conventions to allow its users to communicate with one another about the NCBO’s hosted ontologies in a highly interactive manner. The outcome of these capabilities is that BioPortal offers the equivalent of online, open, community-based peer review for the BioPortal ontology content.⁹

We are developing BioPortal so that computer-based ontology-development tools can access all its content programmatically—including the mappings between ontology terms and the notes about the ontology content contributed by members of the user community. Thus, users of the Web-based version of the Protégé ontology editor¹¹ can view BioPortal content directly from within the Protégé browser window, copy terms and other content from existing ontologies into new ontologies, review the notes and comments about previous versions of ontologies uploaded to BioPortal by their users, and act on those notes as they develop new versions. This integration of ontology authoring with community-based access to ontologies through BioPortal has been particularly important to groups developing large ontologies in an open, distributed fashion. For example, the World Health Organization is now using NCBO technology routinely in its global effort to develop the next edition of the International Classification of Diseases (ICD-11).¹²

The screenshot shows the BioPortal interface. On the left is a hierarchical tree of ontologies, with 'osseous' selected under 'disc composition'. The main area displays a comment thread for the 'osseous' term. The comment is titled 'does not fit to taxonomy' and was submitted by 'daumkep' on 03/08/2011. The comment text reads: 'Hi all, to my view, "osseus" does not really fit as a child of "disc composition" as it can appear in many other contexts. A good terminology practice would be to name the preferred terms in a unique way, e.g. "osseus disc composition" and put the term "osseous" on a place such as Qualifier "value->Appearance->osseous". What do you think? All the best Philipp'. Below the comment is a 'Responses' section with one reply from 'beverlycollins' dated 5 months ago, which says: 'Thanks, Philipp, we agree that the terms in this category more correctly belong in the "composition" class. We will move them there and move the class "disc composition" to Obsolete Radlex terms, since this term does not seem necessary. Beverly'. There is an 'Add Reply' button at the bottom of the response area.

FIGURE 2: Notes in BioPortal. Registered users of BioPortal can comment on any of the ontologies in the repository. They can point out what they believe to be errors or can make suggestions for changes. Other users can respond to these comments and begin a threaded discussion. In the Figure, a user has left a note in the RadLex Ontology suggesting that the term *osseous* may be misclassified. Another user has left a note agreeing that the term needs to be relocated in a future version of RadLex.

Tools and Web services

In addition to providing a comprehensive library of biomedical ontologies and terminologies, the NCBO develops tools and services that use those ontologies to aid biomedical investigators in their work. Although these tools are all available through a Web-browser interface, most users access our software programmatically via Web services.

NCBO Annotator. Perhaps the most widely used tool created by the NCBO is one that maps arbitrary keywords and natural-language text to standardized ontological terms. The NCBO Annotator thus takes as input some specified text and generates as output a set of terms derived from BioPortal-stored ontologies, such that the terms refer to concepts that the NCBO Annotator identifies in the text.¹³ It provides a mechanism to determine what the text is “about” in terms of standardized, ontological entities. The structure of the ontologies in BioPortal permits the NCBO Annotator to associate the text not only with particular terms (e.g., *adenocarcinoma of the lung* from the NCI Thesaurus), but also with more general terms (e.g., *neoplasm*). As a result, users are offered an extremely rich set of descriptors for the corresponding text, at different levels of granularity and generality.

Resource Index Search [View Demo!](#)

Rash (preferred name) from: MedDRA x [Search](#)

To **begin**, type in **text** and select a **matching term**.
 Examples: **melanoma**, **lupus**, **breast cancer**, ...
 Click **search** and **select a repository**. [Clear](#)

Ontology filters

18	ARRS GoldMiner	64585	Adverse Event Reporting System Data
15	ArrayExpress	0	Biositemaps
834	ClinicalTrials.gov	7	Conserved Domain Database (CDD)
0	Database of Genotypes and Phenotypes	14	DrugBank
20	Gene Expression Omnibus DataSets	1	MICAD
28	Online Mendelian Inheritance in Man	0	Pathway Commons
8	PharmGKB [Disease]	4	PharmGKB [Drug]
7	PharmGKB [Gene]	0	PubChem
1750	PubMed	0	Reactome
265	ResearchCrossroads	36	Stanford Microarray Database
0	UniProt KB	0	WikiPathways
0	caNanoLab		

FIGURE 3: The user interface for the NCBO Resource Index. The Resource Index is a database that links each term of each ontology in BioPortal to online data and knowledge resources that may reference that term. In the Figure, the user has entered a particularly vague term—*rash*, as used in MedDRA. The system uses the Resource Index and the underlying ontological structure in which the term appears to allow the user to locate some 18 images in the American Roentgen Ray Society’s GoldMiner repository of radiographs; 20 microarray datasets in the Gene Expression Omnibus; 28 records in Online Mendelian Inheritance in Man; and so on. Each of the associated data sets refers to patients with some kind of rash. Clicking on a particular resource description in the user interface allows the user to navigate to the actual data records that have been indexed.

NCBO Resource Index. A common use of the NCBO Annotator is to ascribe ontological terms to the textual metadata that are associated with experimental data sets. The NCBO automatically runs the Annotator on a large collection of online data sets, linking the textual metadata associated with those data to all relevant ontological terms in BioPortal. The result is an enormous database of all the terms (and abstractions of those terms) that relate to the textual metadata (or text descriptions) found in a growing set of online data resources (such as the microarray data sets in the Gene Expression Omnibus or the individual protein descriptions in UniProt). We refer to this database as the NCBO Resource Index.¹⁴ A Web-based interface allows investigators to search the Index, using terms in BioPortal-stored ontologies to locate relevant data sets from online repositories (*Figure 3*). The Index offers the biomedical community a common interface for information retrieval, linking the dozens of ontologies in BioPortal to dozens of biomedical data resources. Thus, if an investigator is interested in learning what experimental data might have been archived online in public repositories that might be relevant to a particular term or set of terms, she can use the Index to search for the relevant data. The NCBO development team is linking new online data resources to the Resource Index on an ongoing basis.

NCBO Ontology Recommender Service. Investigators often are unsure which of the dozens of ontologies in BioPortal provide the best coverage for capturing the entities in a particular application area. The NCBO Ontology Recommender Service¹⁵ takes as input representative textual data relevant to a domain of interest and returns as output an ordered list of ontologies available in BioPortal the terms of which would be most appropriate for annotating the corresponding text.

NCBO Lexicon Builder. Users frequently turn to the terms of biomedical ontologies to create the *value sets* that constitute the basis of “pick lists” that allow users to make selections from menus when filling in computer-based forms. The NCBO Lexicon Builder¹⁶ also allows users to obtain more manageable subsets of large ontologies that are amenable to particular analyses, and to combine portions of different ontologies to create specialized collections of terms. The latter functionality is of particular interest to members of the natural-language processing community, who often need hand-crafted lexicons to drive named-entity recognition in particular domains.

Web widgets. Many NCBO services are called automatically through small collections of HTML program code that our Center makes available to Web developers who wish to take advantage of our offerings. Developers can embed these “widgets” in their code so that their Web pages can immediately access BioPortal ontologies, value sets, mappings, and other resources.

Detailed information for developers who wish to access and employ NCBO tools, services, and widgets is available on a wiki maintained by the Center.¹⁷

Education and outreach

A full-time NCBO outreach coordinator has multifaceted responsibilities that include serving as a liaison to collaborating projects, shepherding new collaborations, and presenting NCBO technology to the scientific community. Our outreach coordinator hosts a very well attended, biweekly “Webinar” series, in which members of the NCBO and the larger biomedical ontology community discuss their research; video recordings of past Webinars are archived on the NCBO Web site.¹⁸

The NCBO has an active dissemination program of custom-tailored workshops and tutorials. The Center is also a major sponsor of the International Conference on Biomedical Ontology (ICBO).

COLLABORATIVE PROJECTS

The NCBO’s tools and services are designed for use in support of the informatics activities of biomedical researchers. As a result, our collaborators tend to be well versed in biomedical informatics and understand the power that ontologies can offer their work in data annotation and indexing, natural language processing, data mining, and decision support.

The NCBO has supported a series of Driving Biological Projects that have provided important use cases for the NCBO Annotator (e.g., in the annotation of rat genome data¹⁹) and for the NCBO Resource Index (e.g., to enable retrieval of information about therapeutic nanoparticles²⁰). The use of ontology-driven analytics has allowed collaborators to interpret high-throughput data in novel ways,^{21,22} making both methodological and biological contributions. Other collaborators have used the rich content availability in BioPortal as a starting point for quality assurance of ontologies²³ and for further enrichment of biomedical ontologies by processing text from electronic health records.²⁴ Finally, our collaborations have led to the development of a burgeoning number of new ontologies for use by the biomedical community.^{10,12,25,26}

The vast majority of biomedical investigators who take advantage of NCBO technology are not explicit collaborators, however. Most of the Center’s users simply browse the BioPortal Web site or invoke NCBO Web services as a routine element of their investigative work. Currently, some

16,000 users browse ontologies via BioPortal each month. During the same period of time, NCBO servers respond to more than 3-million programmatic Web-service requests. It is extremely gratifying to the members of the Center that the NCBO apparently has become an indispensable technology resource for such a large community of biomedical scientists, and that the vast majority of these users have come to take the Center's services for granted.

FUTURE GOALS

A major initiative of the NCBO in the coming years will concentrate on ensuring the scalability of our technology. As BioPortal acquires increasing numbers of ontologies (with increasing numbers of inter-ontology mappings and end-user notes), as more and more biomedical data resources are linked to the NCBO Resource Index, and as the number of users who access our technology via Web browsers or via Web services continues to grow, the NCBO must be able to accommodate the corresponding demand. Much of the Center's activity concerns ensuring a robust infrastructure for its technology, accommodating more content, more users, and more demands in as seamless a manner as possible.

The scientific work of the Center will focus on support for the management of the complete ontology life cycle, allowing users of ontology-development systems (such as Protégé,¹¹ OBO-Edit,²⁷ and LexWiki²⁸) to integrate with BioPortal. The authors of ontologies will be able to publish their work directly in BioPortal, and to take advantage of end-user notes when revising their ontologies in subsequent versions.²⁹ We will merge the processes of ontology authoring and ontology dissemination, and investigate whether the open peer-review process offered by BioPortal can lead to improvements in biomedical ontologies and to enhanced adoption of the resultant ontologies by the community.

Other work will concentrate on new uses of the ontologies in BioPortal in the interpretation of high-throughput experimental data.³⁰ We also are optimistic that the ontology-oriented techniques that we are developing will enable investigators to analyze data from electronic patient records in novel ways.³¹

More information about the NCBO is available from the Center's Web site.³²

REFERENCES

1. Shah, N.H., and Musen, M.A. Ontologies for formal representation of biological systems. In: Studer, S., and Staab, S. (eds) *Handbook on Ontologies*. New York:Springer-Verlag, 2009, pp. 445–462.
2. Uschold, M., and Gruninger, M. Ontologies: principles, methods, and applications. *Knowledge Engineering Review* **11**(2):93–136, 1996.
3. Lussier, Y.A., and Bodenreider, O. Clinical ontologies for discovery applications. In: Baker, C.J.O., and Cheung, K.-H. (eds.) *Semantic Web: Revolutionizing Knowledge Discovery in the Life Sciences*. New York:Springer-Verlag, pp. 101–119.
4. Noy, N.F., Shah, N.H., Whetzel, P.L., et al. BioPortal: Ontologies and integrated data resources at the click of a mouse. *Nucleic Acids Res* **37**(Database issue):W170–3, 2009.
5. Whetzel, P.L., Noy, N.F., Shah, N.H., et al. BioPortal: enhanced functionality via new Web services from the National Center for Biomedical Ontology to access and use ontologies in software applications. *Nucleic Acids Res* **39**(Web Server issue):W541–5, 2011.
6. Ghazvinian, A., Noy, N.F., Jonquet, C., et al. What four million mappings can tell you about two hundred ontologies. In: *Proceedings of the Eighth International Semantic Web Conference (ISWC 2009)*, Washington, DC, September, 2009, New York:Springer-Verlag.

7. Ghazvinian, A., Noy, N. F. and Musen, M. A. How orthogonal are the OBO Foundry ontologies? *J Biomed Semantics*, in press.
8. Bodenreider, O. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Res* **32**(Database issue):D267–70, 2004.
9. Noy, N.F., Dorf, M.V., Griffith, N., et al. Harnessing the power of the community in a library of biomedical ontologies. In: *Proceedings of the Workshop on Semantic Web Applications in Scientific Discourse*. Eighth International Semantic Web Conference (ISWC-2009), Chantilly, VA, October, 2009.
10. Smith, B., Ashburner, M., Rosse, C., et al. The OBO Foundry: Coordinated evolution of ontologies to support biomedical data integration. *Nature Biotech* **25**(11):1251-1255, 2007.
11. Tudorache, T., Vendetti, J., Noy, N.F. WebProtégé: A lightweight OWL ontology editor for the Web. In: *Proceedings of the Fifth OWLED Workshop on OWL: Experiences and Directions*. Seventh International Semantic Web Conference (ISWC-2008), Karlsruhe, Germany, October, 2008.
12. Tudorache, T., Falconer, S.M., Nyulas, C.I., Storey, M.-A., Üstün, T.B., Musen, M.A. Supporting the Collaborative Authoring of ICD-11 with WebProtégé. In: *AMIA Annu Symp Proc*, Nov 2010:802–806.
13. Shah, N.H., Nipun, B., Jonquet, C., et al. Comparison of concept recognizers for building the Open Biomedical Annotator. *BMC Bioinformatics* **10**(Suppl 9):S14, 2009.
14. Shah, N.H., Jonquet, C., Chiang, A.P., et al. Ontology-driven indexing of public datasets for translational bioinformatics. *BMC Bioinformatics* **10**:Suppl 2:S1, Feb 5 2009.
15. Jonquet C., Musen M.A., Shah, N.H. Building a biomedical ontology recommender web service. *J Biomed Semantics* **10**:Suppl 1:S1, 2010.
16. Parai, G. K., Jonquet, C. M., Xu, R., et al. The Lexicon builder Web service: Building custom lexicons from two hundred biomedical ontologies. In: *AMIA Annu Symp Proc*, 2010 Nov 14:587–591.
17. http://www.bioontology.org/wiki/index.php/Main_Page [accessed July 30, 2011]
18. <http://www.bioontology.org/webinar-series> [accessed July 30, 2011]
19. <http://qtlhighlighter.hmgc.mcw.edu/> [accessed July 30, 2011]
20. Thomas D.G., Pappu R.V., Baker N.A. NanoParticle Ontology for cancer nanotechnology research. *J Biomed Informatics* **44**(1):59–74, 2011.
21. Mort, M., Evani, U.S., Krishnan, V.G., et al. In silico functional profiling of human disease-associated and polymorphic amino acid substitutions. *Hum Mutat* **31**(3):335–46, 2010.
22. Sarkar I.N. Leveraging biomedical ontologies and annotation services to organize microbiome data from Mammalian hosts. In: *AMIA Annu Symp Proc*. Nov 2010:717–21.
23. Verspoor, K., Dvorkin, D., Cohen, K.B., and Hunter, L. Ontology quality assurance through analysis of term transformations. *Bioinformatics* **25**(12):i77–84, 2009.
24. Liu, K., Hogan, W.R., and Crowley, R.S. Natural language processing methods and systems for biomedical ontology learning. *J Biomed Inform* **44**(1):163–79, 2011.
25. Winslow, R. L., Saltz, J., Foster, I., et al. The CardioVascular Research Grid (CVRG) project. In: *AMIA Annu Symp Proc*. Nov 2010:77–81.

26. Tenenbaum, J., Whetzel, P. L., Anderson, K., et al. The Biomedical Resource Ontology (BRO) to enable resource discovery in clinical and translational research. *J Biomed Informatics* **44**(11):137–45, 2011.
27. Day-Richter J., Harris, M.A., Haendel, M., et al. OBO-Edit—an ontology editor for biologists. *Bioinformatics* **23**(16):2198–200, 2007.
28. Jiang G., and Solbrig H. LexWiki framework and use cases. In: *The First Meeting of Semantic MediaWiki Users*. Boston, MA. Nov 22–23, 2008.
29. Noy, N.F., Tudorache, T., Nyulas, C.I. and Musen, M.A. The ontology life cycle: Integrated tools for editing, publishing, peer review, and evolution of ontologies. In: *AMIA Annu Symp Proc*, Nov 2010:552–556.
30. Tirrell, R., Evani, U., Berman, A. E., et al. An ontology-neutral framework for enrichment analysis. In: *AMIA Annu Symp Proc*, Nov 2010:797–801.
31. LePendou, P., Racunas, S., Iyer, S., et al. Annotation analysis for testing drug safety signals. The Fourteenth Annual Bio-Ontologies Meeting. Conference on Intelligent Systems for Molecular Biology (ISMB-2011). Vienna, AT, July 2011. (<http://www.bio-ontologies.org.uk/programme>; [accessed July 30, 2011])
32. <http://bioontology.org> [accessed July 30, 2011]

Funding: The National Center for Biomedical Ontology is supported by the NIH Common Fund, the National Human Genome Research Institute, and the National Heart, Lung, and Blood Institute through grant U54 HG004028.

Competing Interests: None