An Ontology-Driven Knowledge Environment For Subcellular Neuroanatomy

Lisa L. Fong[1], Stephen D. Larson[1], Amarnath Gupta[2], Christopher Condit[2], William J. Bug[3], Li Chen[2], Ruth West[1], Stephan Lamont[1], Masako Terada[1], and Maryann E. Martone[1]

lfong@ncmir.ucsd.edu, slarson@ncmir.ucsd.edu, gupta@sdsc.edu, condit@sdsc.edu, william.bug@drexelmed.edu, lichen@sdsc.edu, spl@ncmir.ucsd.edu, mterada@ncmir.ucsd.edu, mmartone@ucsd.edu

[1] National Center for Microscopy and Imaging Research,
Center for Research in Biological Structure and Development of Neurosciences,
University of California, San Diego, La Jolla, CA 92093-0608, USA
[2] San Diego Supercomputer Center,
University of California, San Diego, La Jolla, CA 92093-0505, USA
[3] Laboratory for Bioimaging and Anatomical Informatics,
Department of Neurobiology and Anatomy
Drexel University College of Medicine, Philadelphia, PA 19129, USA

Abstract. The nervous system, perhaps more so than any other tissue, requires the fine spatial resolution of high resolution microscopy to resolve the intricate structure of individual nerve cells and supracellular domains. The analysis of such imaging data often involves the extraction of cellular surfaces and cell components. The resulting geometries are used to derive measurements or, when combined with additional information, as the basis of a computational model. In an effort to streamline the process of model creation, we are creating a knowledge environment where researchers can access well-annotated geometrical models and a knowledge base of neuronal structure. Underlying this system is a formal ontology describing the subcellular anatomy of the nervous system (SAO), covering nerve cells, their parts and interactions between these parts. In this report, we describe the structure and content of the SAO along with SAO-based tools for annotation and query of electron microscopic data.

Keywords: Protege, nervous system, database, reasoning, descriptive logic, rules-based logic, data integration, upper-level ontology

1 INTRODUCTION

One of the most remarkable properties of the nervous system is its capacity for plasticity across multiple temporal, structural and biochemical scales. The ability of the nervous system to change its structure in response to environment changes is believed to provide the substrate for the brain's ability to store and process information. Studies of nervous system plasticity focus on processes and structures existing in what has come to be called the "mesoscale." We define mesoscale as those structures situated between more gross anatomical scales and the level of individual protein and gene structure. It comprises the dimensional range of microns (μm) down to nanometers (nm), encompassing cellular networks, subcellular microdomains, and their macromolecular constituents.

To elucidate the complex structural and dynamic inter-relations critical to nervous system function, entities at this scale must be observed, described, and manipulated.

Numerous studies have documented the exquisite structural and molecular heterogeneity of nervous tissue, where each cell domain is characterized by precise distributions of macromolecules. With the completion of the Allen Brain Atlas (Lein et al., 2007), a fairly complete representation of gene product distribution in the mouse brain exists at the level of the cell. Due to extensive post-translational processing and transport, however, RNA expression patterns do little to specify the subcellular location of the protein product. The sheer size and complexity of nerve cells and their microdomains renders large scale specification of protein distributions and their functional consequences challenging. The challenge arises because the nerve cell dimensions require a combination of technologies to resolve. Investigations of cellular physiological and structural dynamics, overall cellular morphology and coarse molecular distributions are accomplished by optical microscopies (LM). Resolving fine details of internal structure, cytoskeletal organization, precise localization of molecular constituents, location of synaptic contacts, and detailed views of the synaptic and cellular microdomains requires 3D electron microscopic (EM) imaging.

To discern the functional consequences of structural and molecular variation in the nervous system one must combine nerve cell descriptions derived from multiple techniques. Neuroscientists now have powerful simulation environments incorporating experimentally-derived properties into realistic structural and physiological models of cells and cell ensembles. Simulations bridge the structural and the molecular domains to indicate how higher order structure on the scale of cells and tissues can influence molecular-scale dynamics, as well as how alteration of molecular constituents affects the emergent properties of specific types of nerve cells and circuits.

Computational modeling relies on high resolution 3D models of cellular structures to provide biological realism and to constrain possible solutions (Hunter and Borg, 2003). Of current technologies, electron tomography, the derivation of 3D structure from a series of 2D electron micrographs, holds the most promise for providing such models (Figs. 1-2; Lucic and Baumeister, 2005). Creation of the geometry currently involves segmentation of cellular surfaces on which the distributions of macromolecules are "painted". As simulation packages have no accepted standard for capturing complex distributions of mesoscopic elements, much of the effort invested when constructing models is difficult to re-use and extend. This lack of a shared, formal semantic description also weakens the link between the models and the experimental evidence from which they derive their validity, especially when disparate models are considered collectively.

In this paper, we describe the SAO, an ontology used to represent the subcellular anatomy of the nervous system (`http://ccdb.ucsd.edu/SAO/1.0/SAO.owl`). The SAO emerged from our efforts to design databases and tools for describing mesoscale structures and to provide the necessary infrastructure to link these sub-cellular elements with those from molecular and gross anatomical scales. We illustrate a suite of programs built around the SAO to aid in annotating data from EM imaging and the creating spatially realistic models of mesoscale structures. Finally we describe design and implementation issues encountered along the way during creation and application of the SAO.

2 METHODS AND RESULTS

The SAO was built using Protege version 3.2.x in OWL (Web Ontology Language) 1.0. In most parts, the SAO conforms to OWL-DL, with DL expressiveness of SHIN(D). The core content of the SAO is based largely on Peters, Palay, and Webster, The Fine Structure of the Nervous System,

Ed. 2 (1991), the standard reference for neuronal ultrastructure, with additions from more recent literature. The SAO foundation is the Basic Formal Ontology (BFO; Grenon et al., 2004; Rosse et al., 2005) which divides entities into continuants (e.g. objects, qualities, sites, etc.) and occurrents (processes). The SAO elaborates the regional parts of cells and their associated cell components constituting the normal adult nervous system. The aggregation of these cell parts into supracellular structures such as the Node of Ranvier is also supported. Examples of image data containing these structures can be found at `http://ccdb.ucsd.edu`.

Cell Descriptions: The SAO contains a list of cell types found in the nervous system. The SAO does not contain a comprehensive list of neuron types, because these entities fall under the scope of other ontologies, e.g., Cell Type Ontology (Bard et al., 2005). Rather, because the SAO is designed as an application ontology for annotation of biological data, the parent cell types are expected to be added to the SAO as they are encountered. The SAO lists neurons according to common names reflecting a mixture of classification criteria, e.g., morphology ("pyramidal neuron"), proper names ("Purkinje neuron"). The SAO utilizes these names merely as labels and does not further classify cell types into subtrees, except in instances where the hierarchy is fairly straightforward, e.g., layer 3 cortical pyramidal neuron is a cortical pyramidal neuron. We deliberately kept the cell subsumptive graph shallow, as we intend to infer neuron type based on their specific properties, e.g., primary neurotransmitter, number of processes, anatomical location of cell parts, according to the needs of a given user (see below).

Cell Part Descriptions: The SAO comprises two main classes of cell parts: *regional part* and *component part*. *Regional part of cell* is elaborated under the BFO concept *Fiat Object Part* defined as a part of an object not fully demarcated by physical discontinuities (`http://www.ifomis.org/bfo/1.0/snap\#FiatObjectPart`). Regional parts of neurons include dendrites, axons, the cell soma and protrusions such as dendritic spines. Each of these regional parts may be further subdivided into finer parcellations. *Component parts* are considered to be BFO independent objects (`http://www.ifomis.org/bfo/1.0/snap\#Object`) and represent the building blocks common to all cells, e.g., organelles. Components are largely adapted from the Gene Ontology cell component hierarchy (Gene Ontology Consortium, 2002), with additional neuron-specific components and sub-components added where necessary. Macromolecules are represented as component parts, but are listed in the SAO as children of a bfo:Object class distinct from cellular component. As with cell types, the macromolecule types are not exhaustive, as many are covered in other public semantic resources.

Supracellular Domain Descriptions: The *supracellular structure* class, though seemingly in conflict with the subcellular orientation of SAO, represents multicellular domains defined by subcellular parts of neurons such as neuropil, synapses, and the Node of Ranvier. Supracellular indicates the subcellular parts are derived from at least two different cells. To classify supracellular domains according to the BFO, we used both aggregate object (`http://www.ifomis.org/bfo/1.0/snap\#ObjectAggregate`) and site (`http://www.ifomis.org/bfo/1.0/snap\#Site`). *Supracellular aggregate objects* represent a somewhat ad hoc grouping of cell parts into higher order structures matching the bfo:ObjectAggregate defintion: "a mereological sum of separate objects possessing non-connected boundaries." For example, the neuropil is a term applied to regions of the nervous system characterized by a dense tangle of intertwined cell processes each having distinct non-connected boundaries. *Supracellular sites* locate particular functions and/or cell part configurations, e.g., the synapse is the site at which neurotransmission occurs; inferred because of the presence of synaptic vesicles and

post synaptic densities. For some aggregate structures, we create an *aggregate object* and then a *site* where the object is located, e.g., the chemical synapse may be considered an abstract aggregate entity consisting of a pre-synaptic part, a post-synaptic part and a junctional part. Each of these parts have cell components (e.g., synaptic vesicles) located within them circumscribing the extents of these parts. In this way, we restrict a synaptic site to portions of cells and cell parts occupied by *synaptic components*.

OWL Properties in SAO: Properties in the SAO are asserted against objects via morphological, mereological, and spatial relationships. Regional parts are assigned to each cell class using restrictions, e.g., neurons may only have neuronal regional parts. Each regional part is assumed to belong to a parent cell; topological relations specify spatial relatedness amongst cell parts such as *continuous_with*, e.g., dendrites are *continuous_with* the cell somata. Morphological relations describe subcellular structure shape or layout, e.g., bundled, multipolar. Although some properties assign to cell classes, e.g., morphological_type, most are asserted against regional cell parts. Cell components and macromolecules are asserted to localize to the part of the nerve cell in which they are found. Nerve cells being large, often spanning many brain regions, *has_anatomical_location*, which situates a cell within a nervous system region is assigned separately to each part of the cell. Thus SAO represents brain connectivity based on the contiguous paths formed by individual neuronal components.

3 SAO ENABLED TOOLS

One of our initial motivations for building the SAO was to annotate the content of LM and EM imaging data. We have built an on-line database, the Cell Centered Database (CCDB; `http://ccdb.ucsd.edu`)(Martone et al, 2003) to house such datasets. Because the interpretation of imaging data requires a thorough understanding of the methodological details, the data model of the CCDB describes the workflow required to prepare a specimen, obtain a set of microscopic images, derive a 2D, 3D or 4D reconstruction from those images and then segment meaningful objects from the reconstruction (Fig. 1). The CCDB thus mostly describes the "data" rather than the biological reality represented by the data. The process of segmenting electron tomographic data is essentially a form of annotation. Using either manual or semi-automatic methods, users define cellular andsubcellular components in the volume. The process of segmentation can be very time consuming, taking weeks to months for detailed tomographic volumes. To ensure standard annotation of segmented structures, we incorporated the SAO into the process of segmentation (Fig. 1). Each segmented object contained in CCDB is annotated as an SAO instance, and deposited into an instance store called the Cellular Knowledge Base (CKB).

The CKB and CCDB together contain much of the information to create detailed cellular models situating macromolecules and subcellular structures within regional parts of nerve cells. SAO acts as unifying, semantic "glue" in this environment (Fig. 2). The ontology browsing and query interface layer, OntoQuest, supports easy data exploration via the ontology. Unlike in Protege, a user navigates along any transitive relation exploring the various flavors of parthood and containment. OntoQuest creates a single back-end index across multiple instance stores allowing users to assemble these disparate elements via instance-fetching queries. The current OntoQuest version written in Java and querying SAO stored in a PostgreSQL-hosted OWL store updates an earlier prototype (Chen et al., 2006).

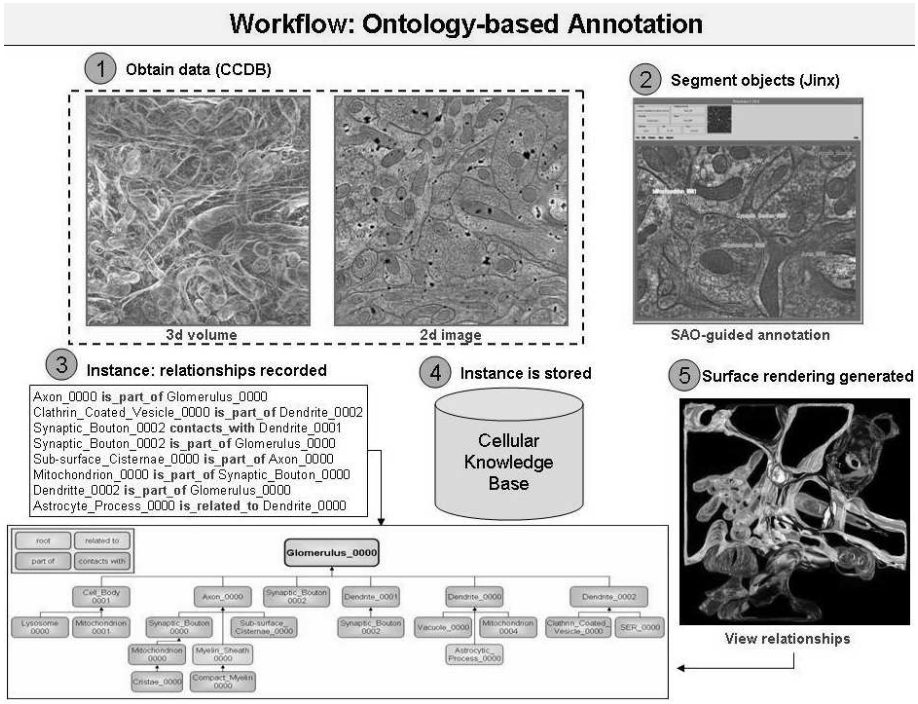

Fig. 1. Workflow for ontology-based annotation of electron tomographic data. Users select a dataset from the CCDB and perform manual segmentation using a custom segmentation tool developed in our laboratory, Jinx. In the ontology-enhanced version of Jinx, the user selects from an entity in the SAO to name an object. Each segmented object is named as a numbered instance of an SAO class, e.g., Axon_0000. Entities may be related to other entities through a simple set of relationships, e.g., clathrin-coated vesicle_0000 is part_of dendrite_0002. Only a subset of SAO is presented to avoid placing too much complexity before naive users. Selecting a class returns subtypes and mereologically related classes, guiding users to classes asserted to be related. Users may also add novel terms as needed to describe their data. These terms are automatically gathered, but currently migrated into the ontology manually by ontology curators. The segmentation tool's legacy XML export has been modified to include the SAO instances of related parts. This XML is currently imported to the CKB, though Jinx is being modified to write directly to the Protege API (see below).

OntoQuest provides a means to add, remove, and switch between ontologies in the data store and to add instances. We have also implemented a function whereby users can aggregate properties from a set of instances. For example, users may issue a query "What organelles are found in dendrites?" These aggregate properties are computed by first collecting instances (in this case organelles declared part of a dendrite) through an internal SPARQL query to the local index; the indexed instances are then fetched from the instance stores and the aggregate operations (i.e., the count and group-by) are performed locally inOntoQuest. Most of the queries currently supported by OntoQuest do not require a separate reasoner (e.g., Pellet).

Ontoquest can interact with our SAO-enabled segmentation tool, Jinx. For example, Jinx can be used by a neuroanatomist in segmenting vesicles and the active zone in an TEM view of a synapse.

They would label each segmented object with the appropriate SAO object class. Jinx would then send both the SAO instance labels and geometry data to a process that would automatically do morphometric analysis to determine actual distance between the objects. All of this information would be fed back to OntoQuest where rules established to determine the maximum distance for vesicle availability would be used to assert *sao:isAvailableForRelease* for that vesicle instance. It should be noted that using the increased DatatypeProperty expressivity in OWL 1.1, Jinx itself will be able to more succinctly represent the distance associated with this proximity relation.

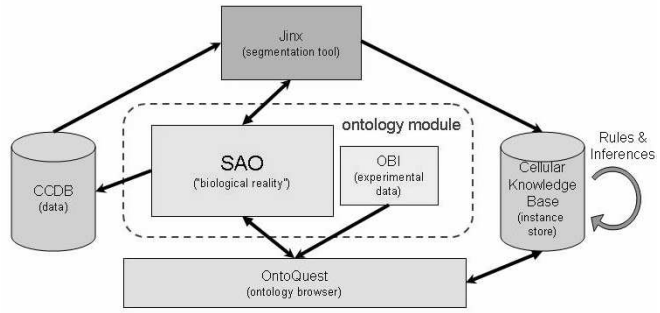

Fig. 2. High-level architectural diagram of SAO-enabled environment. OntoQuest is a stand-alone application also accessible via Java Web Start; we expect to distribute OntoQuest as a Protege plug-in in the near future.

4 DISCUSSION

We created an OWL ontology representing the subcellular anatomy of the nervous system to provide the necessary scaffold for integrating molecular and anatomical data through accurate description of mesoscale anatomy. This scaffold is amenable both to tool development and to semantically-driven information exchange across the field. It also provides individual researchers a means to perform reasoner-based quality control and inferential analysis of annotated neuroimages. Applying formal semantic representation techniques to neuroanatomical structure has been preliminarily addressed in the macroscopic domain (Martin et al., 2001; Mechouche et al., 2006); little exists in the mesoscopic neuroanatomical domain as yet. A Synapse Ontology was recently constructed (Zhang et al., 2007), but it does not situate synapses in their cellular contexts, nor is it built on top of community-shared foundational ontologies or provide the full expressivity of the combination of OWL-DL (current DL complexity merely AL) and SWRL rules logic we describe here and in Larson et al. (submitted). The SAO is publicly available and has been deployed in a prototype system developed around the CCDB, an on-line database for cellular and subcellular imaging data. In the following, we discuss our experiences in constructing the SAO, its integration into our database tools and our initial attempts at utilizing the SAO for rule based reasoning to bridge multiple anatomical scales.

Constructing the SAO: We made a good faith effort to construct the SAO using the best practices put forward by the OBO Foundry project. We provide here an assessment of our experience. At the

outset, we were admittedly naive with respect to ontology creation and tools. We thus recognize the problems encountered may be due to our inexperience. However, problems may also arise because the core ontologies and tools are not yet sufficiently hardened for widespread use. In some cases, we had to compromise between strict adherence to recommended practice and practicality.

Recommended practices, e.g., the re-use of *gold-standard* public ontologies, provision of consistent definitions, and the avoidance of multiple inheritance in the class hierarchy (Smith et al., 2005), are designed to promote re-use, semantically-driven alignment and programmatic interoperability. Though use of an Upper Level Ontology is not explicitly included in the Foundry principles, use of BFO as the foundation for SAO not only facilitates such interoperability with the other BFO-based resources, it also guides the SAO representation of biological reality based on well-founded, formal definitions that have been applied effectively in the biomedical domain (Grenon et al., 2004). We believe both these features aid in the perniciously difficult informatics task of formally describing neuron cell types (Migliore and Shepherd, 2005) in an algorithmically useful manner.

In SAO v1.0, we describe only bfo:continuant types, because EM studies require the analysis of static images. We found the BFO continuant class parcellations very useful, particularly for supracellular structures such as synapses. These structures are defined by parts of two or more cells. In earlier versions of the SAO, we described these structures as objects equivalent to a cell component like an organelle. However, unlike mitochondrion or a dendrite which are distinct subcellular entities, supracellular aggregates represent a functional parcellation of mesoscopic elements. They are sites where evidence suggests specific functions take place - a fact we believe is well captured by the BFO:Site class.

Despite its utility, we've confronted several difficulties in trying to use BFO in an OWL framework. The current BFO.owl implementation doesn't begin to match in expressivity (DL complexity: ALC) the formal definitions given in BFO publications. Though BFO provides useful class examples to guide implementation, its class definitions lack clarity and the logical weight provided by necessary and sufficient restrictions. This is a serious drawback to current BFO utility, as the required foundational class assertions must be assembled by application ontology developers (e.g., providing foundational OWL ObjectProperties defining appropriate domains and range to link *bfo:Sites* to their appropriate *bfo:Function* and *bfo:IndependentContinuants* types). Continuant granularity and boundary vagueness (Rector, A, et al. 2006) also must be addressed more effectively. For instance, the location of a *sao:Synaptic_vesicle* instance relative to *sao:Presynaptic_active_zone* has particular functional implications (e.g., only within a certain proximity to the active zone, a vesicle is considered available for release).

The OBO Foundry project promotes using the OBO Relations Ontology (OBO RO; Smith et al., 2005) for this task, but because it does not directly refer to the classes it is designed to constrain, it too has limited expressivity (DL complexity: ALR+HI). It would be desirable for BFO to come *pre-loaded* with OBO RO entities not restricted to the biomedical domain and asserted against the appropriate foundational classes - an effort currently under consideration. For example, a *bfo:Site* is a SpatialRegion where other *bfo:Objects* can be located. This implied structural relation should be made explicit using the RO ObjectProperty "located_in." We have experimented with importing the OBO RO into the SAO, but without integration into BFO, we cannot assign the relations without making inconsistent assertions. We therefore added such relations to SAO, which we expect to transfer to RO-based ObjectProperties declared in BFO once these are available.

We recognize other foundational layers providing a biological context exist - e.g., Simple Upper Bio Ontology (`http://www.cs.man.ac.uk/~rector/ontologies/simple-top-bio/`); however, given our

current invested effort in mapping to BFO, we must leave it to others for now to determine how to reconcile any such foundational interoperability issues. More recently Biotop v1.0 (`http://www.ifomis.uni-saarland.de/biotop`) has been released to provide a shared biomedical contextual layer on top of BFO. Here, too, it is not completely clear, yet, how to use this in a BFO context and still maintain decidability, and BioTop also has limited expressivity (DL complexity: SH) making it difficult to employ as a shared biomedical semantic framework. Though foundational ontologies ought to be kept thin in terms of class structure, they still need to provide sufficient expressivity if they are to play the critical semantic interoperability role for which they are intended.

We found Protege-OWL very useful for developing ontologies and the community very responsive in helping to resolve difficulties we encountered. The mechanism of importing OWL ontologies into Protege accommodates both OWL files on a local disk or those remotely located on the web. We opted to import ontologies by URL, since our aim was to distribute the SAO on the web ourselves. For some ontologies, e.g., the BFO, the process went smoothly. Problems arose when ontologies for import were: 1) not yet online; 2) online but malformed. The first issue arose because the annotation properties of the SAO come from the BIRN project ontology BIRNLex (see below) only very recently made available via a stable URL. The second issue arose with an online ontology, SKOS, with a badly formed OWL tag. It turned out this problem had been corrected in the RDF version of SKOS, but that version did not include the OWL AnnotationProperties we intended to re-use. These issues demonstrate the fragility of the import mechanism of OWL ontologies in practice. We expect the expanded annotation capability in OWL 1.1 (e.g, annotations using restricted instance lists and semantically neutral comments) will be a significant help in avoid this problem in the future.

We will continue to use DL consistency checks and species classification to maintain the decidability of SAO. The OWL mechanisms for inter-relating distinct ontologies (e.g., owl:import for whole ontologies and rdf:description for individual nodes), however, can each bring their own particular disadvantages interfering with our efforts to maintain DL decidability. We hope as we move to use of the OWL API with tools such as Protege 4, the E-Connections mechanism (Cuenca Grau, et al. 2005) present in that framework will help to address some of these issues. Other general limits in OWL such as the inherent binary nature of relations, the generic nature of property descriptions, and the related lack of temporal expressivity are gradually being addressed (see the OWL 1.1 specification), and we intend for our continued work on SAO to provide use cases to test the proposed solutions.

SAO-based tools: Data contained in the CCDB have been annotated as instances of the SAO and stored in a knowledge base accessed by our custom application OntoQuest. The SAO was designed to provide a fairly generic view of neuronal ultrastructure. Although certain cell classes are characterized by distinct features, e.g., number of primary dendrites, we chose not to create many of these restrictions on SAO cell classes. Unlike gross anatomy, where we have many examples of a given structure from which to create a generic model, our sampling of the ultrastructural world is very sparse. We nearly always find exceptions to any rule of thumb used to identify a given class of cell and, as such, they cannot be represented in a constrained decidable descriptive logic. Instead, we intend the instances stored in the CKB to be used to derive inductive "rules" for cellular assembly. Through OntoQuest, users can view all instances of the entities described according to SAO. Query results are linked to the data stored in the CCDB, thus providing a rudimentary cell-centered query capability.

The necessity for thorough annotation of CCDB data with the SAO was the motivation behind the creation of ontology-based segmentation tools. Projects like the CCDB derived from the Human

Brain Project (Huerta et al., 1993), unlike many of the model organism database projects, do not employ data curators and annotators. Rather, these projects assume annotation will be performed by the research scientist contributing the data. Because there is currently little incentive for researchers to spend much time annotating their data for inclusion in a resource like CCDB, we reasoned annotation with ontologies should be integrated into the scientist's workflow rather than occur as a separate step. Although the segmentation tool was designed primarily to aid us in capturing information derived from electron tomography, we envision it will also assist the annotator/scientist by providing guidance about possible structures to be found within a cellular environment.

Multiscale Integration using the SAO: Unlike many ontologies developed for the nervous system starting at the gross anatomy level, the SAO uses the cell as the rallying point for understanding individualized gene expression and molecular detail, the activity of populations of cells, and the behavior of interacting brain regions. Both molecular entities and anatomical regions in the SAO are drawn from existing ontologies, e.g., the NeuroNames hierarchy (Bowden and Dubach, 2002) and the BAMS system (Bota et al., 2005). Thus, data in the CCDB annotated with the SAO becomes interoperable with atlases and databases covering these scales. The CCDB participates in a large scale data federation project, the Biomedical Informatics Research Network (BIRN; Grethe et al., 2005). BIRN provides a collaborative infrastructure for groups who are imaging the nervous system at multiple scales. Ontologies like the SAO provide the means to annotate evidentiary data repositories so as to infer their relation to other data collected at additional scales (Gupta et al., 2003). The ontology is then employed by agents such as the BIRN mediator (Astakhov et al, 2005) to bring together data contained in distributed data resources.

As described in Larson and Martone (submitted), the SAO serves as the basis for performing rule-based reasoning to bridge ultrastructural and gross anatomy. Because the SAO is built on a model of the cell, we can use this reasoning to infer from a local observation derived from EM - e.g., thalamic axon synapses on spine of cortical neuron - that thalamus projects to cortex. Although preliminary, we believe our approach of using OWL-DL together with BFO highlights the potential power of using formal ontologies for filling in missing gaps in information across scales and for realizing the full benefit of hard won biological data. As we expand the SAO, we will be looking to apply some of the new features of OWL 1.1 to both simplify and extend expressibility - we will report this development in future publications.

5 CONTRIBUTIONS

In this paper we have detailed an OWL-based ontology for subcellular anatomy that serves as the center of a knowledge-driven environment for neuroscience. We have described several motivations for its construction, and explained the basic organizational ideas behind its construction. We have constructed tools that take advantage of this ontology for the purposes of image segmentation and annotation, as well as for data analysis across large data sets. Additionally, we have reviewed upper-level ontologies in the community that our ontology imports and described some of their benefits and weaknesses. Finally we describe some key issues involved in the annotation of neuroscientific data for the purposes of aggregating across knowledge both within and across scales.

Acknowledgments. This work was supported by NIH grants NIDA DA016602 (CCDB), NCRR RR04050 (National Center for Microscopy and Imaging Research) and RR08605 (National Biomedi-

cal Computing Resource). The Bioinformatics Research Network is supported by NIH grants RR08605-08S1 (BIRN-CC) and RR021760 (Mouse BIRN). The Protege resource is supported by grant LM007885 from the United States National Library of Medicine.

References

1. Astakhov, V., Gupta, A., Santini, A., Grethe, J. S.: Data Integration in the Biomedical Informatics Research Network (BIRN). DILS 2005: 317-320, 2005
2. Bard J, Rhee SY, Ashburner M. An ontology for cell types. Genome Biol. 6(2):R21, 2005.
3. Bota M, Dong HW, Swanson LW. Brain architecture management system. Neuroinformatics. 3:15-48, 2005.
4. Bowden DM, Dubach MF (2003) NeuroNames 2002. Neuroinformatics. 1(1):43-59.
5. Chen, L., Martone, M.E., Gupta, A., Fong, L., Wong-Barnum, M. (2006). OntoQuest: Exploring Ontological Data Made Easy, Abstract, Very Large Databases Conference. 1183-1186.
6. Cuenca Grau, C, Parsia, B, and Sirin, E (2005) Combining OWL Ontologies Using E-Connections. Journal of Web Semantics. 4:1, 42 pp.
7. Gene Ontology Consortium (2002). Creating the gene ontology resource: design and implementation. Genome Research, 11, 1425-1433.
8. Grenon P, Smith B, Goldberg L (2004) Biodynamic ontology: applying BFO in the biomedical domain. Stud Health Technol Inform. 102: 20-38.
9. Gupta, A., Ludascher, B., Martone, M., Rajasekar, A., Ross, E., Qian, X., Santini, S., He, H. and Zaslavsky, I. (2003) BIRN-M: A Semantic Mediator for Solving Real-World Neuroscience Problems. Proceedings of SIGMOD, June 9-12, 2003, San Diego, CA,
10. Huerta, M., Koslow, S.H., Leshner, A. I. (1993) The Human Brain Project: an international resource. TINS 16: 436-438.
11. Hunter, P. J. and Borg, T. K. (2003) Integration from proteins to organs: the Physiome Project. Nat Rev Mol Cell Biol, 4: 237-243.
12. Larson SD, Martone ME, Rule-Based Reasoning With A Multi-Scale Neuroanatomical Ontology. OWL: Experiences and Directions workshop 2007, submitted
13. Lein ES, Hawrylycz MJ, Ao N, Ayres M et al. (2007) Genome-wide atlas of gene expression in the adult mouse brain. Nature 445: 168-176.
14. Lucic V, Forster F, Baumeister W. Structural studies by electron tomography: from cells to molecules. Annu Rev Biochem. 74:833-65, 2005.
15. Martin, RF, Mejino, JL Jr, Bowden, DM, Brinkley, JF 3rd, Rosse, C. (2001) Foundational model of neuroanatomy: implications for the Human Brain Project. Proc AMIA Symp 438-42.
16. Martone ME, Zhang S, Gupta A, Qian X, He H, Price DL, Wong M, Santini S, Ellisman MH. The cell-centered database: a database for multiscale structural and protein localization data from light and electron microscopy. Neuroinformatics. 1:379-95, 2003
17. Mechouche, A, Golbreich, Gibaud, B. Towards an hybrid system for annotating brain MRI images. OWL ED 2006.
18. Migliore M, Shepherd GM. Opinion: an integrated approach to classifying neuronal phenotypes. Nat Rev Neurosci. 2005 Oct;6(10):810-8.Perkins, G., Renken, C., Martone, M. E., Young, S. J., Ellisman, M. and Frey, T. (1997) Electron tomography of neuronal mitochondria: three-dimensional structure and organization of cristae and membrane contacts. J Struct Biol, 119: 260-272.
19. Peters, A. L., Palay, S. and Webster, H. d. F. (1991) The fine structure of the nervous system. New York: Harper and Row.
20. Rector, A, Rogers, J, Bittner T (2006) Granularity, scale and collectivity: when size does and does not matter. Journal of Biomedical Informatics (Special issue: Biomedical ontologies). 39:3, p333.
21. Rosse C, Kumar A, Mejino JLV Jr, Cook DL, Detwiler LT, Smith B (2005). A strategy for improving and integrating biomedical ontologies. AMIA 2005 Symposium Proceedings, 639-643.
22. Smith B, Ceusters W, Klagges B, Kohler J, Kumar A, Lomax J, Mungall C, Neuhaus F, Rector AL, Rosse C. (2005) Genome Biol. 2005;6(5):R46
23. Zhang, W, Zhang, Y, Zheng, H, Zhang, C, Xiong, W, Olyarchuk, JG, Walker, M, Xu, W, Zhao, M, Zhao, S, Zhou, Z, Wei, L. (2007) SynDB: a Synapse protein DataBase based on synapse ontology. Nucleic Acids Res. 35: D737-41.