

Annual Report for Period:08/2009 - 07/2010**Submitted on:** 08/04/2010**Principal Investigator:** Jones, Matthew B.**Award ID:** 0743429**Organization:** U of Cal Santa Barbara**Submitted By:**

Jones, Matthew - Principal Investigator

Title:

Semantic Enhancements for Ecological Data Management

Project Participants**Senior Personnel****Name:** Jones, Matthew**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Jones is the project lead and coordinator for the project. He contributes to the design and conceptualization of software extensions to support semantic data management, and supervises development staff on the project.

Name: Schildhauer, Mark**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Schildhauer is a co-PI on the project focusing on Knowledge Representation approaches, particularly with respect to observational data models. He is a leader in the development of OBOE, and is the principal liaison with the Scientific Observations Network (SONet).

Name: Bowers, Shawn**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Bowers is a co-PI on the project, helping to specify project goals especially with respect to semantic data representation and query capabilities.

Name: Madin, Joshua**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Joshua Madin is a co-PI on the project and has helped to specify the overall semantic data management product suite. He also contributes to ontology development in the biological science area.

Name: O'Brien, Margaret**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Margaret O'Brien is a co-PI on the project and is conceptualizing and leading the two scientific use case projects that will determine the capabilities needed by the proposed semantic data management tools.

Post-doc**Graduate Student****Undergraduate Student****Technician, Programmer****Name:** Berkely, Chad**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Berkley served as a software engineer on the project to develop the initial system prototype for semantic extensions to the Metacat system. Berkley was funded from the SEEK project but was working jointly on this project that extends SEEK concepts of semantic data query.

Name: Leinfelder, Ben

Worked for more than 160 Hours: Yes

Contribution to Project:

Ben Leinfelder has assumed the role of software engineer on the project, and is involved in designing and implementing the semantic extensions to Morpho, Metacat, and other related software subsystems.

Other Participant

Research Experience for Undergraduates

Organizational Partners

SONet

We have been collaborating with the Scientific Observations Network (SONet), which is an NSF-funded INTEROP project focused on advancing a core semantic model of scientific observations. Semtools is basing much of its development effort on the observation model being evaluated in the SONet project. Initially we are using the OBOE model as the core and building an annotation framework around that model. As the SONet project progresses we hope to produce a generic solution that allows interoperability with different observation models and ontology authoring approaches. The SONet team has provided valuable insight into authoring domain-specific ontologies and defining best-practices regarding ontology construction.

Other Collaborators or Contacts

SONet: We have been collaborating with the Scientific Observations Network (SONet), which is an NSF-funded INTEROP project focused on advancing a core semantic model of scientific observations. The initial Semtools proposal outlined work on just such a semantic model of observations, but was dropped during budget cuts as we felt the work would be addressed by SONet. Nevertheless, the tools that we propose in Semtools depend on the core model from SONet being completed, and so Schildhauer, Bowers, Madin, and Jones are all collaborating to create a core model for SONet. The first step of this process is to produce effective comparisons of the existing semantic observations models, and so we have organized a workshop for Fall 2009 that will include an Interoperability Challenge that can be used to compare the various models in use in the science communities. We expect this workshop to produce a straw-man core scientific observations model that can be used as the basis of observational data modeling in the Semtools project.

Activities and Findings

Research and Education Activities: (See PDF version submitted by PI at the end of the report)

Activities 2008-2009

Data for ecological and environmental studies quantify, among other things, the distribution and abundance of organisms; the processes that influence biological populations, communities, and ecosystems; and the environmental and anthropogenic drivers of these processes. Scientists increasingly rely on accessing and analyzing these

diverse data collected by cross-disciplinary communities of researchers to achieve synthetic, crosscutting insights into the environment that can address issues of fundamental importance to science and society.

Despite these needs, discovering these data is difficult. The precision and recall of data searches in data repositories is not satisfactory even at current collection sizes. Data archives like the Knowledge Network for Biocomplexity (KNB), the National Biological Information Infrastructure (NBII) Metadata Clearinghouse, and the Global Change Master Directory (GCMD) rely on semi-structured metadata with fields containing largely natural-language descriptions to provide search and browsing capabilities and to allow human use and interpretation of the data. These metadata enable simple keyword searches that return results generally related to the topics of interest, but they cannot be used to perform precise searches of the data archives. Ironically data sets with more extensive (natural language) metadata are included in search results simply due to the incidental mention of a term in an ancillary part of the metadata document. These extraneous results decrease the precision of the search, seriously reducing the efficiency in researchers' finding the data they need. In addition, because natural-language metadata does not generally rely on controlled vocabularies, researchers typically classify their data sets using ad-hoc descriptive terms, reducing recall. Given the number of synonyms and overlapping terms used in scientific disciplines, searches frequently miss relevant data because the search terms do not exactly match the terms used to classify the documents.

Activities during the first year of the project have been focused on the refinement of a core model for scientific observations (in collaboration with SONet), and on the development of a prototype semantic search system for Metacat. This prototype represents a proof-of-concept for semantic search approaches and will allow us to compare multiple search strategies. As shown in Figure 1, we have added support to Metacat for storing and managing OWL-DL ontologies and semantic annotations, and for reasoning and search services to support different semantic-search strategies.

To implement these extensions to Metacat, we used approaches that exploit the use of formal reasoning over an ontology designed to facilitate the semantic description of scientific observations (OBOE). In our current implementation, the Jena API is used to access ontologies and ontology terms within Metacat, and Pellet is used to provide reasoning services over these ontologies (e.g., to compute class subsumption hierarchies and to ensure ontologies added to Metacat are consistent). We also extend Metacat's XML management capabilities with support for managing semantic annotations. Ontologies and annotations added to Metacat are assigned unique identifiers (URIs), allowing both to be easily accessed through external applications (e.g., Protégé). Further, ontologies and annotations can be versioned using this URL scheme.

The Extensible Observation Ontology (OBOE) provides a high-level abstraction of scientific observations and measurements that facilitate the creation of domain-specific vocabularies for defining observation and measurement semantics. OBOE is represented using OWL-DL and enables data (or metadata) structures to be linked to domain-specific ontology concepts so that critical aspects of scientific observations can be documented: such as what kind of Entity was measured, which Characteristics of that entity were measured and by which Measurement Standards (e.g., kilograms/m²), and what other observations provide Context for interpreting those measurements. In our approach, semantic annotations are then used to map these critical parts of a scientific observation to the data instances that are stored in a data repository (see Figure 2).

In addition to plain-text keyword search, we implemented three different search methodologies to investigate the utility of semantic methods for scientific data discovery: (i) simple term expansion against ontologies to broaden the search terms against the metadata corpus; (ii) term expansion against semantic annotations; and (iii) structured searches that pose queries against the components of an observation described via OBOE.

Activities 2009-2010

After focusing on the development of a prototype semantic search system for Metacat in the first year of the project, we've turned our attention to the process of actually annotating observational data so that it can be discovered using these specialized semantic search strategies. Through a weekly feedback loop between the PIs and development team, we've created a new plug-in for the Morpho data management application that allows users to annotate packages and also to locate data packages that have previously been annotated. Morpho development has coincided with authoring a domain-specific ontology that effectively describes data collected at the Santa Barbara Coastal LTER.

At the core of our development is the Semtools API that can be used in both Metacat and Morpho for managing annotations and ontologies. The common features of the library provide an opaque interface for leveraging semantics. The abstracted and centralized API buffers us from the vagaries of rapidly evolving Knowledge Representation technologies. For instance, we were able to switch from Jena API to OWL API as the ontology management library without having to refactor other systems that were utilizing the existing Semtools API.

The Annotation Plugin for Morpho augments the existing data table view with a tabbed interface that highlights different aspects of the annotation process: Summary, Column, and Context. Given the complexity inherent in formally describing observational data, the team has tried to minimize confusion by keeping the focus of the annotation activity on the actual data table being annotated. Furthermore, we've adopted a fill-in-the-blank format for collecting the crucial facts that comprise a complete annotation. These self-explanatory, natural language sentences prompt the user for information and also provide built-in help.

The search interface for locating annotated data packages has been directly informed by enhancements to the fill-in-the-blank approach. Sharing similar user interface motifs makes the transition from semantic data consumer to semantic data provider all the more natural. We expect users will initially interact with the plugin's query facilities before tackling their own data as annotators.

The search interface strikes a balance between expressivity and simplicity. Inspired by a familiar "smart playlist" criteria building interface, the search panel can be used to create compound nested queries that maximize search precision, and reduce false positive matches. To complement search precision, we use result ranking to maintain broad recall with the "best" or "closest" matches appearing at the beginning of the search results. Having a solid interface for defining search criteria allows us to continue our focus on evaluating structured, ontology-based query strategies that exploit both explicit and inferred concept relationships.

Findings:

Findings 2008-2009

Our semantic search system adds to Metacat the ability to store OWL-DL ontologies in addition to semantic annotations that link data set attributes to ontology terms. Our approach also extends Metacat to improve metadata search in multiple ways: (i) by expanding standard keyword searches with ontology term hierarchies; (ii) by allowing keyword searches to be applied to annotations in addition to traditional metadata; and (iii) by allowing more structured searches over annotations via ontology terms. We describe our implementation of these extensions, and compare and contrast these different types of

search for a corpus of annotated documents. As data repositories continue to grow, these tools will be instrumental in helping scientists precisely locate and then interpret data for their research needs.

Figure 1 shows the primary components of our semantic-discovery framework. The bottom of Figure 1 consists of two simple, example data sets. Although different types of data are often used within ecological analysis (e.g., raster, GIS, etc.), data sets are predominantly tabular (relational) and denote sets of related observations and measurements that were either directly collected or were the result of aggregation or analysis. Although not obvious, the example data sets in Figure 2 contain largely similar information consisting of spatial locations divided into sub-locations (i.e., a plot or quadrat), fertilization treatment information, and weight measurements.

Metadata schemes such as the Ecological Metadata Language (EML) provide standard ways of describing implicit aspects of data sets. In Figure 2, we show a fragment of EML for describing the ?wt? and ?LL? attributes of the data sets. EML can be used to represent the basic structural aspects of data—the number of attributes, their names, and their allowable values—but the semantics of the data set—the types of entities observed, the characteristics of these entities that were measured, and how these entities were observed in relation to each other—is either indirectly described (e.g., within the methods section of the metadata document) or are altogether missing. Metadata alone would not reveal the closely related semantics of the highlighted attributes from our sample data sets.

Semantic annotations extend EML by providing a mechanism to describe data set attributes in terms of OBOE concepts. An annotation is a formal structure, which represents a mapping from data set values to ontology instances (i.e., individuals), and an XML-based syntax is used to represent annotation mappings. As shown in Figure 2, we can see that the two annotated attributes: (i) represent observations of leaf-litter entities; (ii) measure the weight of leaf-litter (although using different weight characteristics); and (iii) use compatible but different measurement units (kilograms and grams). Annotations can be used to find all data sets related to a particular concept, determine all of the concepts related to particular data set attributes, and compare data sets based on their corresponding OBOE structures (which can facilitate data integration). XML is used as an interchange format for representing annotations; in general, annotation providers will annotate data sets using higher-level metadata editors and interfaces provided through tools such as Morpho.

A more detailed example showing the various XML syntaxes used for representing EML attributes (bottom), semantic annotations (middle), and an OBOE ontology extension (top) are shown in Figure 3.

To improve overall precision and recall of Metacat searches we prototyped three new search strategies.

Keyword-Based Term Expansion. In this approach, we “intercept” keyword queries issued to Metacat and expand them according to the term hierarchies of the stored ontologies. Specifically, if a given search keyword matches a class name (i.e., as specified by the `rdf:label` property of the class), then the search is expanded to include the synonyms of the class as well as the names of subclasses. This form of search alleviates the problem with simple keyword searches of not returning data sets described with synonyms or more specialized terms of the user-entered keywords. In our implementation, when a user enters a keyword search, Metacat locates synonyms and corresponding subclasses for each keyword using an ontology manager. The query is then augmented by Metacat with the expanded terms according to the given search constraints (i.e., whether terms should exactly match document terms and whether all given keywords must be present in a document) and executed against the current Metacat keyword search service. Although this strategy improves recall for documents that may have been omitted with simple

keyword searching, it can also cause additional false positives due to the addition of keywords. Thus, this approach generally increases recall, but not necessarily precision. Specifically, the set of metadata documents returned is always a superset of what is returned using the traditional Metacat keyword search.

Annotation-Enhanced Term Expansion. Semantic annotations allow individual data set attributes to be linked to one or more ontology classes. By applying keyword searches only to annotations, search results can potentially improve precision by returning fewer false positives. In annotation-enhanced search, each search term is first expanded using the ontology similar to the keyword-based term expansion. Here, when a search term matches an ontology class, we use the class and all subclasses to find matching annotations. An annotation is considered a match if it contains the corresponding classes according to the search constraints (see above). The metadata document linked to the matched annotation is returned by the search. For example, a metadata description of a data attribute described textually as ?counts of grasshoppers? would be annotated as a ?count per square meter? of ?Romalea guttata? by linking the attribute to the ontology classes that define these concepts. When the user searches for ?grasshopper?, the term is expanded to ?Romalea guttata? via the ontology?s class hierarchy, and the annotation linking the metadata attribute to the class ?Romalea guttata? becomes a match. Since the annotation is linked to a specific attribute reference within the metadata, data sets containing comments about ?grasshoppers? in other fields would not be matched. Moreover, recall is improved due to matches facilitated by descending the ontology?s class hierarchy.

Observation-Based Structured Query. Though the annotation-enhanced term expansion approach limits search to the relevant portions of metadata that describe data content (via attribute annotations), it does not take advantage of observation and measurement structures and relationships. In the observation-based query approach, users can search for data sets via their observed entities (organism, site, etc.) and the characteristics and standards used to measure them. In an observation-based search, queries are specified by explicitly filling in an observation ?template? where ontology classes are given for the observed entity, measurement characteristic, and measurement standard. To search for data sets containing tree lengths, we would fill in an observation query template, using the tree class as the observed entity type and the length class as the measurement characteristic type. Search is performed by finding matches between the observation types of annotations and the query template, where a match includes searching subclasses of the template classes. This type of search has both good recall ? hitting all relevant data through appropriate use of term expansion ? and good precision by exploiting the structure of OBOE annotations to find exactly the entity, characteristic, and context of interest to the user.

The search interfaces in Metacat were implemented rapidly to explore the implications of different search strategies. Our next stage of development is to design an effective user interface for composing semantic queries and then to use the semantic search engine to execute those queries.

Findings 2009-2010

In developing the Annotation Plugin for Morpho, we?ve found that the inherent complexity of fully describing an observational data table begs for a compact visualization of the annotation in-progress. Graph-based representations of the annotation ? used extensively in the SONet for illustrating OBOE ? show promise for effectively conveying the structure of the observation model as applied to a specific data table. Moreover, the succinct fill-in-the-blank summaries of each observation and their relationship[s] to one another encourage a casual description that is in fact formally rigorous and adheres to our strict Entity/Characteristic/Standard/Protocol annotation mapping. We employ this same format

in the search template used for observation-based structured searches.

Figure 1 shows a sample data set that has been partially annotated against the OBOE model. The biomass column ?wt? is annotated: ?The WetBiomass of the GiantKelp was recorded in Grams.?

Figure 2 graphs this annotation and also contextualizes the biomass observation Within a Subplot that is itself located Within a Plot. Context is an integral piece to the observation model as it determines how the data should be interpreted and analyzed.

Observation-Based Structured Query. With a corpus of well-described, fully annotated data packages, observation-based searches allow for very precise matching with a small likelihood of false positives commonly introduced by free-text descriptive fields used in keyword-only queries. By populating a template that specifies the observed Entity, the Characteristic measured, the Standard (unit) for that measurement, and the Protocol used for collecting that measurement, a query can unambiguously target data tables that contain those observations. Context relationships can be specified to further restrict the packages returned.

Domain-specific ontologies. In order to evaluate the effectiveness of the Morpho Annotation Plugin we've developed an extension to the OBOE ontology that can be applied to existing real-life datasets collected by the Santa Barbara Coastal LTER. O'Brien has selected a subset of data that investigates the nitrogen supply in giant kelp forests. The ontology organizes and formalizes such concepts as the observation Entity (GiantKelp), Characteristic (WetBiomass), Measurement Standard (Gram), and Protocol (wet vs. dry). Applying concepts to observations captured in tabular data columns requires a fairly comprehensive ontology. When annotating from the ontology, it is desirable to select the narrowest (most specific) concept available such that broad-concept searches can include all subclasses under that concept's hierarchy. Axioms defined in domain ontologies can be exploited by reasoners during both annotation authoring and annotation searching. We are poised to use Pellet for inferring indirect class subsumption hierarchies and also to validate imported ontologies as being logically consistent.

Training and Development:

Outreach Activities:

Outreach activities for the project have principally been through talks at scientific conferences and workshops where we have discussed our approaches to semantically modeling scientific observations and the benefits of doing so, and common use cases. O'Brien is consulting with SBC LTER ecologists and oceanographers in the development of the domain-specific ontology. Presentations on Semtools and SONet related work included:

Jones, M. 2008. Directions in observational data organization: from schemas to ontologies. Biodiversity Information Standards (TDWG) Annual Conference. Freemantle, Australia. 19-25 October, 2008.

Schildhauer, M. 2008. Facilitating data interoperability within the environmental and ecological sciences through advanced semantic approaches. Biodiversity Information Standards (TDWG) Annual Conference. Freemantle, Australia. 19-25 October, 2008.

Schildhauer: Improving Data Discovery in Metadata Repositories through Semantic Search. CISIS/iSEEK. Fukuoka, Japan. March 18, 2009

Jones: Semantic Data Integration for Heterogeneous Scientific Data. Lifewatch WP5 Workshop on Semantic Data Integration. Amsterdam, Netherlands. May 18, 2009.

David Thau and Shawn Bowers. Best Effort Data Exchange of Taxonomically Organized Data. International Workshop on New Trends in Information Integration (NTII), In conjunction with ICDE, 2010.

Journal Publications

Berkley C, Bowers S, Jones MB, Madin JS, Schildhauer M, "Improving Data Discovery in Metadata Repositories through Semantic Search", Proceedings of iSEEK'09. IEEE Computer Society., p. , vol. , (2009). Published,

Shawn Bowers, Joshua S. Madin, Mark P. Schildhauer., "Owlifier: Creating OWL-DL ontologies from simple spreadsheet-based knowledge descriptions.", Ecological Informatics, p. 19, vol. 5 (1), (2010). Published, doi:10.1016/j.ecoinf.2009.08.010

David Thau, Shawn Bowers, Bertram Ludwiger, "Merging Sets of Taxonomically Organized Data Using Concept Mappings under Uncertainty", OTM Conferences, p. 1103, vol. 2, (2009). Published,

Books or Other One-time Publications

Web/Internet Site

URL(s):

<http://semtools.ecoinformatics.org>

Description:

This is the main web site for the Semtools project. It is a collaborative wiki that was established to aid communication among project participants and help with organization and outreach for the project.

Other Specific Products

Contributions

Contributions within Discipline:

Through our work on Semtools, we have demonstrated improvements in the effectiveness of data discovery for large, heterogeneous data collections such as the Knowledge Network for Biocomplexity (KNB). These advances have been possible through the use of an semantic model of scientific observations (Extensible Observation Ontology) and an annotation language that is used to map relational data sources to the concepts in OBOE. The prototype system that we developed will form the basis for future work this year on a production system that will have broad applicability in the ecological and environmental sciences.

The current movement within the ecological sciences to develop ontologies for organizing and formalizing what was observed and how provides the semtools team with a good opportunity to exchange ideas about creating these ontologies. Our initial work with OBOE and the Morpho Annotation Plugin has illuminated questions about how disparate ontologies can be unified without having to conform to any one approach. As the annotation plugin matures and existing EML metadata is augmented with formal observational descriptions, locating and synthesizing heterogeneous data sources will be more efficient, and downstream analysis more accurate.

Contributions to Other Disciplines:

The relative newness of knowledge representation and the use of ontologies to express and formalize information in a way that machines can "understand" puts our real-life use of the technology at the forefront of the art. We are involved in the OWL API user community "a forum for both providing and soliciting support for the rapidly evolving software. Similarly, our extensive use of Protégé" increases the user base directly and indirectly as we encourage collaborators to view and author domain ontologies within this application.

Contributions to Human Resource Development:

Through collaboration with the SONet project, we have worked with the SONet postdoc (Huiping Cao) on the development of semantic tools for ecological data management. Cao helped develop an approach to materializing data and ontologies together in an integrated view for data subsetting. Cao has now accepted a position as a faculty member at NMSU, and SONet will be recruiting additional postdoc for the remainder of the project.

Contributions to Resources for Research and Education:

The project is helping to build the extensive Knowledge Network for Biocomplexity (KNB) repository, which provides thousands of data sets for use in research and educational contexts. Data from the KNB will become more accessible as the semantic search facilities that we have developed become incorporated into the production Metacat software used by the KNB. This will enable educators and researchers to more readily access KNB data and therefore facilitate science and education advances in many disciplines.

In addition, we have begun discussions with DataONE and the Data Conservancy, two major NSF projects for data preservation, about the incorporation of semantic observational data approaches into those networks. Over the next year, we expect to develop a tight collaboration with DataONE on co-development of semantic search tools that can be used in the DataONE system.

Contributions Beyond Science and Engineering:**Conference Proceedings****Special Requirements**

Special reporting requirements: None

Change in Objectives or Scope: None

Animal, Human Subjects, Biohazards: None

Categories for which nothing is reported:

Activities and Findings: Any Training and Development

Any Book

Any Product

Contributions: To Any Beyond Science and Engineering

Any Conference




Figure 1: Semantic extensions (highlighted in blue) to the Metacat data and metadata repository support improved precision and recall in searches for scientific data sets.

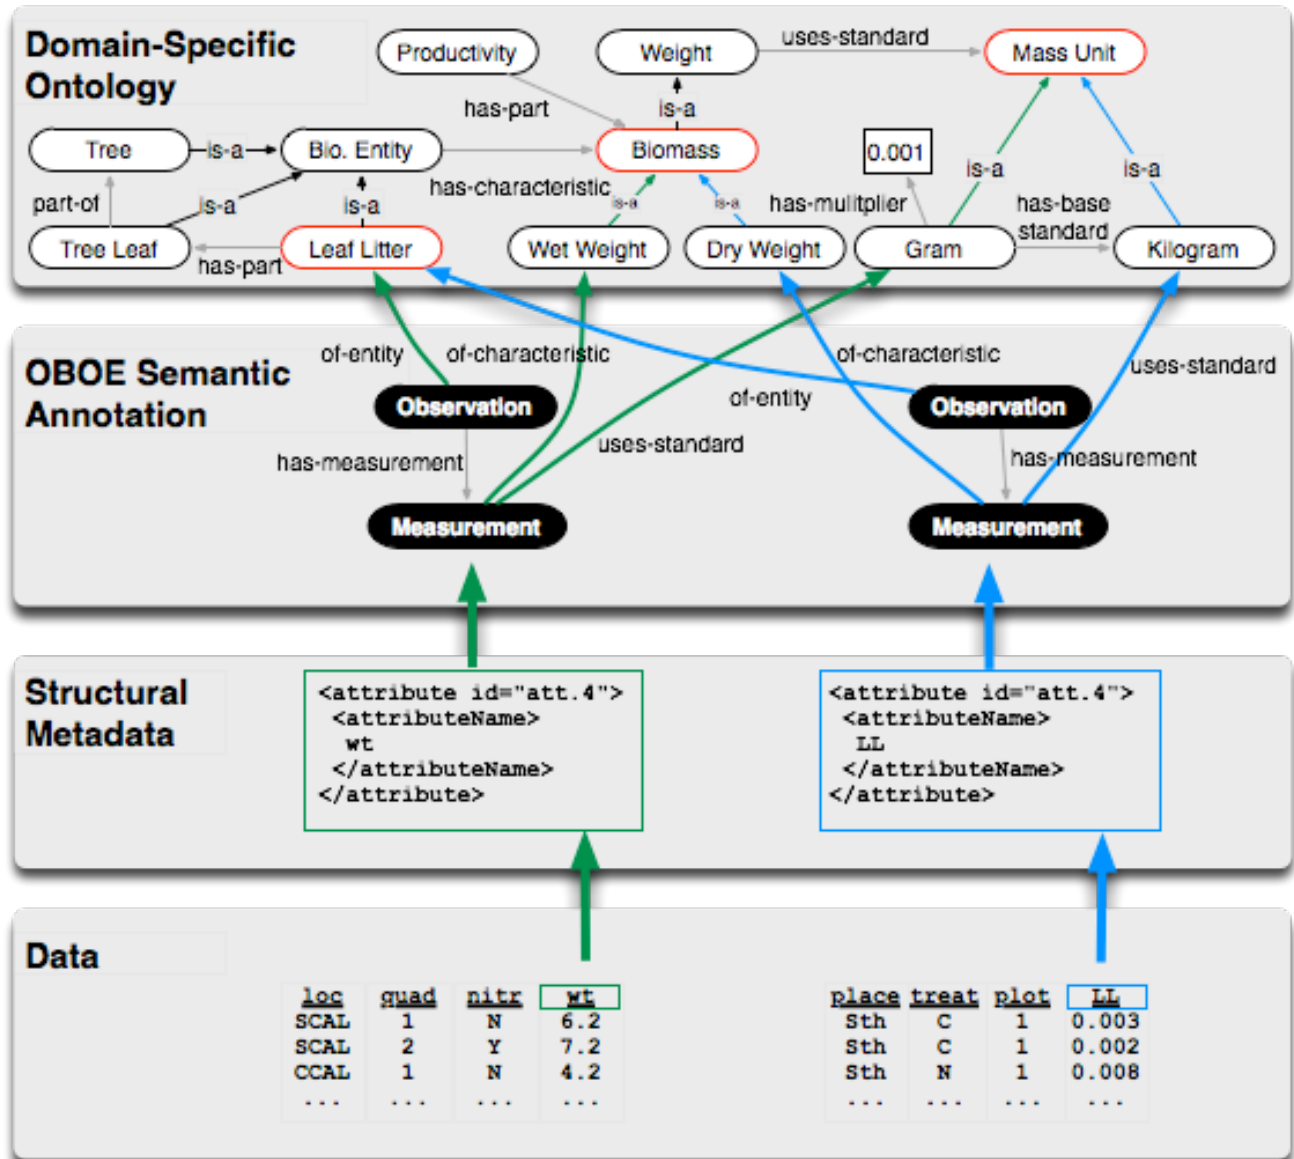
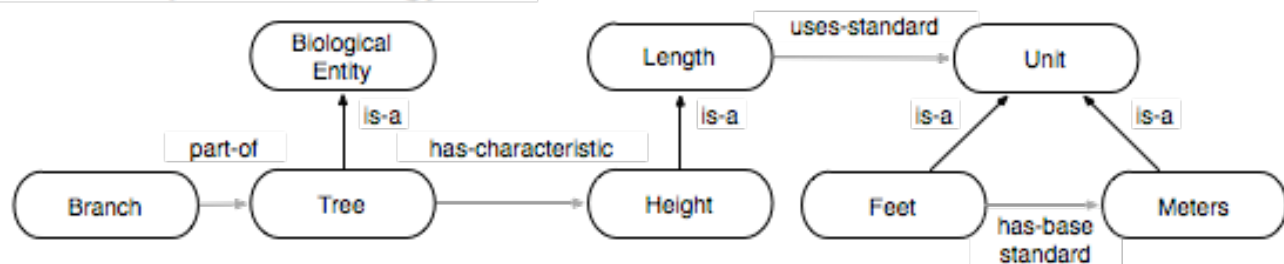



Figure 2. The components of our semantic-search framework including relational data, EML-based metadata, semantic annotations based on OBOE, and OBOE domain-ontology extensions.

Domain-Specific Ontology



Semantic Annotation



Data

loc	tree	ht
1	1	8.1
1	2	7.3
2	1	4.5
...

site	tree	len
A	1	1.1
B	2	1.3
B	1	0.5
...

Figure 3. Example annotations demonstrating more precise search results for observation-based structured query.