Ontology-Based Knowledge Acquisition for Neuromotor Functional Recovery in Stroke

Christopher Townsend*, Jingshan Huang*, Dejing Dou†, Haishan Liu†,
Lei He‡, Patrick Hayes§, Robert Rudnick*, Hardik Shah*, Dennis Fell¶, and Wei Liu¶
*School of Computer and Information Sciences, University of South Alabama
†Computer and Information Science Department, University of Oregon
‡Department of Information Technology, Armstrong Atlantic State University
§Florida Institute for Human and Machine Cognition
¶Department of Physical Therapy, University of South Alabama

***Abstract*—Hemiparesis is the most common impairment after stroke, the leading cause of adult disability in the United States. The initial severity of hemiparesis had been the strongest predictor of neuromotor functional recovery level. However, the intervention response of stroke survivors does not always correlate with their initial level of impairment. This implies the existence of other factors that may significantly affect stroke survivors' recovery process. In order to design targeting intervention therapy strategies, it is critical to consider these factors in a principled, comprehensive way so that physical rehabilitation (PR) researchers may predict which stroke survivors will respond best to therapy and subsequently, determine if a particular type of therapy is a more optimal match. Currently, such prediction is primarily a manual process and remains a challenging task to PR researchers and clinicians. We propose a computing framework based upon a domain-specific ontology. This framework aims to facilitate knowledge acquisition from existing sources via semantics-enhanced data mining (SEDM) techniques. As a result, it will assist PR researchers and clinicians in better predicting stroke survivors' neuromotor functional recovery level, and will help physical therapists customize most effective intervention therapy plans for individual stroke survivors.**

***Index Terms*—stroke; hemiparesis; extremity dysfunction; neuromotor functional recovery; intervention therapy; ontology; formal semantics; data mining**

I. Introduction

Stroke is the leading cause of adult disability in the United States. According to the American Heart Association's heart disease and stroke statistics [35], the prevalence of stroke in the United States was 5,500,000 and rising in the year of 2007. In addition, as reported in [38], more than 750,000 Americans are affected by a new or recurrent stroke each year. The estimated direct and indirect cost of stroke in the year of 2006 was as high as 57.9 billion US dollars [32]. Hemiparesis, i.e., weakness on one side of the body, is the most common impairment after stroke. Two-thirds of stroke survivors have persistent impairments that contribute to functional limitation and disability [12]. Among all stroke survivors, about 95 percent have some upper extremity (UE) dysfunction and 20 percent regain no functional use of their hemiparetic arm [11]. Although previous research, [17] for example, had demonstrated that the initial severity of hemiparesis, i.e., the initial level of impairment, is the strongest predictor of functional recovery level in sub-acute patients, A. Luft et al. [18] have recently discovered a range of responders who do not always correlate with their initial level of impairment. This research implies that there may exist other factors that play significant roles in determining stroke survivors' neuromotor functional recovery. In order to design targeting intervention therapy strategies, it is critical to consider these factors comprehensively. Only in this manner is it possible to predict which stroke survivors will respond best to therapy and, subsequently, to determine if a particular type of therapy is a more optimal match given certain characteristics of a specific stroke survivor. Currently, such prediction is primarily a manual process, which relies heavily on human efforts. The manual prediction is not only time-consuming and error-prone, but also subject to humans' limited prior knowledge. In addition, existing data sources from geographically distributed physical rehabilitation (PR) groups are usually developed in isolation, and with ad-hoc styles, e.g., relational databases, Excel spreadsheets, text files, etc. Due to the lack of agreed-upon and formal semantics, these structured, semi-structured, or unstructured data repositories have witnessed inconsistent terminologies, different schemas, and incompatible data formats. Consequently, it will be very difficult, if not impossible, for PR researchers and clinicians to obtain information of interest from such data sources, let alone to share and reuse existing information and assemble it into a more meaningful and useful block of knowledge. In order to handle this challenge, a formal knowledge model is needed in addition to traditional data mining techniques.

In computer science, ontologies are declarative knowledge representation models, playing a key role in defining formal semantics in traditional knowledge engineering and the emerging Semantic Web [2]. Noy and McGuinness [20] highlighted several reasons for developing ontologies: 1) to share domain information among people and software; 2) to enable reuse of domain knowledge; 3) to analyze domain knowledge and make it more explicit; and 4) to separate domain knowledge from its implementation. Each of these reasons plays an important role in our decision to develop an ontology that formally represents knowledge rather than a relational database that merely integrates data. Specifically, we argue the following:

- Relational databases focus on syntactic representation of data, lacking the ability to explicitly encode semantics, which is critical in automated knowledge acquisition.
- There are powerful tools available for capturing and managing ontological knowledge, including an abundance of reasoning tools that are readily supplied for ontological models. These reasoners (also known as inference engines) make it much more convenient to query, manipulate, and reason over available data sets. In particular, semantics-based (sometimes known as logic-based) queries, instead of traditional SQL queries, are made possible.
- Advances in the PR domain require that changes be made on a regular basis with regard to underlying data models. In addition, more often than not, it is preferable to represent data at different levels and/or with different abstractions. There are no straightforward methods for performing such updates if relational models are adopted.
- Ontologies, along with Semantic Web technologies, better enable PR researchers and clinicians to append additional data sets in a more flexible way. More importantly, the formal semantics encoded in ontologies makes it possible to reuse the data in unplanned and unforeseen ways, especially in cases where data users are not data producers, which is now very common.

Based on the above analysis and insights, we propose an innovative computing framework to facilitate (semi-)automated knowledge acquisition for PR researchers and clinicians to address the aforementioned challenge of predicting stroke survivors' improvements in extremity motor function. The framework is built upon a domain-specific ontology, the NeuroMotor Recovery (NeuMORE) Ontology, which is *the very first formal knowledge model* in the domain of stroke recovery. With the formal semantics defined in the NeuMORE ontology, and the logic-based inference (also known as ontological reasoning) that follows, it is possible to enhance existing data mining techniques. Resultant techniques from this unique combination are referred to as *semantics-enhanced data mining (SEDM) techniques*, which facilitate effective knowledge discovery, sharing, and reuse from existing sources. Ultimately, our research objective is **to assist PR researchers and clinicians in identifying factors that determine stroke survivors' neuromotor functional recovery process after stroke; thus, more optimal intervention therapy plans can be customized for individual stroke survivors**. We aim to synthesize data from existing sources into a comprehensive conceptual model, i.e., the NeuMORE ontology, which permits an emphasis on data semantics rather than on the forms in which the data was originally represented. Consequently, PR researchers and clinicians can acquire a more accurate, complete view of the set of factors that influence neuromotor functional recovery after stroke.

The rest of this work-in-progress paper is organized as follows. A brief overview of the related research is provided in Section II. The system framework and a series of tasks to be performed, along with proposed approaches and anticipated challenges, are discussed in Section III. Details of the NeuMORE ontology and its software implementation are described in Section IV. Finally, Section V concludes with future work.

II. Background and Related Research

A. Background Knowledge of Ontologies

An **Ontology** is a computational model of some portion or domain of the world [29]. The model describes the semantics of the terms used in the domain. An ontology is often captured in some form of a semantic network, i.e., a graph whose nodes are concepts or individual objects and whose arcs represent relationships or associations among the concepts. The semantic network is augmented by properties and attributes, constraints, functions, and rules that govern the behavior of the concepts. In brief, an ontology consists of a finite set of concepts (also known as "terms" or "classes"), along with these concepts' properties and relationships. In addition, most real-world ontologies have very few or no instances, i.e., they possess only the aforementioned graphical structure (also known as "schema"). Also notice that heterogeneity is an inherent characteristic of ontologies developed by different parties, therefore, ontology matching, which is the process of determining correspondences between concepts from heterogeneous ontologies, has become an active research area.

B. Factors of neuromotor Functional Recovery in Stroke

As indicated in [4], [33], and [37], several promising methods of repetitive, task-oriented UE stroke rehabilitation have been shown to be successful at motor function improvement in chronic patients. Currently, the initial severity of hemiparesis still remains the strongest predictor of stroke recovery. A longitudinal study [17] demonstrated that the time course of the change in the Barthel index (a measurement of disability) during stroke recovery is well approximated by a logistics regression model. This model indicated that the earlier patients showed recovery, the better the outcome at six months. The first week of Barthel index measurements explains around 50 percent of the variance in the outcome at six months. Another study [7] showed that the Fugl-Meyer scale (a measurement of function impairment) at 30 days predicted 86 percent of the variance in recovery of motor function at six months.

C. Ontological Techniques in Medical/Biological Research

Ontological techniques have been widely applied to medical and biological research. The most successful example is the Gene Ontology (GO) project [10], a major bioinformatics initiative with the aim of standardizing the representation of gene and gene product attributes across species and databases. Unified Medical Language System (UMLS) [34], the National Center for Biomedical Ontology (NCBO) [19], and NeuroLex (formerly BirnLex) [3] are some other successful examples.

In neuroscience–the electroencephalographic (EEG) domain in particular–D. Dou et al. [6], [9] propose data-driven and

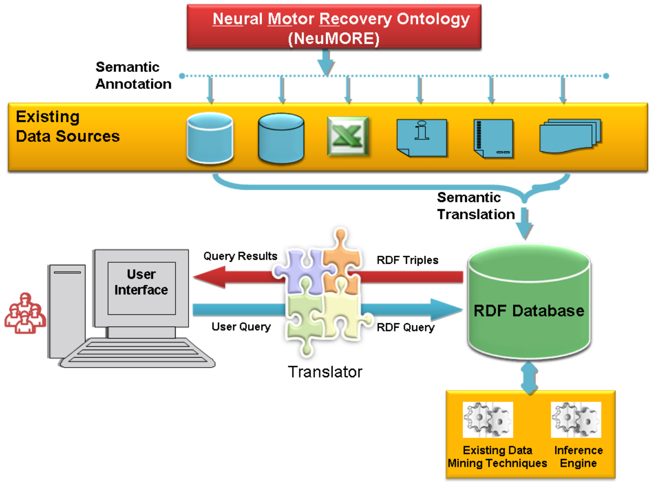

Fig. 1. NeuMORE System Framework

knowledge-driven ways to develop the Neural Electromagnetic Ontologies (NEMO). The NEMO is intended to be used for data sharing across several EEG labs worldwide. In the microRNA (miRNA) domain, J. Huang et al. [15] report the OMIT (Ontology for MicroRNA Target) framework to assist biologists in unraveling important roles of miRNAs in human cancer. The OMIT ontology was designed from GO concepts and relationships, the domain expertise of biologists in the project team, and six collaborating labs from around the world.

III. Methodologies

As shown in Figure 1, the NeuMORE framework consists of a series of tasks that are closely related to each other:

1) to develop a domain-specific ontology that contains a set of NeuMORE concepts, along with these concepts' properties and the relationships among these concepts;
2) to annotate existing data sources with NeuMORE concepts, and to enrich these sources with formal semantics;
3) to translate NeuMORE-annotated data sources into a RDF database; and
4) to perform complex search/query in a unified style so that *deep* knowledge can be obtained, shared, and reused out of a wealth of data sets.

A. Phase One: Domain-Specific Ontology

1) Ontology Development Language and Tool: There are different formats for describing an ontology, all of which are popular and based on different logics: Web Ontology Language (OWL) [24]; Open Biological and Biomedical Ontologies (OBO) [21]; Knowledge Interchange Format (KIF) [16]; and Open Knowledge Base Connectivity (OKBC) [22]. We have chosen the OWL format, which is a standard recommended by the World Wide Web Consortium (W3C) [36]. OWL is designed for use by applications that need to process the content of information instead of just presenting information to humans. As a result, OWL facilitates greater machine interpretability of Web contents. As for our development tool, we have chosen Protégé [26] over other available tools such as CmapTools [5] and OntoEdit [23].

2) Design Methodologies: While developing the NeuMORE ontology, we have observed the seven practices proposed by the OBO Foundry Initiative [30]. In this iterative, top-down, knowledge-driven approach, both ontology engineers and domain experts (PR researchers and clinicians) have been involved, working together to capture domain knowledge, develop a conceptualization, and implement the conceptual model.

In order to design top-level concepts in the NeuMORE ontology, we have reused and extended a subset of concepts from the Basic Formal Ontology (BFO) [1], which focuses on the task of providing a genuine upper ontology that can be used in support of domain ontologies developed for scientific research. In addition, we have also adopted some concepts from the Open Bio Ontologies' Upper Bio Ontology (OBO-UBO) developed by Berkeley Bioinformatics Open-Source Projects[1]. Note that the OBO-UBO itself was (partially) built upon the BFO.

Our development process has taken place over many iterations, involving a series of interviews, exchanges of documents, evaluation strategies, and refinements. In addition, revision-control procedures have been adopted to document the whole process for future reference.

B. Phase Two: Annotation of Existing Data Sources

Semantic annotation is the process of tagging source files with predefined metadata, which usually consists of a set of ontological concepts. Our annotation will take two steps:

- The first step is to annotate data schemas, resulting in a set of mapping rules between NeuMORE concepts and elements from existing data schemas. These mapping rules will be specified in W3C Rule Interchange Format-Production Rules Dialect (RIF-PRD) [28].
- The next step is to annotate original data sets, and the annotated data sets will be published in the resource description framework (RDF) [27]. Being a structure based on the directed acyclic graph model, the RDF defines statements about resources and their relationships in triples. Such generic structure allows structured and semi-structured data to be mixed, exposed, and shared across different applications.

The above annotation process is referred to as a "deep" annotation, a term coined by C. Goble in the Semantic Web Workshop in 2002, and further investigated in [13], [14]. It is necessary to annotate more than just data schemas because there are situations where the opposite "shallow" annotation, i.e., an annotation on schemas alone, cannot provide users with desired knowledge. The annotation outcomes, i.e., (1) a set of mapping rules between NeuMORE concepts and data schemas, along with (2) annotated data sets published in the RDF format, will become the input to the next phase.

[1] http://www.berkeleybop.org/projects.html

C. Phase Three: RDF Database in Semantic Translation

Instead of traditional databases (i.e., relational databases), we propose to create a RDF database for use in the semantic translation phase. By making use of a RDF repository to store RDF triples, our system makes it possible to query, manipulate, and reason about existing data through available tools, including, but not limited to, SPARQL Query Language for RDF [31] and OWL reasoners (e.g., Pellet [25] and FaCT++ [8]). In this manner, the logic-based knowledge acquisition can be realized, and existing data mining techniques will be enhanced and result in a unique combination that we propose as *semantics-enhanced data mining (SEDM) techniques*. With the correspondence between data sources and the global schema (the NeuMORE ontology) established in *Phase Two*, user queries will be translated into RDF-based queries expressed in NeuMORE concepts. As a result, the SEDM techniques should have a better performance.

D. Phase Four: Complex Query/Search in a Unified Style

As illustrated in Figure 1, a typical knowledge acquisition consists of four steps:

- *Step 1:* the user initiates a search/query in a nonprocedural specification format;
- *Step 2:* the original query is translated into a NeuMORE concept-based RDF query;
- *Step 3:* RDF triples are retrieved from the RDF database;
- *Step 4:* original RDF triples are clustered and returned to the user via a friendly graphical user interface (GUI).

An example RDF-based query is shown as follows, where: (1) "NeuMORE" refers to the global schema defined for our RDF database; (2) "comMeasures" is a concept in the NeuMORE ontology; (3) "comVelocity," "comDisplacement," "emgSignal," and "tdMeasure" are four properties of the concept "comMeasures," corresponding to the velocity measure of whole body center of mass (COM), the displacement measure of whole body COM, the electromyography (EMG) signal amplitudes for subjects of interest, and the temporal-distance measures for subjects of interest, respectively.

```
SELECT    DISTINCT NeuMORE: comVelocity, NeuMORE: comDisplacement
FROM      NeuMORE: comMeasures
WHERE     NeuMORE: comMeasures_emgSignal = "XXX"
AND       NeuMORE: comMeasures_tdMeasure = "YYY"
USING     NAMESPACE
NeuMORE = <http://neumore.cis.usouthal.edu/ontology/NeuMORE.owl>.
```

Notice that it is beyond the expertise of system users (PR researchers and clinicians) for them to compose such queries. Instead, the NeuMORE system will provide a friendly graphical user interface (GUI) to guide users through the procedure in creating complex queries. In addition, we propose to integrate OWL inference engines into the SEDM techniques mentioned in *Phase Three*, when interacting with the RDF database. Such a unique combination will enable us to take advantage of emerging Semantic Web technologies for knowledge acquisition. Furthermore, we aim to provide users with a single, *unified* interface that takes their requests in a nonprocedural specification format. While it is possible for existing data sources to have heterogeneous semantics, the NeuMORE system presents users a *uniform* view of information, along with newly obtained (deduced) knowledge via data mining enhanced by the automated inference, whenever possible.

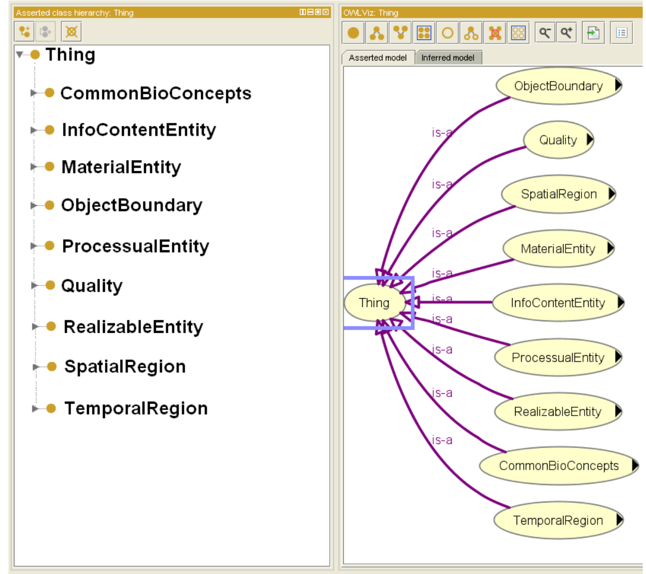

Fig. 2. Top-Level NeuMORE Concepts

IV. The NeuMORE Ontology

A. Ontology Schema

An initial version of the NeuMORE ontology was designed using Protégé 4.0, having 351 concepts in total, along with 29 object properties and 16 data properties.

Figure 2 is a screen shot from the Protégé GUI, demonstrating NeuMORE top-level concepts, which are also listed in Table I in alphabetical order. Also listed in the table are human-readable explanations for each top-level concept, along with some representative second-level concepts, i.e., direct children of top-level concepts.

B. Software Implementation

1) Class Architecture: The back end of the NeuMORE system is implemented in the C# language, following an object-oriented approach. An overall class diagram is depicted in Figure 3, whereas the detailed design of classes *OWLClass* and *OWLOntology* are shown in Figure 4. Relationships such as *superClassOf*, *subClassOf*, *superPropertyOf*, and *subPropertyOf* are placed within an inheritance structure under the common base class *OWLRelation*, which is a generic relationship between two types that specifies both the domain and the range of that relationship. For example, we use OWLRelation<OWLClass, OWLClass> to denote a *subClassOf* relationship between two concepts (with the first

TABLE I
TOP-LEVEL NEUMORE CONCEPTS AND PART OF THEIR DIRECT CHILDREN

Top-Level Concept	**Annotation (Human Readable Definition)**	**Representative Direct Children**
CommonBioConcepts	*A collection of concepts that are common in Medical and Biological Area*	SignsOrSymptoms, Tissue, Treatment
InfoContentEntity	*An entity dependent on some artifact and stands in relation of aboutness to some entity*	DataSet, Measurement, PlanSpecification
MaterialEntity	*An independent continuant whose identity is independent of that of other entities*	FiatObjectPart, Object, ObjectAggregation
ObjectBoundary	*A lower dimensional part of a spatial entity, normally a closed two-dimensional surface*	ScalpSurface
ProcessualEntity	*An occurrent that has temporal parts and involves and depends on some entity*	Process, ProcessAggregation, ProcessBoundary
Quality	*A specifically dependent continuant that is exhibited as a categorical property*	ClinicalCharacteristics, MovementMeasure
RealizableEntity	*A dependent continuant whose exhibition occurs under certain circumstances*	Disposition, Function, Plan
SpatialRegion	*A continuant that neither is bearer of quality entities nor inheres in any other entities*	AnatomyPoint, AnatomyAxis, AnatomySurface
TemporalRegion	*An occurrent (sometimes known as perdurant) that is part of time*	TemporalInstant, TemporalInterval

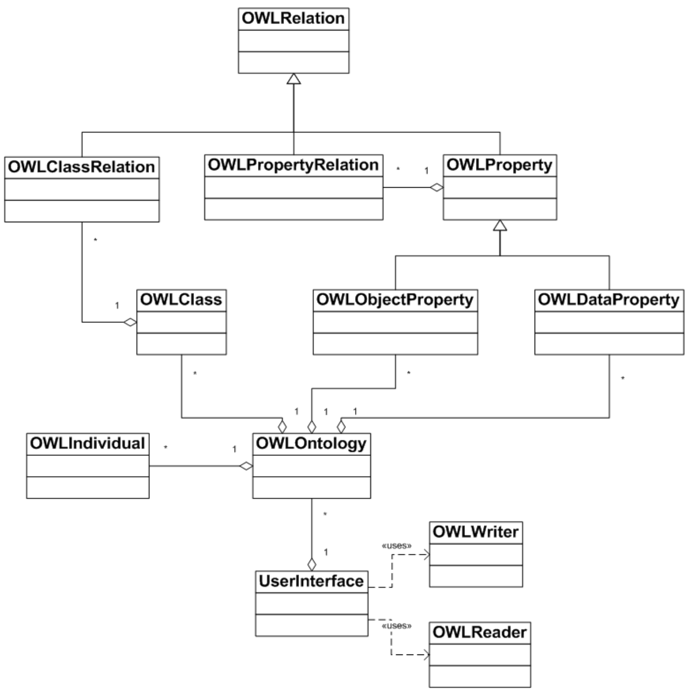

Fig. 3. NeuMORE Software Class Diagram

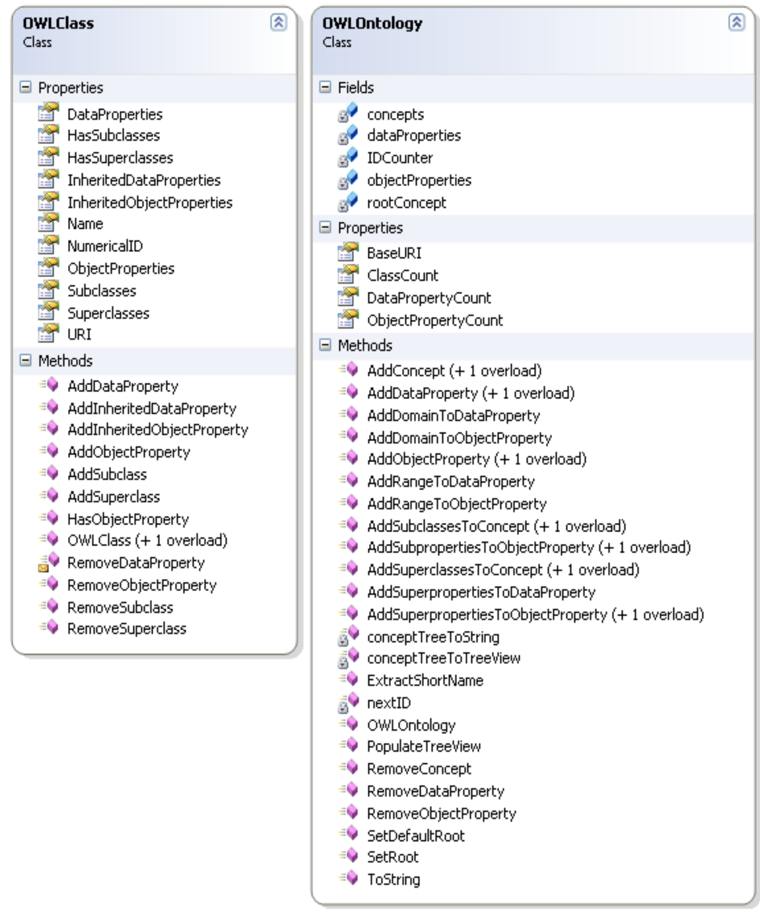

Fig. 4. Detailed Design of Classes *OWLClass* and *OWLOntology*

concept as the domain and the second one as the range), along with the methods needed to modify these structures. The inheritance structure, together with common subclasses such as *OWLClassRelation* and *OWLPropertyRelation* greatly simplifies the implementation of increasingly complex operations to be performed on the ontology. We represent each OWL entity of interest, i.e., concepts, object properties, and data properties, as its own class. Our *OWLOntology* class constitutes the external interface to this data structure, and it stores each entity in a private hash table for quick lookup by names. In this manner, the *OWLOntology* class encapsulates all operations on its members, and it allows any operation to be efficiently performed on an entity by its name.

2) Parsing the Ontology File: We have developed our own parser for OWL ontology files based on the built-in XML parsing capabilities in C#. Using the System.XML library, our *OWLReader* class is able to move from element to element in the OWL document. It checks the name of each element against the available OWL schema features and processes its attributes accordingly. Take the following lines for example:

<owl:Class rdf:about="#MuscleTissue">
<rdfs:subClassOf rdf:resource="#AdultTissue"/>
</owl:Class>

The parser recognizes the name of the element, "owl:Class." It then obtains the actual name of the class (*MuscleTissue* in this example) from the "rdf:about" attribute and holds that information as the current state. It proceeds to process each element with the name "rdfs:subClassOf," adding the current class as a subclass to each of the classes that are listed.

V. Conclusions and Future Work

We propose an innovative computing framework based on a domain-specific ontology, NeuMORE, to handle the challenge of predicting stroke survivors' improvements in extremity motor function. The NeuMORE framework, created upon *the very first* ontology for stroke recovery, is specifically designed for the physical rehabilitation (PR) domain. It will assist PR researchers and clinicians to better predict stroke survivors' neuromotor functional recovery after stroke. As a result, it will help physical therapists in making sound decisions when they design optimal intervention therapy plans for individual stroke survivors. The research goal will be accomplished via facilitated knowledge discovery, sharing, and reuse from existing data sources, through formal semantic definition and the subsequent logic-based inference that is integrated into traditional data mining techniques. In this work-in-progress paper, we first discuss proposed approaches and anticipated challenges in the NeuMORE framework; then our efforts have focused on the first phase: the development of a domain ontology. We adopt an iterative, top-down, knowledge-driven approach when designing the NeuMORE ontology. The methodologies by which we develop the ontology, along with the software implementation, are described in detail. Besides, the NeuMORE has been submitted to NCBO BioPortal (http://bioportal.bioontology.org/ontologies/44245), and a Web portal has been deployed at http://neumore.cis.usouthal.edu, which contains the up-to-date progress of the project.

Future investigation is envisioned as:

- We will continue our efforts to fine tune the NeuMORE ontology, and we plan to combine a bottom-up, data-driven approach with the current top-down approach.
- We will tag existing data sources with NeuMORE concepts. Both data schemas and data sets will be involved in this *deep* semantic annotation. The results will be specified in the RIF-PRD format.
- We will translate NeuMORE-annotated data sources into a RDF database.
- Finally, a single, unified interface, supported through data mining techniques enhanced by OWL inference engines, will be provided to PR researchers and clinicians, whose requests will be taken in a nonprocedural specification format. A uniform view of original data, along with newly deduced knowledge, will then be presented to PR researchers and clinicians for them to realize an automated knowledge acquisition.

References

[1] Basic Formal Ontology. *http://www.ifomis.org/bfo/*, November 2010.

[2] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, 2001.

[3] W. Bug, G. Ascoli, J. Grethe, A. Gupta, C. Fennema-Notestine, A. Laird, S. Larson, D. Rubin, G. Shepherd, J. Turner, and M. Martone. The NIFSTD and BIRNLex vocabularies: building comprehensive ontologies for neuroscience. *Neuroinformatics*, 6:175–194, 2008.

[4] C. Burgar, P. Lum, P. Shor, and M. der Loos HF. Development of robots for rehabilitation therapy: the Palo Alto VA/Stanford experience. *Journal of Rehabilitation Research and Development*, 37(6):663–673, 2000.

[5] CmapTools. *http://www.ihmc.us/cmaptools.html*, November 2010.

[6] D. Dou, G. Frishkoff, J. Rong, R. Frank, A. Malony, and D. Tuckeru. Development of neuroElectroMagnetic ontologies (NEMO): a framework for mining brainwave ontologies. In *Proc. the 13th ACM International Conference on Knowledge Discovery and Data Mining (KDD 07)*, pages 270–279, 2007.

[7] P. Duncan, L. Goldstein, D. Matchar, G. Divine, and J. Feussner. Measurement of motor recovery after stroke: outcome assessment and sample size requirements. *Stroke*, 23:1084–1089, 1992.

[8] FaCT++ Reasoner. *http://owl.man.ac.uk/factplusplus/*, November 2010.

[9] G. Frishkoff, P. LePendu, R. Frank, H. Liu, and D. Dou. Development of neural electromagnetic ontologies (NEMO): ontology-based tools for representation and integration of event-related brain potentials. In *Proc. the International Conference on Biomedical Ontology (ICBO 09)*, pages 31–34, 2009.

[10] Gene Ontology. *http://www.geneontology.org/index.shtml*, October 2010.

[11] C. Gowland, H. deBruin, J. Basmajian, N. Plews, and I. Burcea. Agonist and antagonist activity during voluntary upper-limb movement in patients with stroke. *Physical Therapy*, 72(9):624–633, 1992.

[12] G. Gresham, T. Fitzpatrick, P. Wolf, P. McNamara, W. Kannel, and T. Dawber. Residual disability in survivors of stroke - the Framingham study. *The New England Journal of Medicine*, 293(19):954–956, November 1975.

[13] S. Handschuh, S. Staab, and R. Volz. On deep annotation. In *Proc. the 12th International World Wide Web Conference (WWW 03)*, May 2003.

[14] S. Handschuh, R. Volz, and S. Staab. Annotation for the deep Web. *IEEE Intelligent Systems*, 18(5):42–48, 2003.

[15] J. Huang, M. Tan, D. Dou, L. He, C. Townsend, and P. Hayes. Ontology for microRNA target prediction in human cancer. In *Proc. the First ACM International Conference on Bioinformatics and Computational Biology (ACM-BCB 10)*, August 2010.

[16] KIF. *http://logic.stanford.edu/kif/*, November 2010.

[17] G. Kwakkel, B. Kollen, J. van der Grond, and A. Prevo. Probability of regaining dexterity in the flaccid upper limb: impact of severity of paresis and time since onset in acute stroke. *Stroke*, 34:2181–2186, 2003.

[18] A. Luft, S. McCombe-Waller, J. Whitall, L. Forrester, R. Macko, J. Sorkin, J. Schulz, A. Goldberg, and D. Hanley. Repetitive bilateral arm training and motor cortex activation in chronic stroke: a randomized controlled trial. *Jama*, 292:1853–1861, 2004.

[19] NCBO. *http://www.bioontology.org/*, November 2010.

[20] N. Noy and D. Mcguinness. Ontology development 101: a guide to creating your first Ontology. 2000.

[21] OBO. *http://www.obofoundry.org/*, November 2010.

[22] OKBC. *http://www.ai.sri.com/ okbc/*, November 2010.

[23] Ontoprise. *http://www.ontoprise.de/index.php?id=135*, November 2010.

[24] OWL. *http://www.w3.org/TR/owl-features/*, November 2010.

[25] Pellet Reasoner. *http://clarkparsia.com/pellet/*, November 2010.

[26] Protégé. *http://protege.stanford.edu/*, November 2010.

[27] RDF. *http://www.w3.org/RDF/*, November 2010.

[28] RIF. *http://www.w3.org/TR/2010/REC-rif-prd-20100622/*, October 2010.

[29] M. Singh and M. Huhns. *Service-Oriented Computing - Semantics, Processes, Agents*. Wiley, Chichester, England, first edition, 2005.

[30] B. Smith, M. Ashburner, C. Rosse, J. Bard, W. Bug, W. Ceusters, L. Goldberg, K. Eilbeck, A. Ireland, C. Mungall, N. Leontis, P. Rocca-Serra, A. Ruttenberg, S. Sansone, R. Scheuermann, N. Shah, P. Whetzel, and S. Lewis. The OBO foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology*, 25(11):1251–1255, 2007.

[31] SPARQL. *http://www.w3.org/TR/rdf-sparql-query/*, November 2010.

[32] T. Thomas et al. Heart Disease and Stroke Statistics - 2006 Update. *Circulation*, 113:e85–e151, 2006.

[33] E. Taub, J. Crago, L. Burgio, T. Groomes, E. Cook, C. DeLuca, and N. Miller. An operant approach to rehabilitation medicine: overcoming learned nonuse by shaping. *The Journal of the Experimental Analysis of Behavior*, 61:281–293, 1994.

[34] UMLS. *http://www.nlm.nih.gov/research/umls/*, November 2010.

[35] W. Rosamond et al. Heart Disease and Stroke Statistics - 2007 Update. *Circulation*, 115:e69–e171, 2007.

[36] W3C. *http://www.w3.org/*, November 2010.

[37] J. Whitall, S. Waller, K. Silver, and R. Macko. Repetitive bilateral arm training with rhythmic auditory cueing improves motor function in chronic hemiparetic stroke. *Stroke*, 31(10):2390–2395, 2000.

[38] G. Williams. Incidence and characteristics of total stroke in the United States. *BMC Neurology*, 1(2), 2001.