Comparison of Semantic Similarity Measures for Application Specific Ontology Pruning

Wei-Nchih Lee[1,2], Will Bridewell[1], and Amar K. Das[1]
[1] Stanford Center for Biomedical Informatics Research, Stanford, CA
[2] VA Palo Alto Health Care System, Palo Alto, CA
email: wemt4yl@stanford.edu

***Abstract*—Comparing the effects of one drug to another drug, based on their similarity, is important in clinical research. Ontology-derived measures of drug-drug similarity may help to automate such analyses on large data sets. However, general drug ontologies can contain hierarchical distinctions that are irrelevant to a particular clinical application and thus may lead to inaccurate semantic similarity measures. We propose that ontology pruning be used to remove unneeded concepts so that the resulting ontology better reflects the semantic distinctions of a particular domain. In this paper, we present a novel pruning strategy for drug ontologies. For three clinical domains, we derive previously developed semantic similarity measures for the automatically pruned ontology and the full drug ontology against those for the expert derived ontology. We show that the values of similarity measures based on our pruned approach are closer to those of the expert derived ontology than to those of the full ontology. Our pruning approach thus provides a standardized domain-specific measure of drug-drug similarity for clinical applications.**

Keywords - Semantic similarity;Ontology Pruning;Clinical Informatics

I. INTRODUCTION

A. Background

As researchers in the biomedical sciences gather increasing amounts of digital health information, they need computational tools that can intelligently manage the large volume and complexity of the data. Ontologies—computational models of classification knowledge—can play an important role in the automated interpretation of large data sets [1]. In particular, ontologies can allow the automated determination of *semantic similarity* of two concepts, by measuring the likeness of one concept to another based on the closeness of those concepts within an ontology hierarchy. For example, the drugs *enalapril* and *captopril* share a common class in the mechanism of action of *angiotensin converting enzyme inhibitors*, and are closer together than *enalapril* and *metoprolol*, because *metoprolol's* mechanism of action is in the class of *adrenergic receptor blockade*.

Many approaches to measuring semantic similarity from ontologies have been developed. These can broadly be classified as graph based, term frequency based, or a combination of the two [2]. Graph based methods, such as by Leacock and Chodorow [3], or Wu and Palmer [4], use some variation of a shortest path algorithm, and have the advantage not being dependent on the availability of text corpora. These approaches have commonly been applied to the WordNet ontology, though recently similarity measures have been used and developed with standard biomedical ontologies such as the Gene Ontology [5] and SNOMED-CT [6]. A benefit of using standard biomedical ontologies is that similarity measures between concepts will be the same regardless of the application for which it is intended.

Biomedical ontologies, however, can be much larger in size than WordNet, which contains about 200,000 concepts. SNOMED-CT, by contrast, contains more than a million concepts. This difference in size imposes an adverse performance effect on the computation of any semantic similarity measure. Software applications that require fast results would be hampered by this constraint. One solution to this problem is to pre-compute pairwise concept similarities in a single matrix, which could then be used as a look up table. However, an ontology like SNOMED-CT could result in an unreasonably large matrix. In addition, hierarchical distinctions found in the context of an entire drug ontology may not be relevant for a specific clinical application. For example, in the National Drug File ontology, the concept *Protease Inhibitors* contains the child concept *HIV Protease Inhibitors*. While this distinction may be useful given other protease inhibitors such as *Angiotensin Converting Enzyme Inhibitors*, an application designed for the HIV clinical domain would find it unnecessary. Since semantic similarity measures often incorporate a shortest path algorithm with weighted edges, the computed similarity measures might not reflect the needs of the application designer. We propose to develop a pruning strategy that will provide the most relevant similarity scores for an application, and this topic is an open research question.

The work in this paper grew out of a need for a drug-drug similarity matrix to support our research on temporal sequence alignment algorithms applied to clinical drug treatment databases [7]. The drug-drug similarity matrix, which would serve the purpose of scoring similar drug regimens in a treatment history, needed to satisfy three conditions. First, the matrix needed to be derived from a standard ontology. Second, the pruning strategy needed to be deterministic, so that for a given set of concepts within an ontology, the same similarity scoring matrix would always be derived. Lastly, the pruning strategy needed to have an ontology structure similar to what an expert might have derived for a particular application.

We present an automated pruning strategy of available large ontologies to create an domain specific ontology suitable for measuring semantic similarity. Using three different clinical domains, we present an evaluation of how well existing semantic similarity measures for a reference ontology correlate with those for our automatically pruned ontology and those for a full drug ontology.

978-0-7695-4407-6/11 $26.00 © 2011 IEEE
DOI 10.1109/HISB.2011.17

B. Related Work

Relevant work is published in ontology views and ontology pruning. An ontology view is a subset of the concepts for a full ontology—the specific subset is chosen based on the needs of the application developer. They are analogous to views in relational databases, with the difference that an ontology view includes concept properties and their antecedents, and that it is easily represented as a hierarchical graph structure. Ontology views have been implemented with query languages and graph-based traversals of the ontology structure [8,9]. In either case, however, the sub-ontology that is extracted in the view maintains the graph structure it had in the full ontology. Therefore, the problem of having hierarchical levels that are irrelevant to the application domain, such as with the *Protease Inhibitor* example described previously, still exists.

Ontology pruning methods remove concepts irrelevant to the application domain from the full ontology. A number of ontology pruning methods exist, with varying degrees of manual selection of concepts and automatic pruning [10]. Conesa recently described an automated approach that requires an initial selection of concepts that then outputs a minimum yet complete set of relevant concepts from the original full ontology [11]. Conesa's approach adds to the ontology view literature by including a method for pruning unnecessary parent concepts, thus potentially removing un-needed distinctions in an ontology hierarchy. However, pruning methods still assume that the user is aware of the relevant upper level ontology concepts in the initial selection of pruning parameters. In the LATCH example, by contrast, we are only interested in the similarity among drugs, which would presumably be leaf nodes in any ontology structure we use. Knowledge of upper level concepts is not essential. Also, Conesa's approach can result in unconnected ontology graphs. Semantic similarity measures that rely on shortest path distances, and therefore assume fully connected graphs, would not function with these methods.

C. Study Aims and Hypothesis

Using a standard ontology from the National Center for Bio-Ontologies (NCBO) [12], we present an algorithmic approach to pruning the ontology hierarchy for three different clinical domains – hypertension, congestive heart failure, and HIV infection. We evaluate our pruning method on three published, graph-based, ontology semantic similarity measures. With each similarity measure, we construct drug-drug similarity matrices for the full and pruned ontologies. We then apply the measures to expert derived, publicly available ontologies in each clinical domain. We hypothesize that for three different clinical domains, our ontology pruning method will provide drug-drug semantic similarity scores with stronger correlations to an expert derived ontology than a full, non-pruned ontology would.

II. METHODS

A. Term Selection and Reference Standard Ontologies

Table 1 shows the drug terms that we used for congestive heart failure, hypertension, and HIV. The drugs share a moderate overlap, since many anti-hypertensive medications are also useful for congestive heart failure. Also, the terms are selected from publicly available, expert derived and application specific ontologies for ease of evaluation. We used expert derived ontologies as reference standards for congestive heart failure HEARTFAID [13] and hypertension (from Mabotuwana) [14]. In each, we manually selected the drug ontologies such that the terms of interest were all leaf nodes in the hierarchy. Branches of the hierarchy that did not contain a term of interest were not included in the reference ontology. For the HIV domain, we built the ontology based on expert information from the Stanford HIV Database regarding appropriate drug classes and the respective drug instances [15].

B. Initial Extraction of Ontology View

For this paper, we used the National Drug File, a publicly available resource from the Veterans Administration Health Care System that is available via the NCBO in the Web Ontology Language [16]. We first had to make a decision on the appropriate ontology view from which to begin the pruning process. Concepts for drug terms in the NDF are classified based on an alphabetical ordering, which does not contain any sort of useful semantics. On the other hand, there is the concept class *Drug Products by VA Class*, which organizes drugs based on generally accepted drug classes in the medical community. However, this concept class uses the drug formulations themselves as instances – there are no concept designations for the drug itself. For example, the child concepts of *ACE Inhibitors* under *Drug Products by VA Class* include *Benazepril HCL 10mg tab, Benazepril HCL 20mg tab,* and so on. However, there is no concept class for

TABLE 1. Drug terms used for the clinical domains

Congestive heart failure	Hypertension	HIV
Enalapril	Quinapril	Abacavir
Lisinopril	Captopril	Didanosine
Carvedilol	Enalapril	Emtricitabine
Metoprolol	Candesartan	Lamivudine
Spironolactone	Losartan	Stavudine
Candesartan	Clonidine	Tenofovir
Losartan	Diazoxide	Zidovudine
Amlodipine	Prazosin	Delavirdie
Felodipine	Atenolol	Efavirenz
Digoxin	Metoprolol	Nevirapine
Furosemide	Carvedilol	Atazanavir
Metolazone	Amlodipine	Darunavir
Amiloride	Nifedipine	Amprenavir
Triamterene	Felodipine	Indinavir
Hydrochlorothiazide	Bumetanide	Lopinavir
Indapamide	Amiloride	Nelfinavir
Clopidogrel	Triamterene	Saquinavir
Warfarin	Hydrochlorothiazide	Tipranavir
Amrinone	Bendroflumethiazide	Ritonavir
Dobutamine	Diltiazem	Enfuvirtide
Milrinone	Verapamil	

the drug *Benazepril*. For our purposes, we would have had to add the *Benazepril* concept class to the ontology, or calculate similarity measures for all the different formulations of benazapril (because we are seeking to do an automated approach). Since adding concepts goes against our goals of pruning an ontology, we instead looked for an ontology view that most closely reflected our needs in a drug similarity matrix.

For this reason, we used the *has_mechanism_of_action* property to infer an *is_a* relationship for the drug term and its parent class. For example, the relationship,

has_mechanism_of_action (ENALAPRIL, ACE-I)

is used to infer the relationship

is_a (ENALAPRIL, ACE-I)

The parent classes in this ontology view reflect the common medical practice of classifying drugs by their mechanism of action.

With this approach, we determined the least common ancestor (LCA) for the parents of the drug terms of interest. Using the LCA as the root of the hierarchy, we reconstructed a sub-ontology from the NDF as the initial view in which all the leaf nodes refer to the drug terms of interest.

C. Pruning Algorithm

We start with the assumption that parent nodes in the ontology contain information based on their number of child edges. Therefore, a node with 2 or more child edges implicitly states that each child node can be differentiated from the other. For example, the concept class *Adrenergic Beta Antagonist* contains two child concepts – *Adrenergic beta1* and *Adrenergic beta2 Antagonists* – each of which is relevant to hypertension. We presume that the hierarchical distinction here serves a useful purpose as each child branch leads to a hypertension drug of interest. The corollary is that a parent node that has only one child edge offers no implicit information that differentiates it from its single child. Thus, if *Adrenergic beta2* did not include in its path a drug of interest, it would be pruned initially as an extraneous concept, leaving *Adrenergic beta1* as the only child of *Adrenergic Beta Antagonist.* In this case, *Adrenergic Beta Antagonist* can be pruned, as the added hierarchical distinction is unnecessary. In this way, parent nodes with one child can be removed from the ontology.

Our approach, then, has three steps. The first is to remove extraneous concepts to the application. To accomplish this task, we simply removed any branch of the hierarchy that did not contain a drug concept of interest in its path. Secondly, we determined the least common ancestor (LCA) for all the drug concepts of interest in the NDF, and constructed a sub-ontology view with the LCA as the root node of the hierarchy. This assured us of a fully connected ontology graph against which we could apply semantic similarity measures. Third, we used a bottom up recursive algorithm which we call OntoPrune, starting at the leaf nodes of the sub-ontology, and then removing any parent node with only 1 child node. This was continued until no further parent nodes could be removed, resulting in the final pruned ontology hierarchy. Fig. 1 shows the pseudo-code for the pruning algorithm, OntoPrune. The complexity of the algorithm is driven by the nested for-loop when a parent node is pruned and the child concepts for the parent are shifted up, resulting in a complexity of $O(n^2)$, where n is the number of drug terms of interest.

D. Evaluation of Pruning Algorithm

For each clinical domain, we applied three semantic similarity measures to the full, pruned, and expert derived ontology hierarchies. The measures we chose are by Wu and Palmer (WP) [4], Leacock and Chodorow (LC) [3], and Al-Mubaid (AM) [17]. Each uses some variation of the shortest path between concepts to determine the similarity measure. Their approaches differ in the method chosen to scale the shortest path distance.

WP defined the similarity measure as:

$$sim_{WP}(c_1, c_2) = \max\left(\frac{2 \times depth[LCA(c_1, c_2)]}{len(c_1, c_2) + 2 \times depth[LCA(c_1, c_2)]} \right) \quad (1)$$

where the LCA is the least common ancestor between two concepts, c_1 and c_2. In this case, the shortest path distance is scaled by the larger of the two concept depths. As a result, the calculated similarity score depends only on the relative positions of the concept pairs, and not on the complexity of the ontology from which it came. LC, by contrast, scaled the shortest path distance by the entire depth of the ontology hierarchy:

$$sim_{LC}(c_1, c_2) = \max\left(-\log \frac{ShortestPath(c_1, c_2)}{2 \times Taxonomy\ Depth} \right) \quad (2)$$

An ontology with a greater number of hierarchical levels, then, would reduce the calculated distance between two concepts because of the presence of more concepts in the knowledge base. However, this approach scales all the

```
OntoPrune(Tree, Parents)
  If Parents empty, return
  Else
    New ParentList
    For par in Parent
       If par has ≥ 2 children
         Add parent of par to ParentList
       Else
         Add parent of par to ParentList
         Remove par from Tree
         Shift children of par up
    OntoPrune(Tree,ParentList)
```

Figure 1 Pseudo-code for the Ontology Pruning Algorithm. Terms of interest are assumed to be the leaf nodes of Tree

distances equally, without concern for the relative positions of the concept pairs within the ontology. To address this deficiency, AM introduced a cluster-based approach to semantic similarity, measured as:

$$sim_{AM}(c_1, c_2) = \log[(Path(c_1, c_2))^{\alpha} \times (CSpec)^{\beta} + k] \quad (3)$$

where CSpec, the common specificity, denotes the position of c_1 and c_2 relative to the root of the ontology hierarchy. AM then went on to describe four distinct patterns for the relative locations of any concept pair within an ontology, and the definition for CSpec in each pattern. The parameters α and β describe the contributions of path length and hierarchy position respectively while k is a constant ≥ 1. For the purposes of this work, we set α, β and k to 1 as described by AM.

With the full, pruned, and expert derived ontologies in each clinical domain, we then constructed a symmetric, drug-drug similarity matrix using each of the similarity measures that we described above. To standardize the scores for analysis, we scaled all of the similarity scores to between 0 and 1 with an exponential function, such that 0 implies no similarity between drug pairs, and 1 signifies an exact match in drug pairs [18]. To simplify analysis, we selected from the reference hierarchies one pair of drugs such that they were 1) siblings of the same class, 2) in the same hierarchical level, and 3) from different hierarchical levels (diff = 1). Statistical tests for comparisons among the similarity matrices were done with Spearman's rank correlation, a non-parametric rank-ordering test, in which a coefficient closer to 1 or (-1) implies a stronger correlation between matrices, and a value close to 0 implies no correlation [19].

III. Results

For illustrative purposes and in the interest of space, we show only the hierarchical structures for hypertension, but will report on the analysis for all three clinical domains. Fig. 2 shows the full ontology hierarchy for the hypertensive medications from the NDF. As seen in the figure, drugs are at different hierarchical levels of varying depths throughout the ontology. Also, we can see that only branches that end with drug terms of interest are included in the full hierarchy – the other branches have already been excluded for the purposes of analysis.

Fig. 3 shows the hierarchy that resulted from the pruning algorithm, along with the expert derived reference as comparison. As is true with congestive heart failure and HIV, the expert derived ontologies had fewer hierarchical levels than in the full ontology. The pruning algorithm reduced the number of hierarchy levels, and resulted in a flattening of the ontology, that on visual inspection makes it similar to the expert derived reference. Table 2 shows the values for the three semantic similarity measures for siblings, same hierarchy, and different hierarchy levels for each clinical domain. In general, the pruned ontology appears to have a similarity score that is closer to the reference ontology

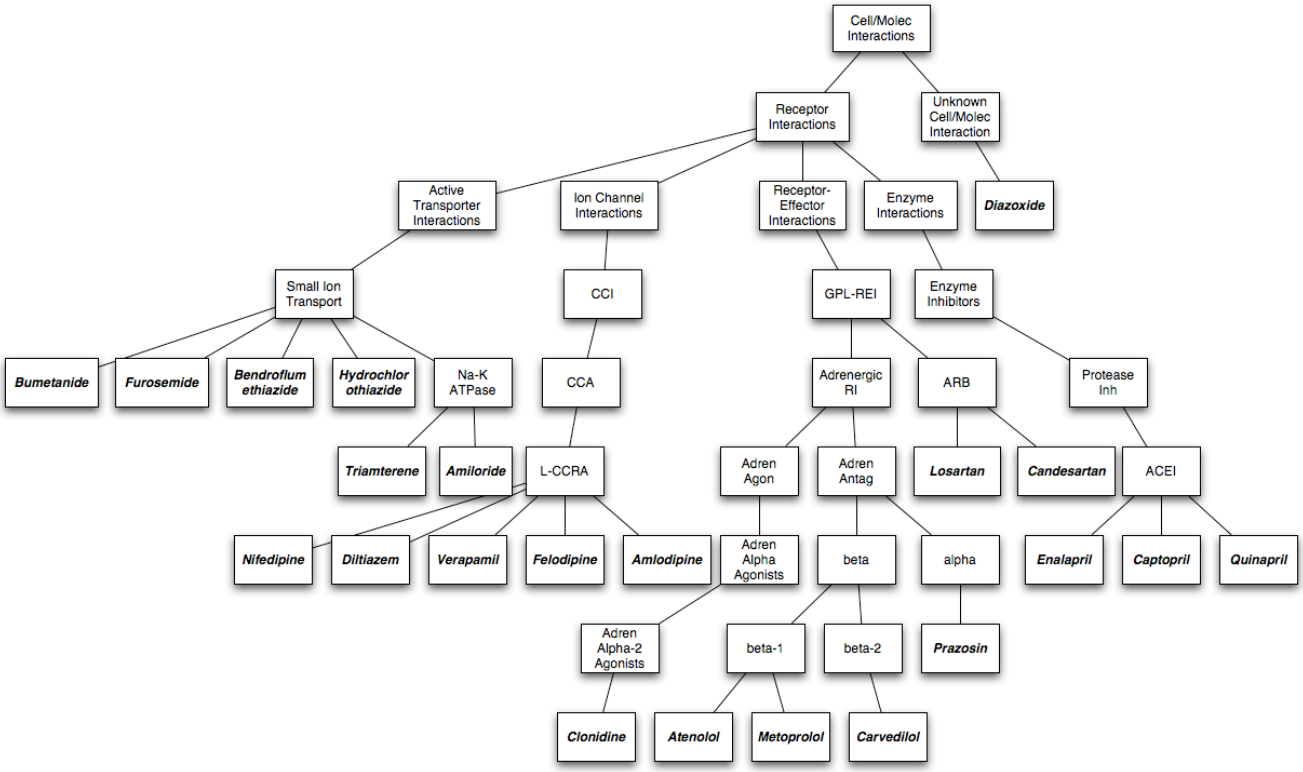

Figure 2 Full hierarchy for hypertension from the NDF ontology

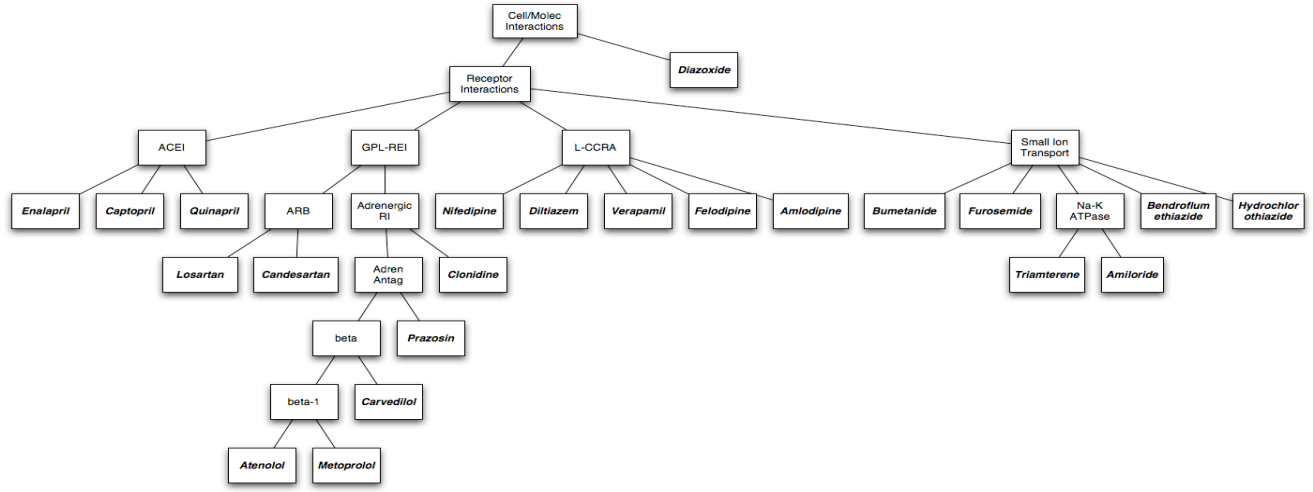

Pruned Ontology

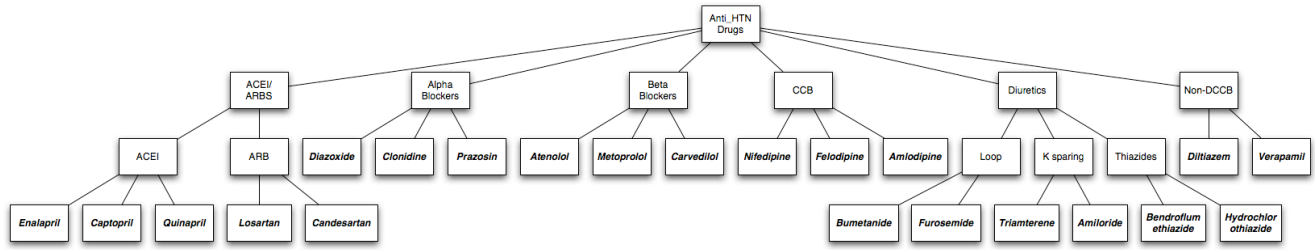

Reference Ontology

Figure 3 Pruned Ontology versus Expert Derived Reference Standard for Hypertension

as opposed to the full ontology when comparing sibling drugs and drugs at the same hierarchical level. The performance of the similarity measures when comparing drugs at different hierarchical levels appears more variable.

Fig. 4 shows the Spearman correlations for the different semantic similarity measures in each clinical domain. In each domain, the correlations for the full versus reference ontology, and the pruned versus reference ontology, are shown. As seen in the figure, the pruned ontology generally has a stronger correlation with the expert derived reference than with the full ontology. The one exception is with the hypertension ontology when using the measure by AM. The correlations between the pruned and reference ontologies are all statistically significant: The p-value for the LC measure in congestive heart failure is 0.017, and for the WP measure in hypertension 0.002. Otherwise, all other p-values < 0.001.

IV. Discussion

In this work, we took a large publicly available ontology, the NDF, and developed a pruning strategy that results in sub-ontologies that are tailored to particular application needs of deriving drug-drug semantic similarity. Specifically, we needed a fully connected ontology graph that would allow us to utilize shortest path algorithms in the computation of these measures. Prior work does not meet these requirements. Some methods will maintain portions of the hierarchy that are not needed for shortest distance. Other methods, like Conesa, may result in unconnected graphs.

Also, our approach can be automated without user input. In prior work, users must choose higher-level concepts in the concepts to prevent them from being pruned; in other words, the user must have some prior insight of the ontological *commitments* for the application. In our work, however, the user is only interested in the *leaf* terms of the hierarchy—that is, the application specific drug terms. The specific parent concepts to the drug terms are less important than the overall *structure* of the ontology.

We demonstrated the robustness of our pruning approach along two dimensions. First, we evaluated the pruning algorithm using three different semantic similarity measures. The measures that we selected are published, and established graph-based methods. For each of the selected similarity measures, we found that the drug similarity matrix for the pruned ontology more closely correlated with the expert derived references than with the full ontology. With the LC and WP measures, in particular, the pruned ontologies were substantially more correlated in the domains of hypertension and congestive heart failure. The one exception was with the AM measure in the hypertension domain. This may be because that measure takes into account the position of concept pairs within the ontology, thereby reducing the effect of multiple hierarchical levels in an ontology. Even

TABLE 2. Semantic Similarity Measures in Clinical Domains

	Leacock			Wu			Al-Mubaid		
	Full	*Pruned*	*Gold*	*Full*	*Pruned*	*Gold*	*Full*	*Pruned*	*Gold*
Hypertension									
Siblings	0.62	0.60	0.47	0.80	0.67	0.67	0.45	0.38	0.64
Same Hierarchy	0.11	0.21	0.22	0.24	0.31	0.29	0.15	0.11	0.35
Different Hierarchy	0.11	0.21	0.14	0.24	0.31	0.25	0.15	0.11	0.32
Congestive Heart Failure									
Siblings	0.42	0.52	0.47	0.71	0.72	0.57	0.41	0.64	0.52
Same Hierarchy	0.21	0.15	0.23	0.18	0.22	0.29	0.19	0.27	0.35
Different Hierarchy	0.21	0.30	0.14	0.18	0.29	0.25	0.15	0.32	0.32
HIV									
Siblings	0.58	0.52	0.47	0.80	0.72	0.67	0.64	0.64	0.64
Same Hierarchy	0.09	0.30	0.27	0.15	0.29	0.29	0.19	0.15	0.35
Different Hierarchy	0.21	0.22	0.14	0.43	0.40	0.25	0.24	0.29	0.28

then, the correlation scores between pruned and full ontologies with the references were close to each other.

Second, we evaluated the pruning algorithm for robustness along different clinical domains. As seen in Fig. 4, the pruned ontologies more closely correlated with the reference ontology in each clinical domain. Our data suggests that the pruned ontology performs best when evaluating drugs at the same hierarchical level in the ontology. Because our pruning algorithm removes extraneous levels in an ontology hierarchy, we will obtain a sub-ontology that is flatter, with fewer hierarchical levels—consistent with the ontology structure that was developed by experts in each of the clinical domains we evaluated.

There are limitations to this study. We used expert derived ontologies as a reference standard, instead of obtaining similarity measures for drugs from a sample of domain experts. As such, we are evaluating the expected similarity measures from domain experts, rather than the actual measures from the experts directly. By using published expert derived ontologies, however, we more closely mimic the process of ontology development that a developer may go through with a specific application. Another limitation is that we only evaluated graph-based semantic similarity measures, rather than those that used term-frequency or a hybrid term-frequency/graph-based approach. We chose, however, not to evaluate these approaches for two main reasons. First, we felt that many application developers may not have access to the large text corpora necessary for those methods to be implemented. Second, while standard ontologies are available in ontology repositories, standard text corpora, particularly with respect to the biomedical domain, do not exist.

In conclusion, we developed a pruning algorithm that can be applied to large, standard ontologies for the purposes of

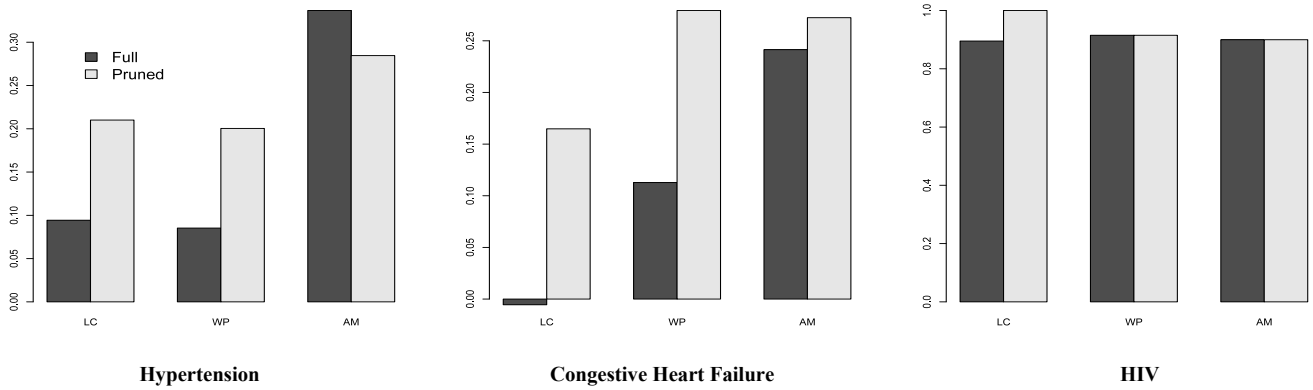

Figure 4 Correlation to reference hierarchies of the full and pruned ontologies for each clinical domain. The y-axis is the value of Spearman's correlation coefficient for the semantic similarity measures by LC, WP, and AM respectively.

deriving drug-drug similarity measures that are standard, reproducable, and do not require substantial domain expertise. We evaluated our algorithm with different semantics similarity measures in different clinical domains and have found it to be robust in both dimensions. We intend to further evaluate this ontology pruning approach to our current work with temporal sequence alignment in the clinical domain, which requires the availability of clinical similarity matrices for drug treatments.

ACKNOWLEDGMENT

WL is supported as a medical informatics fellow in the Veterans Administration Health Care System of Palo Alto, Palo Alto, California. Work by the authors WB and AD is funded by grant R01LM09607 and the Richard and Susan Levy Fund at Stanford University. Views expressed are those of the authors and not necessarily those of the Department of Veterans Affairs.

REFERENCES

[1] Chandrasekaran JRJ. What are ontologies, and why do we need them? IEEE Xplore. 2008:1-7.

[2] C Pesquita, D Faria, AO Falcao, P Lord, FM Couto, Semantic similarity in biomedical ontologies. PLoS Comput Biol 5: e1000443. (2009)

[3] C Leacock, M Chodorow, Combining local context and wordnet similarity for word sense identification. In: Fellbaum C, editor. WordNet: An Electronic Lexical Database. MIT Press; 1998. pp. 265–283.

[4] Z Wu, Palmer M, Verb semantics and lexical selection. Assoc Comput Linguist Proc. Las Cruces, NM: 1994. pp. 133–138.

[5] Othman R, Deris S, Illias R. A genetic similarity algorithm for searching the gene ontology terms and annotating anonymous protein sequences. J Biomed Inform 23:529-38.

[6] T Pedersen, SV Pakhomov, S Patwardhan, CG Chute. Measures of semantic similarity and relatedness in the biomedical domain. J Biomed Inform. 2007 Jun;40(3):288-99.

[7] WN Lee, AK Das, Local Alignment Tool for Clinical History: Temporal Semantic Search of Clinical Databases. AMIA Annu Symp Proc. 2010. ,in press.

[8] N Noy and M Musen. Specifying ontology views by traversal. In Proc. of ISWC, 2004.

[9] LT Detwiler, JF Brinkley. Querying non-materialized ontology views. AMIA Annu Symp Proc. 2007, p 934.

[10] JW Kim, J Conesa Caralt, JK Hilliard, Pruning Bio-Ontologies. 40th Annual Hawaii International Conference on System Sciences (HICSS'07), 2007, hicss, pp.196c

[11] J Conesa, A Olive: A Method for Pruning Ontologies in the Development of Conceptual Schemas of Information Systems. J. Data Semantics V: 64-90 (2006)

[12] DL Rubin, D Moreira, PP Kanjamala, MA Musen, Bioportal: A web portal to biomedical ontologies. In Symbiotic Relationships between Semantic Web and Knowledge Engineering (AAAI Spring Symposium Series). Stanford Univ. Press, 2008.

[13] D Conforti, D Costanzo, F Perticone, G Parati, K Kawecka-Jaszcz, A Marsh, et al. HEARTFAID: A knowledge based platform of services for supporting medical-clinical management of heart failure within elderly population. Stud Health Technol Inform 2006;121:108– 25.

[14] T Mabotuwana, J Warren, An ontology-based approach to enhance querying capabilities of general practice medicine for better management of hypertension. Artificial Intelligence in Medicine, 47(2): 87-103, 2009.

[15] http://hivdb.stanford.edu

[16] SH Brown, PL Elkin, ST Rosenbloom, C Husse, BA Bauer, MJ Lincoln et al., VA national drug file reference terminology: a cross-institutional content coverage Study, Medinfo 2004 (2004), pp. 477–481.

[17] H Al-Mubaid, HA Nguyen, A cluster-based approach for semantic similarity in the biomedical domain. IEEE Int Conf GR Proc. 2006:2713–2717.

[18] RN Shepard, ""Towards a Universal Law of Generalisation for Psychological Science,"" Science, vol. 237, pp.1317-1323, 1987.

[19] DG Altman, Practical Statistics for Medical Research. Chapman & Hall, London, (1991), pp. 285-288