

Measuring Similarity in Ontologies: A New Family of Measures

Tahani Alsubait, Bijan Parsia, and Uli Sattler

School of Computer Science, The University of Manchester, United Kingdom
{alsubait,bparsia,sattler}@cs.man.ac.uk

Abstract. Without a doubt, similarity measurement is important for numerous applications (e.g., information retrieval, clustering, ontology matching). Several attempts have been already made to develop similarity measures for ontologies. We noticed that some existing similarity measures are ad-hoc and unprincipled. In addition, there is still a need for similarity measures which are applicable to expressive Description Logics (i.e., beyond \mathcal{EL}) and which are terminological (i.e., do not require an *ABox*). To address these requirements, we have developed a new family of similarity measures. To date, there has been no thorough empirical investigation of similarity measures. This has motivated us to carry out two separate empirical studies. First, we compare the new measures along with some existing measures against a gold-standard. Second, we examine the practicality of using the new measures over an independently motivated corpus of ontologies (BioPortal library). In addition, we examine whether cheap measures can be an approximation of some more computationally expensive measures.

1 Introduction

The process of assigning a numerical value reflecting the degree of resemblance between two ontology concepts or the so called conceptual similarity measurement is important for numerous applications. Be it classical information retrieval, ontology matching [9], ontology learning [4] or various other applications. It is also known that similarity measurement is difficult. This can be easily seen by looking at the several attempts that have been made to develop similarity measures, see for example [6, 23, 33, 24, 15, 25, 32]. The problem is also well-founded in psychology and a number of psychological models of similarity have been already developed, see for example [8, 31, 27, 21, 30, 10, 14]. Rather than adopting a psychological model for similarity as a foundation, we noticed that some existing similarity measures for ontologies are ad-hoc and unprincipled. In addition, there is still a need for similarity measures which are applicable to expressive Description Logics (DLs) (i.e., beyond \mathcal{EL}) and which are terminological (i.e., do not require an *ABox*). To address these requirements, we have developed a new family of similarity measures which are founded on the feature-based psychological model [30]. The individual measures vary in their accuracy/computational cost based on which *features* they consider.

To date, there has been no thorough empirical investigation of similarity measures. This has motivated us to carry out two separate empirical studies.

First, we compare the new measures along with some existing measures against a gold-standard. Second, we examine the practicality of using the new measures over an independently motivated corpus of ontologies (BioPortal¹ library) which contains over 300 ontologies. In addition, we examine whether cheap measures can be an approximation of some more computationally expensive measures.

To understand the major differences between similarity measures w.r.t. the task in which they are involved in, consider, for example, the following three tasks:

- Task1: Given a concept C , retrieve all concepts D s.t. $\text{Similarity}(C, D) > 0$.
- Task2: Given a concept C , retrieve the N most similar concepts.
- Task3: Given a concept C and some threshold Δ , retrieve all concepts D s.t. $\text{Similarity}(C, D) > \Delta$.

We expect most similarity measures to behave similarly in Task 1 because we are not interested in the particular similarity values nor any particular ordering among the similar concepts. However, the Task 2 gets harder as N gets smaller. In this case, a similarity measure that underestimates the similarity of some very similar concepts and overestimates the similarity of others can fail the task. In Task 3, the actual similarity values matter. Hence, using the most accurate similarity measure is essential.

2 Preliminaries

We assume the reader to be familiar with DL ontologies. In what follows, we briefly introduce the relevant terminology. For a detailed overview, the reader is referred to [1]. The set of terms, i.e., concept, individual and role names, in an ontology \mathcal{O} is referred to as its signature, denoted $\bar{\mathcal{O}}$. Throughout the paper, we use N_C , N_R for the sets of concept and role names respectively and $\mathcal{C}_{\mathcal{L}}$ to denote a set of possibly complex concepts of a concept language $\mathcal{L}(\Sigma)$ over a signature Σ and we use the usual entailment operator \models .

3 Desired properties for similarity measures

Various psychological models for similarity have been developed (e.g., Geometric [27, 21], Transformational [19, 10] and Features [30] models). Due to the richness of ontologies, not all models can be adopted when considering conceptual similarity in ontologies. This is because many things are associated with a concept in an ontology (e.g., atomic subsumers/subsumees, complex subsumers/subsumees, instances, referencing axioms). Looking at existing approaches for measuring similarity in DL ontologies, one can notice that approaches which aim at providing a numerical value as a result of the similarity measurement process are mainly founded on feature-based models [30], although they might disagree on which features to consider.

¹ <http://biportal.bioontology.org/>

In what follows, we concentrate on feature-based notions of similarity where the degree of similarity S_{CD} between objects C, D depends on features common to C and D , unique features of C and unique features of D . Considering both common and distinguishing features is a vital property of the features model.

Looking at existing approaches for measuring similarity in ontologies, we find that some of these approaches consider common xor unique features (rather than both) and that some approaches consider features that some instances (rather than all) of the compared concepts have. To account for all the features of a concept, we need to look at all (possibly complex) entailed subsumers of that concept. To understand this issue, we present the following example:

Example 1 Consider the ontology:

$$\begin{aligned} \{ & \textit{Animal} \sqsubseteq \textit{Organism} \sqcap \exists \textit{eats}.\top, & \textit{Plant} \sqsubseteq \textit{Organism}, \\ & \textit{Carnivore} \sqsubseteq \textit{Animal} \sqcap \forall \textit{eats}.\textit{Animal}, & \textit{Herbivore} \sqsubseteq \textit{Animal} \sqcap \forall \textit{eats}.\textit{Plant}, \\ & \textit{Omnivore} \sqsubseteq \textit{Animal} \sqcap \exists \textit{eats}.\textit{Animal} \sqcap \exists \textit{eats}.\textit{Plant} \} \end{aligned}$$

Please note that our “Carnivore” is also known as *obligate* carnivore. A good similarity function $Sim(\cdot)$ is expected to derive that $Sim(\textit{Carnivore}, \textit{Omnivore}) > Sim(\textit{Carnivore}, \textit{Herbivore})$ because the first pair share more **common** subsumers and have fewer **distinguishing** subsumers. On the one hand *Carnivore*, *Herbivore* and *Omnivore* are all subsumed by the following **common** subsumers (abbreviated for readability): $\{\top, \textit{Org}, \textit{A}, \exists e.\top\}$. In addition, *Carnivore* and *Omnivore* share the following **common** subsumer: $\{\exists e.\textit{A}\}$. On the other hand, they have the following **distinguishing** subsumer: $\{\exists e.\textit{P}\}$ while *Carnivore* and *Herbivore* have the following **distinguishing** subsumers: $\{\exists e.\textit{P}, \forall e.\textit{P}, \exists e.\textit{A}, \forall e.\textit{A}\}$. Here, we have made a choice to ignore (infinitely) many subsumers and only consider a select few. Clearly, this choice has an impact on $Sim(\cdot)$. Details on such design choices are discussed later.

We refer to the property of accounting for both common and distinguishing features as rationality. In addition, the related literature refer to some other properties for evaluating similarity measures (e.g., equivalence closure, symmetry, triangle inequality, monotonicity, subsumption preservation, structural dependence). For a detailed overview, the reader is referred to [6, 18].

4 Overview of existing approaches

We classify existing similarity measures according into two dimensions as follows.

Taxonomy vs. ontology based measures Taxonomy-based measures [23, 33, 24, 20, 16] only consider the taxonomic representation of the ontology (e.g., for DLs, we *could use* the inferred class hierarchy); hence only atomic subsumptions are considered (e.g., *Carnivore* \sqsubseteq *Animal*). In fact, this can be considered an approximated solution to the problem which might be sufficient in some cases. However, the user must be aware of the limitations of such approaches. For example, direct siblings are always considered equi-similar although some siblings might share more features/subsumers than others.

Ontology-based measures [6, 15, 18] take into account more of the knowledge in the underlying ontology (e.g., $Carnivore \sqsubseteq \forall eats.Animal$). These measures can be further classified into (a) structural measures, (b) interpretation-based measures or (c) hybrid. Structural measures [15, 18] first transform the compared concepts into a normal form (e.g., \mathcal{EL} normal form or \mathcal{ALCN} disjunctive normal form) and then compare the syntax of their descriptions. To avoid being purely syntactic, they first unfold the concepts w.r.t. the $TBox$ which limits the applicability of such measures to cyclic terminologies. Some structural measures [18] are applicable only to inexpressive DLs (e.g., \mathcal{EL}) and it is unclear how they can be extended to more expressive DLs. Interpretation-based measures mainly depend on the notion of canonical models (e.g., in [6] the canonical model based on the $ABox$ is utilised) which do not always exist (e.g., consider disjunctions).

Intensional vs. extensional based measures Intensional-based measures [23, 33, 15, 18] exploit the terminological part of the ontology while extensional-based measures [24, 20, 16, 6] utilise the set of individual names in an $ABox$ or instances in an external corpus. Extensional-based measures are very sensitive to the content under consideration; thus, adding/removing an individual name would change similarity measurements. These measures might be suitable for specific content-based applications but might lead to unintuitive results in other applications because they do not take concept definitions into account. Moreover, extensional-based measures cannot be used with pure terminological ontologies and always require representative data.

5 Detailed inspection of some existing approaches

After presenting a general overview of existing measures, we examine in detail some measures that can be considered “cheap” options and explore their possible problems. In what follows, we use $S_{Atomic}(C)$ to denote the set of atomic subsumers for concept C . We also use $Com_{Atomic}(C, D)$, $Diff_{Atomic}(C, D)$ to denote the sets of common and distinguishing atomic subsumers respectively.

5.1 Rada et al.

This measure utilises the length of the shortest path [23] between the compared concepts in the inferred class hierarchy. The essential problem here is that the measure takes only distinguishing features into account and ignores any possible common features.

5.2 Wu and Palmer

To account for both common and distinguishing features, Wu & Palmer [33] presented a different formula for measuring similarity, as follows:

$$S_{Wu \& Palmer}(C, D) = \frac{2 \cdot |Com_{Atomic}(C, D)|}{2 \cdot |Com_{Atomic}(C, D)| + |Diff_{Atomic}(C, D)|}$$

Although this measure accounts for both common and distinguishing features, it only considers atomic concepts and it is more sensitive to commonalties.

5.3 Resnik and other IC measures

In information theoretic notions of similarity, the information content $IC_C = -\log P_C$ of a concept C is computed based on the probability (P_C) of encountering an instance of that concept. For example, $P_{\top} = 1$ and $IC_{\top} = 0$ since \top is not informative. Accordingly, Resnik [24] defines similarity $S_{\text{Resnik}}(C, D)$ as:

$$S_{\text{Resnik}}(C, D) = IC_{LCS}$$

where LCS is the least common subsumer of C and D (i.e., the most specific concept that subsumes both C and D). IC measures take into account features that some instances of C and D have, which are not necessarily neither common nor distinguishing features of all instances of C and D . In addition, Resnik’s measure in particular does not take into account how far the compared concepts are from their least common subsumer. To overcome this problem, two [20, 16] other IC-measures have been proposed:

$$S_{\text{Lin}}(C, D) = \frac{2 \cdot IC_{LCS}}{IC_C + IC_D}$$

$$S_{\text{Jiang\&Conrath}}(C, D) = 1 - IC_C + IC_D - 2 \cdot IC_{LCS}$$

6 A new family of similarity measures

Following our exploration of existing measures and their associated problems, we present a new family of similarity measures that addresses these problems. The new measures adopt the features model where the features under consideration are the subsumers of the concepts being compared. The new measures are based on Jaccard’s similarity coefficient [13] which has been proved to be a proper metric (i.e., satisfies the properties: equivalence closure, symmetry and triangle inequality). Jaccard’s coefficient, which maps similarity to a value in the range $[0,1]$, is defined as follows (for sets of “features” A', B' of A, B , i.e., subsumers of A and B):

$$J(A, B) = \frac{|(A' \cap B')|}{|(A' \cup B')|}$$

We aim at similarity measures for general OWL ontologies and thus a naive implementation of this approach would be trivialised because a concept has infinitely many subsumers. To overcome this issue, we present some refinements for the similarity function in which we do not simply count all subsumers but consider subsumers from a set of (possibly complex) concepts of a concept language \mathcal{L} . More precisely, for concepts C, D an ontology \mathcal{O} and a concept language \mathcal{L} , we set:

$$S(C, \mathcal{O}, \mathcal{L}) = \{D \in \mathcal{L}(\tilde{\mathcal{O}}) \mid \mathcal{O} \models C \sqsubseteq D\}$$

$$\text{Com}(C, D, \mathcal{O}, \mathcal{L}) = S(C, \mathcal{O}, \mathcal{L}) \cap S(D, \mathcal{O}, \mathcal{L})$$

$$\text{Union}(C, D, \mathcal{O}, \mathcal{L}) = S(C, \mathcal{O}, \mathcal{L}) \cup S(D, \mathcal{O}, \mathcal{L})$$

$$\text{Sim}(C, D, \mathcal{O}, \mathcal{L}) = \frac{|\text{Com}(C, D, \mathcal{O}, \mathcal{L})|}{|\text{Union}(C, D, \mathcal{O}, \mathcal{L})|}$$

To design a new measure, it remains to specify the set \mathcal{L} . In what follows, we present some examples:

$$\begin{aligned} \text{AtomicSim}(C, D) &= \text{Sim}(C, D, \mathcal{O}, \mathcal{L}_{\text{Atomic}}(\tilde{\mathcal{O}})), \text{ and } \mathcal{L}_{\text{Atomic}}(\tilde{\mathcal{O}}) = \tilde{\mathcal{O}} \cap N_C. \\ \text{SubSim}(C, D) &= \text{Sim}(C, D, \mathcal{O}, \mathcal{L}_{\text{Sub}}(\tilde{\mathcal{O}})), \text{ and } \mathcal{L}_{\text{Sub}}(\tilde{\mathcal{O}}) = \text{Sub}(\mathcal{O}). \\ \text{GrSim}(C, D) &= \text{Sim}(C, D, \mathcal{O}, \mathcal{L}_{\text{G}}(\tilde{\mathcal{O}})), \text{ and } \mathcal{L}_{\text{G}}(\tilde{\mathcal{O}}) = \{E \mid E \in \text{Sub}(\mathcal{O}) \\ &\text{ or } E = \exists r.F, \text{ for some } r \in \tilde{\mathcal{O}} \cap N_R \text{ and } F \in \text{Sub}(\mathcal{O})\}. \end{aligned}$$

where $\text{Sub}(\mathcal{O})$ is the set of concept expressions in \mathcal{O} . $\text{AtomicSim}(\cdot)$ captures taxonomy-based measures since it considers atomic concepts only. The rationale of $\text{SubSim}(\cdot)$ is that it provides similarity measurements that are sensitive to the modeller’s focus. It also provides a cheap (yet principled) way for measuring similarity in expressive DLs since the number of candidates is linear in the size of the ontology. To capture more possible subsumers, one can use $\text{GrSim}(\cdot)$. We have chosen to include only grammar concepts which are subconcepts or which take the form $\exists r.F$ to make the following experiments more manageable. However, the grammar can be extended easily.

7 Approximations of similarity measures

Some of the presented examples for similarity measures might be practically inefficient due to the large number of candidate subsumers. For this reason, it would be nice if we can explore and understand whether a “cheap” measure can be a good approximation for a more expensive one. We start by characterising the properties of an approximation in the following definition.

Definition 1 *Given two similarity functions $\text{Sim}(\cdot), \text{Sim}'(\cdot)$, and an ontology \mathcal{O} , we say that:*

- $\text{Sim}'(\cdot)$ preserves the order of $\text{Sim}(\cdot)$ if $\forall A_1, B_1, A_2, B_2 \in \tilde{\mathcal{O}}: \text{Sim}(A_1, B_1) \leq \text{Sim}(A_2, B_2) \implies \text{Sim}'(A_1, B_1) \leq \text{Sim}'(A_2, B_2)$.
- $\text{Sim}'(\cdot)$ approximates $\text{Sim}(\cdot)$ from above if $\forall A, B \in \tilde{\mathcal{O}}: \text{Sim}(A, B) \leq \text{Sim}'(A, B)$.
- $\text{Sim}'(\cdot)$ approximates $\text{Sim}(\cdot)$ from below if $\forall A, B \in \tilde{\mathcal{O}}: \text{Sim}(A, B) \geq \text{Sim}'(A, B)$.

Consider $\text{AtomicSim}(\cdot)$ and $\text{SubSim}(\cdot)$. The first thing to notice is that the set of candidate subsumers for the first measure is actually a subset of the set of candidate subsumers for the second measure ($\tilde{\mathcal{O}} \cap N_C \subseteq \text{Sub}(\mathcal{O})$). However, we need to notice also that the number of entailed subsumers in the two cases need not to be proportionally related. For example, if the number of atomic candidate subsumers is n and two compared concepts share $\frac{n}{2}$ common subsumers. We cannot conclude that they will also share half of the subconcept subsumers. They could actually share all or none of the complex subsumers. Therefore, the order-preserving property need not be always satisfied. As a concrete example,

let the number of common and distinguishing atomic subsumers for C and D to be 2 and 4 respectively (out of 8 atomic concepts) and let the number of their common and distinguishing subconcept subsumers to be 4 and 6 respectively (out of 20 subconcepts). Let the number of common and distinguishing atomic subsumers for C and E to be 4 and 4 respectively and let the number of their common and distinguishing subconcept subsumers to be 4 and 8 respectively. In this case, $AtomicSim(C, D) = \frac{2}{6} = 0.33$, $SubSim(C, D) = \frac{4}{10} = 0.4$, $AtomicSim(C, E) = \frac{4}{8} = 0.5$, $SubSim(C, E) = \frac{4}{12} = 0.33$. Notice that $AtomicSim(C, D) < AtomicSim(C, E)$ while $SubSim(C, D) > SubSim(C, E)$. Here, $AtomicSim(\cdot)$ is not preserving the order of $SubSim(\cdot)$ and $AtomicSim(\cdot)$ underestimates the similarity of C,D and overestimates the similarity of C,E compared to $SubSim(\cdot)$.

A similar argument can be made to show that entailed subconcept subsumers are not necessarily proportionally related to the number of entailed grammar-based subsumers. We conclude that the above examples of similarity measures are, theoretically, non-approximations of each other. In the next section, we are interested in knowing the relation between these measures in practice.

8 Empirical evaluation

The empirical evaluation constitutes two parts. In Experiment 1, we carry out a comparison between the three measures $GrSim(\cdot)$, $SubSim(\cdot)$ and $AtomicSim(\cdot)$ against human experts-based similarity judgments. In [22], IC-measures along with Rada measure [23] has been compared against human judgements using the same data set which is used in the current study. The previous study [22] has found that IC-measures are worse than Rada measure so we only include Rada measure in our comparison and exclude IC-measures. We also include another path-based measure with is Wu & Palmer [33]. In Experiment 2, we further study in detail the behaviour of our new family of measures in practice. $GrSim(\cdot)$ is considered as the expensive and most precise measure in this study. We use $AtomicSim(\cdot)$ as the cheap measure as it only considers atomic concepts as candidate subsumers. Studying this measure can allow us to understand the problems associated with taxonomy-based measures as they all consider atomic subsumers only. Recall that taxonomy-based measures suffer from other problems that were presented in the conceptual inspection section. Hence, $AtomicSim(\cdot)$ can be considered the best candidate in its class since it does not suffer from these problems. We also consider $SubSim(\cdot)$ as a cheaper measure than $GrSim(\cdot)$ and more precise than $AtomicSim(\cdot)$ and we expect it to be a better approximation for $GrSim(\cdot)$ compared to $AtomicSim(\cdot)$. We excluded from the study instance-based measures since they require representative data which is not guaranteed to be present in our corpus of ontologies.

We have shown in the previous section that the above three measures are not approximations of each other. However, this might not be the case in practice as we will explore in the following experiment. We study the relation between $AtomicSim(\cdot)$ and $SubSim(\cdot)$ and refer to this as AS , the relation between $AtomicSim(\cdot)$ and $GrSim(\cdot)$ and refer to this as AG , the relation between

$SubSim(\cdot)$ and $GrSim(\cdot)$ and refer to this as SG . For each relation, we examine the following properties: (1) order-preservation, (2-3) approximation from above/below, (4) correlation and (5) closeness. Properties 1-3 are defined in Definition 1. For correlations, we calculate Pearson’s coefficient for the relation between each pair of measures. Finally, two measures are considered close if the following property holds: $|Sim_1(C, D) - Sim_2(C, D)| \leq \Delta$ where $\Delta = 0.1$ in the following experiment.

8.1 Infrastructure

With respect to hardware, we used the following machine: Intel Quad-core i7 2.4GHz processor, 4 GB 1333 MHz DDR3 RAM, running Mac OS X 10.7.5. As for the software, firstly, the OWL API v3.4.4 [11] is used. Secondly, to avoid runtime errors caused by using some reasoners with some ontologies, a stack of freely available reasoners were utilised: FaCT++ [29], HermiT [26], JFact ², and Pellet [28].

8.2 Test data

Experiment 1 In 1999, SNOMED-CT was jointly developed by the College of American Pathologists (CAP) and the National Health Service (NHS) in the UK. For the purposes of our comparison study, we use the 2010 version of SNOMED-CT. This ontology has been described as the most complete reference terminology in existence for the clinical environment [3, 2]. It provides comprehensive coverage of diseases, clinical findings, therapies, body structures and procedures. As in February 2014, the ontology has 397,924 concepts. These are organised into 13 hierarchies. The ontology has the highest views amongst all BioPortal ontologies with over 13,600 views.

The reason for choosing this particular ontology is the availability of test data that shows the degree of similarity between some concepts from that ontology as rated by medical experts. Pedersen et al. [22] introduced a test set consisting of 30 pairs of clinical terms. The similarity between each pair is rated by two groups of medical experts: physicians and coders. For details regarding the construction of this dataset, the reader is referred to [22]. We consider the average of physicians and coders similarity values in the comparison. We include in our study 19 pairs out of the 30 pairs after excluding pairs that have at least one concept that has been described as an ambiguous concept in the ontology (i.e., is assigned as a subclass of the concept `ambiguous_concept`) or not found in the ontology.

Experiment 2 The BioPortal library of biomedical ontologies has been used for evaluating different ontology-related tools such as reasoners [17], module extractors [7], justification extractors [12], to name a few. The corpus contains 365 user contributed ontologies (as in October 2013) with varying characteristics such as axiom count, concept name count and expressivity.

² <http://jfact.sourceforge.net/>

A snapshot of the BioPortal corpus from November 2012 was used. It contains a total of 293 ontologies. We excluded 86 ontologies which have only atomic subsumptions as for such ontologies the behaviour of the considered measures will be identical, i.e., we already know that *AtomicSim*(\cdot) is good and cheap. We also excluded 38 more ontologies due to having no concept names or due to run time errors. This has left us with a total of 169 ontologies.

Due to the large number of concept names (565,661) and difficulty of spotting interesting patterns by eye, we calculated the pairwise similarity for a sample of concept names from the corpus. The size of the sample is 1,843 concept names with 99% confidence level. To ensure that the sample encompasses concepts with different characteristics, we picked 14 concepts from each ontology. The selection was not purely random. Instead, we picked 2 random concept names and for each random concept name we picked some neighbour concept names (i.e., 3 random siblings, atomic subsumer, atomic subsumee, sibling of direct subsumer). This choice was made to allow us to examine the behaviour of the considered similarity measures even with special cases such as measuring similarity among direct siblings.

8.3 Experiment workflow

Experiment 1 The similarity of 19 SNOMED-CT concept pairs was calculated using the three methods along with Rada [23] and Wu & Palmer [33] measures. We compare these similarities to human judgements taken from the Pedersen et al.[22] test set.

Experiment 2

Module extraction: For optimisation, rather than working on the whole ontology, the next steps are performed on a \perp -module [5] with the set of 14 concept names as seed signature. One of the important properties of \perp -modules is that they preserve almost all the seed signature’s subsumers. There are 3 cases in which a \perp -module would miss some subsumers. The first case occurs when $\mathcal{O} \models C \sqsubseteq \forall s.X$ and $\mathcal{O} \models C \sqsubseteq \forall s.\perp$. The second case occurs when $\mathcal{O} \models C \sqsubseteq \forall s.X$ and $\mathcal{O} \models \forall s.X \equiv \top$. The third case occurs when $\mathcal{O} \models C \sqsubseteq \forall s.X$ and $\mathcal{O} \not\models C \sqsubseteq \exists s.X$. Since in all three cases $\forall s.X$ is a vacuous subsumer of C , we chose to ignore these, i.e., use \perp -modules without taking special measures to account for them.

Candidate subsumers extraction: In addition to extracting all atomic concepts in the \perp -module we recursively use the method `getNestedClassExpressions()` to extract all subconcepts from all axioms in the \perp -module. The extracted subconcepts are used to generate grammar-based concepts. For practical reasons, we only generate concepts taking the form $\exists r.D$ s.t. $D \in \text{Sub}(\mathcal{O})$ and r a role name in the signature of the extracted \perp -module. Focusing on existential restrictions is justifiable by the fact that they are dominant in our corpus (77.89% of subconcepts) compared to other complex expression types (e.g., universal restrictions: 2.57%, complements: 0.14%, intersections: 13.89, unions: 2.05%).

Testing for subsumption entailments: For each concept C_i in our sample and each candidate subsumer S_j , we test whether the ontology entails that $C_i \sqsubseteq S_j$. If the entailment holds, subsumer S_j is added to the set of C_i ’s subsumers.

Calculating pairwise similarities: The similarity of each distinct pair in our sample is calculated using the three measures.

8.4 Results and discussion

Experiment 1 (How good are the new measures?) *GrSim* and *SubSim* had the highest correlation values with experts' similarity (Pearson's correlation coefficient $r = 0.87, p < 0.001$). Secondly comes *AtomicSim* with $r = 0.86$. Finally comes Wu & Palmer then Rada with $r = 0.81$ and $r = 0.64$ respectively. Clearly, the new expensive measures are more correlated with human judgements which is expected as they consider more of the information in the ontology. The differences in correlation values might seem to be small but this is expected as SNOMED is an \mathcal{EL} ontology and we expect the differences to grow as the expressivity increases. Figure 1 shows the similarity curves for the 6 measures used in this comparison. As we can see in the figure, the new measures along with Wu & Palmer measure preserve the order of human similarity more often than the Rada measure. And, they mostly underestimated the similarity whereas the Rada measure was mostly overestimating the human similarity.

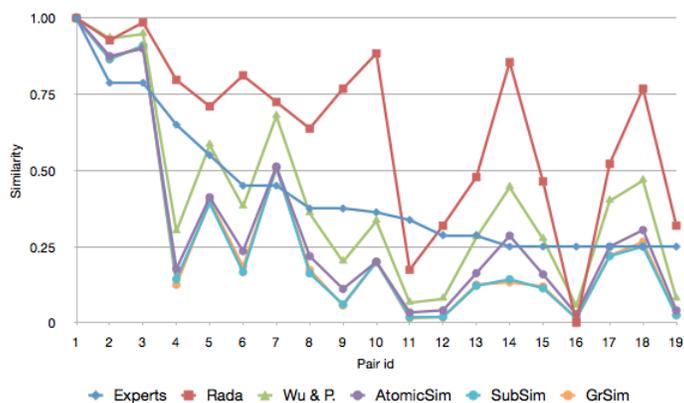


Fig. 1: 6 Curves of similarity for 19 SNOMED clinical terms

Experiment 2 Cost of the new measures One of the main issues we want to explore in this study is the cost (in terms of time) for similarity measurement in general and the cost of the most expensive similarity measure in particular.

The average time per ontology taken to calculate grammar-based pairwise similarities was 2.3 minutes (standard deviation $\sigma = 10.6$ minutes, median $m = 0.9$ seconds) and the maximum time was 93 minutes for the Neglected Tropical Disease Ontology which is a *SRIQ* ontology with 1237 logical axioms, 252 concepts and 99 roles. For this ontology, the cost of *AtomicSim*(\cdot) was only 15.545 sec and 15.549 sec for *SubSim*(\cdot). 9 out of 196 ontologies took over 1 hour to be processed. One thing to note about these ontologies is the high number of

logical axioms and roles. However, these are not necessary conditions for long processing times. For example, the Family Health History Ontology has 431 roles and 1103 logical axioms and was processed in less than 13 sec. Clearly, $GrSim(\cdot)$ is far more costly than the other two measures. This is why we want to know how good/bad a cheaper measure can be.

Approximations and correlations Regarding the relations (AS, AG, SG) between the three measures, we want to find out how frequently can a cheap measure be a good approximation for/have a strong correlation with a more expensive measure. Recall that we have excluded all ontologies with only atomic subsumptions from the study. However, in 21 ontologies (12%), the three measures were perfectly correlated ($r = 1, p < 0.001$) mostly due to having only atomic subsumptions in the extracted module (except for three ontologies which have more than atomic subsumptions). In addition to these perfect correlations for all the three measures, in 11 more ontologies the relation SG was a perfect correlation ($r = 1, p < 0.001$) and AS and AG were very highly correlated ($r \geq 0.99, p < 0.001$). These perfect correlations indicate that, in some cases, the benefit of using an expensive measure is totally neglectable.

In about a fifth of the ontologies (21%), the relation SG shows a very high correlation ($1 > r \geq 0.99, p < 0.001$). Among these, 5 ontologies were 100% order-preserving and approximating from below. In this category, in 22 ontologies the relation SG was 100% close. As for the relation AG , in only 14 ontologies (8%) the correlation was very high.

In nearly half of the ontologies (49%), the correlation for SG was considered medium ($0.99 > r \geq 0.90, p < 0.001$). And in 19 ontologies (11%), the correlation for SG was considered low ($r < 0.90, p < 0.001$) with ($r = 0.63$) as the lowest correlation value. In comparison, the correlation for AG was considered medium in 64 ontologies (38%) and low in 55 ontologies (32.5%).

As for the order-preservations, approximations from above/below and closeness for the relations AG and SG , we summarise our findings in the following table. Not surprisingly, $SubSim(\cdot)$ is more frequently a better approximation to $GrSim(\cdot)$ compared to $AtomicSim(\cdot)$. Although one would expect that the

	Order-preservations	Approx. from below	Approx. from above	Closeness
AG	32	32	37	28
SG	44	49	42	56

Table 1: Ontologies satisfying properties of approximation

properties of an ontology have an impact on the relation between the different measures used to compute the ontology’s pairwise similarities, we found no indicators. With regard to this, we categorised the ontologies according to the degree of correlation (i.e., perfect, high, medium and low correlations) for the SG relation. For each category, we studied the following properties of the ontologies in that category: expressivity, number of logical axioms, number of concepts, number of roles, length of the longest axiom, number of subconcepts. For ontologies

in the perfect correlation category, the important factor was having a low number of subconcepts. In this category, the length of the longest axiom was also low (≤ 11 , compared to 53 which is the maximum length of the longest axiom in all the extracted modules from all ontologies). In addition, the expressivity of most ontologies in this category was \mathcal{AL} . Apart from this category, there were no obvious factors related to the other categories.

How bad is a cheap measure? To explore how likely it is for a cheap measure to encounter problems (e.g., fail one of the tasks presented in the introduction), we examine the cases in which a cheap measure was not an approximation for the expensive measure. AG and SG were not order-preserving in 80% and 73% of the ontologies respectively. Also, they were not approximations from above nor from below in 72% and 64% of the ontologies respectively and were not close in 83% and 66% of the ontologies respectively.

If we take a closer look at the African Traditional Medicine ontology for which the similarity curves are presented in Figure 2, we find that the SG is 100% order-preserving while AG is only 99% order-preserving. Note that for presentation purposes, only part of the curve is shown. Both relations were 100% approximations from below. As for closeness, SG was 100% close while AG was only 12% close. In order to determine how bad are $AtomicSim(\cdot)$ and $SubSim(\cdot)$ as cheap approximations for $GrSim(\cdot)$, we study the behaviour of these measures w.r.t. the Tasks 1-3 presented in the introduction.

Both cheap measures would succeed in performing Task 1 while only $SubSim(\cdot)$ can succeed in Task 2 (1% failure chance for $AtomicSim(\cdot)$). For Task 3, there is a higher failure chance for $AtomicSim(\cdot)$ since closeness is low (12%).

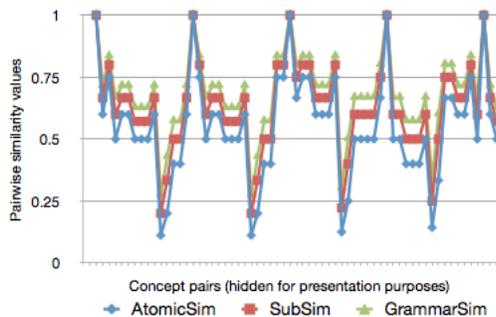


Fig. 2: African Traditional Medicine ontology

As another example, we examine the Platynereis Stage Ontology for which the similarity curves are presented in Figure 3. In this ontology, both AG and SG are 75% order-preserving. However, AG was 100% approximating from above while SG was 85% approximating from below (note the highlighted red spots). In this case, both $AtomicSim(\cdot)$ and $SubSim(\cdot)$ can succeed in Task 1 but not always in Tasks 2 & 3 with $SubSim(\cdot)$ being worse as it can be overestimating in some cases and underestimating in other cases.

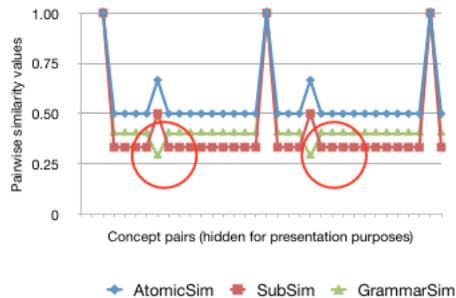


Fig. 3: Platynereis Stage Ontology

In general, both measures are good cheap alternatives w.r.t. Task 1. However, $AtomicSim(\cdot)$ would fail more often than $SubSim(\cdot)$ when performing Tasks 2-3.

9 Threats to validity

9.1 Threats to internal validity

For practical reasons and due to the high runtime of the similarity measurement process, we had to restrict our analysis to a relatively small sample of concepts per ontology. Although the sample is statistically significant in terms of size, it could not be selected in a pure random mechanism. Rather than selecting 14 random concepts, we selected 2 random concepts and 6 neighbour concepts for each random concept. This design option was necessary for understanding the behaviour of similarity measures. Note that non-neighbour concepts tend to have low similarity values. Including a lot of the non-neighbour concepts in our sample could, for example, cause unwanted high percentages for order-preservation.

In addition, relying on only one ontology (i.e., SNOMED-CT) for comparing the new measures and some existing measures against human judgements might limit the generalizability of the results. Rather than dealing with these results as confirmatory, they should be treated as preliminary indicators.

9.2 Threats to external validity

Although ontologies in the BioPortal corpus may not be representative for all available ontologies, it does contain a wide range of ontologies with different properties (e.g., size and expressivity). Moreover, it is built and maintained by a community that has a noticeable interest in the similarity measurement problem. Therefore, it is reasonable to adopt this corpus for testing services that would be provided for its community.

10 Conclusion and future research directions

In conclusion, no obvious indicators were found to inform the decision of choosing between a cheap or expensive measure based on the properties of an ontology.

However, the task under consideration and the error rate allowed in the intended application can help. In general, *SubSim*(\cdot) seems to be a good alternative to the expensive *GrSim*(\cdot). First, it is restricted in a principled way to the modeller's focus. Second, it has less failure chance in practise compared to *AtomicSim*(\cdot).

As for our future research directions, we aim to extend the study by looking deeply at the possible causes of failure (e.g., it can be due to a certain relation between the pair of concepts at the point of failure). And in a broader sense, we aim to extend our research into different notions of similarity and relatedness (e.g., similarity between pairs of concepts usually referred to as relational similarity). Finally, we would like to apply and evaluate the presented measures in a real ontology-based application.

References

1. F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. (eds.) Patel-Schneider. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, second edition, 2007.
2. J.R. Campbell, P. Carpenter, C. Sneiderman, S. Cohn, CG. Chute, and J. Warren. Phase ii evaluation of clinical coding schemes: completeness, taxonomy, mapping, definitions, and clarity. *Journal of the American Medical Informatics Association*, 4:238–251, 1997.
3. C. Chute, S. Cohn, K. Campbell, D. Oliver, and JR. Campbell. The content coverage of clinical classifications. *Journal of the American Medical Informatics Association*, 3:224–33, 1996.
4. T. Cohen and D. Widdows. Empirical distributional semantics: Methods and biomedical applications. *Journal of Biomedical Informatics*, 42(2):390405, 2010.
5. B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler. Modular reuse of ontologies: Theory and practice. *J. of Artificial Intelligence Research*, 31:273318, 2008.
6. C. d'Amato, S. Staab, and N. Fanizzi. On the Influence of Description Logics Ontologies on Conceptual Similarity. In *EKAW '08 Proceedings of the 16th international conference on Knowledge Engineering: Practice and Patterns*, 2008.
7. C. Del Vescovo, P. Klinov, B. Parsia, U. Sattler, T. Schneider, and D. Tsarkov. Syntactic vs. semantic locality: How good is a cheap approximation? In *WoMO 2012*, 2012.
8. W. K. Estes. Statistical theory of distributional phenomena in learning. *Psychological Review*, 62:369–377, 1955.
9. J. Euzenat and P. Shvaiko. *Ontology matching*. Springer-Verlag, 2007.
10. U Hahn, N Chater, and LB Richardson. Similarity as transformation. *COGNITION*, 87 (1):1 – 32, 2003.
11. M. Horridge and S. Bechhofer. The OWL API: A Java API for working with OWL 2 ontologies. In *In Proceedings of the 6th International Workshop on OWL: Experiences and Directions (OWLED)*, 2009.
12. M. Horridge, B. Parsia, and U. Sattler. Extracting justifications from bioportal ontologies. *International Semantic Web Conference*, 2:287–299, 2012.
13. P. Jaccard. Etude comparative de la distribution florale dans une portion des alpes et du jura. *Bulletin de la Societe Vaudoise des Sciences Naturelles*, 37:547–579, 1901.
14. W. James. *The principles of psychology*. dover: New york. (original work published 1890), 1890/1950.

15. K. Janowicz. Sim-dl: Towards a semantic similarity measurement theory for the description logic ALCNR in geographic information retrieval. In *SeBGIS 2006, OTM Workshops 2006*, pages 1681-1692, 2006.
16. J. Jiang and D. Conrath. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proc. of the 10th International Conference on Research on Computational Linguistics, Taiwan*, 1997.
17. Y.-B. Kang, Y.-F. Li, and S. Krishnaswamy. Predicting reasoning performance using ontology metrics. In *ISWC 2012 Lecture Notes in Computer Science. Volume 7649*, 2012.
18. K. Lehmann and A. Turhan. A framework for semantic-based similarity measures for ELH-concepts. *JELIA 2012*, pages 307-319, 2012.
19. V. I. Levenshtein. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, 10:707-710, 1966.
20. D. Lin. An information-theoretic definition of similarity. In *Proc. of the 15th International Conference on Machine Learning*, San Francisco, CA, 1998. Morgan Kaufmann.
21. R. M. Nosofsky. Similarity scaling and cognitive process models. *Annual Review of Psychology*, 43:25-53, 1992.
22. T. Pedersen, S. Pakhomov, S. Patwardhan, and C. Chute. Measures of semantic similarity and relatedness in the biomedical domain. *Journal of Biomedical Informatics*, 30(3):288-299, 2007.
23. R. Rada, H. Mili, E. Bicknell, and M. Blettner. Development and application of a metric on semantic nets. In *IEEE Transaction on Systems, Man, and Cybernetics*, volume 19, page 1730, 1989.
24. P. Resnik. Using information content to evaluate semantic similarity in a taxonomy. In *In Proceedings of the 14th international joint conference on Artificial intelligence (IJCAI95)*, volume 1, pages 448-453, 1995.
25. A. Schlicker, FS. Domingues, J. Rahnenfu hrer, and T. Lengauer. A new measure for functional similarity of gene products based on gene ontology. *BMC Bioinformatics*, 7, 2006.
26. R. Shearer, B. Motik, and I. Horrocks. Hermit: A highly-efficient OWL reasoner. In *Proceedings of the 5th International Workshop on OWL: Experiences and Directions (OWLED-08EU)*, 2008.
27. R.N. Shepard. Toward a universal law of generalization for psychological science. *Science*, 237:1317-1323, 1987.
28. E. Sirin, B. Parsia, B. Cuenca Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2), 2007.
29. D. Tsarkov and I. Horrocks. FaCT++ description logic reasoner: System description. In *Proceedings of the 3rd International Joint Conference on Automated Reasoning (IJCAR)*, 2006.
30. A. Tversky. Features of similarity. *Psychological Review by the American Psychological Association, Inc.*, 84(4), July 1977.
31. A.R. Wagner. Evolution of an elemental theory of pavlovian conditioning. *Learning and Behavior*, 36:253-265, 2008.
32. JZZ. Wang, Z. Du, R. Payattakool, PSS. Yu, and CFF. Chen. A new method to measure the semantic similarity of GO terms. *Bioinformatics*, 2007.
33. Z. Wu and MS. Palmer. Verb semantics and lexical selection. In *Proceedings of the 32nd. Annual Meeting of the Association for Computational Linguistics (ACL 1994)*, page 133138, 1994.