Semantic Relatedness through Ontologies and Wikipedia

Seth Ebner
Washington University in St. Louis
January 3, 2017

1 Introduction

A question of particular interest to tasks such as information retrieval and summarization is how to quantify the relatedness or similarity between two concepts or documents [5, 12]. Information retrieval systems must compare queries to text to find the most relevant documents, and summarization techniques must determine how similar a given sentence is to a larger document. Often, this goes beyond surface-level relationships due to synonymy, hyponymy (subclass relation), hypernymy (superclass relation), and many other phenomena.

A semantic relatedness measure should show that, for example, *dog* and *wolf* are more related to each other than *dog* and *chair* are to each other because *dog* and *wolf* are both *canines*. By convention, a higher semantic relatedness score indicates a stronger relation between two concepts. Semantic relatedness is more general than semantic similarity, as the former also accounts for antonyms as well as concepts that appear together but may not necessarily be similar. *Animal* and *bear* are similar and related, whereas *school* and *teacher* are related but not similar.

Semantic relatedness and semantic similarity require real-world knowledge that often falls outside the text itself [2, 11, 12]. For example, documents about the Emancipation Proclamation and the Gettysburg Address, respectively, may not reference the other concept. However, these two concepts are strongly related through Abraham Lincoln and the American Civil War. Solutions to the problem of quantifying semantic relatedness must then consult external knowledge and common-sense knowledge to be able to accurately determine the relatedness of two concepts [11].

Human knowledge consists of concepts and relations between concepts. In this way, knowledge can be represented as a graph, known as an *ontology*. Nodes in such a graph are concepts, and edges are relationships. Viewing knowledge as a graph allows for semantic relatedness to be reduced to a graph-based problem, which allows for the use of techniques well-founded in graph theory. For various reasons explained later, Wikipedia provides a particularly useful graph for evaluating semantic relatedness measures.

In this report, we address the following questions:

- What relations are expressed in ontologies, and what are their properties?
- What are some common measures of semantic relatedness?
- How are relatedness measures evaluated?
- How can Wikipedia help us determine relatedness?
- What features of Wikipedia are the most useful in determining relatedness?

However, before answering these questions, we first define the terminology used throughout this report.

2 Ontologies

2.1 Ontologies vs. Knowledge Bases

Ontologies are maps of knowledge in a given domain that allow one to reason about concepts. Ontologies are usually visualized as graphs of concepts, where edges represent relations between concepts. Knowledge bases largely consist of instances of concepts. One way in which ontologies differ from knowledge bases (KBs) is that ontologies also incorporate relations between entities, whereas KBs simply store information about entities' properties.

Definition 2.1.1. Ontology [10]
Formally, an **ontology** Θ is a structure that consists of:

- C: a set of **concepts**
- A: a set of **attributes**
- R: a set of **relations**
- P: a set of **primitive data types**
- $\leq_C$: a **concept hierarchy** or **taxonomy**
- $\leq_R$: a **relation hierarchy**

Definition 2.1.2. Knowledge Base [10]
Formally, a **knowledge base** KB is a structure that consists of:

- C_{KB}: a set of **concepts**
- R_{KB}: a set of **relations**
- I: a set of **instances**
- i_C: a function of **concept instantiations**
- i_R: a function of **relation instantiations**

One can see that an ontology is richer in detail than a knowledge base, particularly in its use of attributes and hierarchies. It is important to note that knowledge bases, on the other hand, contain instances of concepts, which are useful for identifying the underlying concept of a named entity (*John Lennon* is an instance of *musician*). However, the hierarchical structure of an ontology is more significant for our purposes. Additionally, an ontology can be supplemented by a knowledge base in order to match instances to concepts. For these reasons, we'll consider only ontologies in this discussion.

The authors of [9] list many common ontologies. One of the most widely used ontologies, WordNet, groups words into "synsets", which are sets of synonyms expressing discrete concepts. It also consists of word-sense pairs arranged into a hierarchical structure. WordNet contains *is-a* and *part-of* relations (described below), and it is often used as a lexical resource (i.e. dictionary or thesaurus) when computing the semantic relatedness of two words. Cyc is a knowledge base that attempts to contain all of humans' common-sense knowledge. Its entries come in the form of terms and assertions, including facts and heuristics. Besides general purpose ontologies, there are domain-specific ontologies, many of which concern medicine. These include the Unified Medical Language System, Medical Subject Headings, and Gene Ontology [9].

2.2 Relations in Ontologies

There are two common relations found in ontologies, but these are by no means the only relations. Some ontologies, such as Cyc, have as many as 17,000 relations [9]. The two relations we'll discuss are the **is-a** and **part-of** relations.

The ***A* is-a *B*** relation denotes that A is a subconcept of B [13]. This relation does not concern instances of concepts, meaning that A is not an instance of B. However, every instance of A is an instance of B. For example, a *cat* is-a *mammal*, *cat* is not an instance of *mammal*, and *Garfield* (an instance of *cat*) is also an instance of *mammal*. The *is-a* relation is transitive.

The ***A* part-of *B*** relation denotes that A is necessarily a part of B [13]. For example, *wheel* is part-of *car*. The *part-of* relation is transitive.

An example of an ontology is given in Figure 1. In this ontology, *dolphin* is-a *mammal*, *flower* is part-of *plant*, and all of these are (is-a) *organisms*.

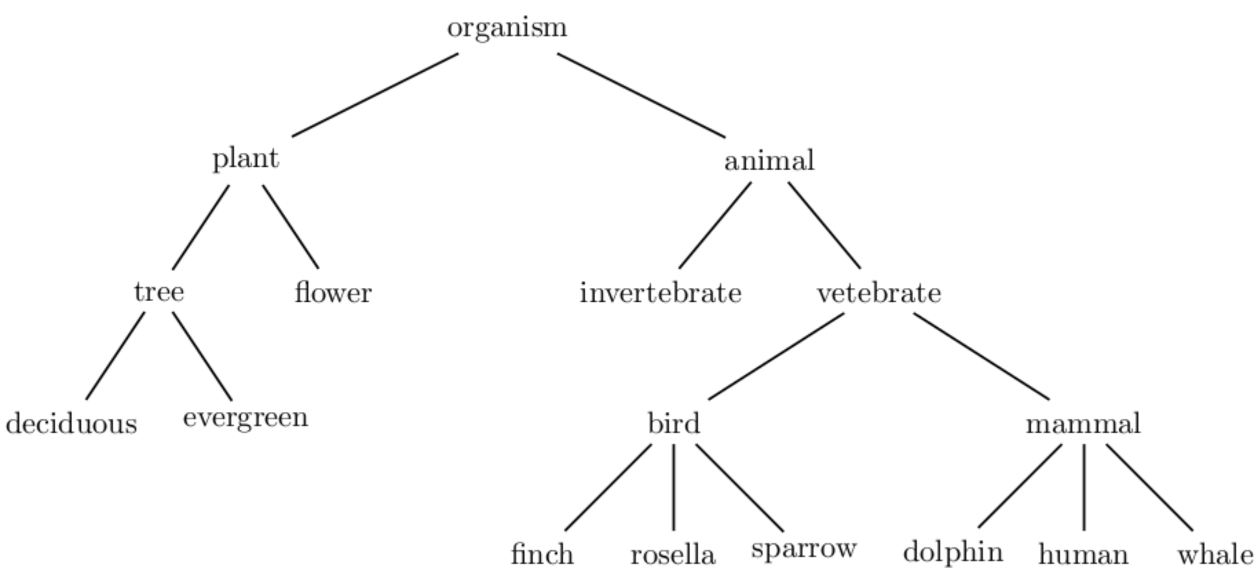

Figure 1: An example ontology. (`https://mvngu.wordpress.com`)

2.3 Graph Properties of Ontologies

When an ontology locally takes the form of a tree, we can discuss the relationship between two concept nodes. Borrowing terms from syntax, we define:

Definition 2.3.1. Dominance, ancestor, descendant [8]
A node X **dominates** a node Y if X appears above Y in the ontology. That is, X is more general than Y. X is called an **ancestor** of Y, and Y is called a **descendant** of X.

Definition 2.3.2. Immediate dominance, parent, child [8]
A node X **immediately dominates** a node Y if X dominates Y and no node that X dominates also dominates Y. In other words, X is the node directly above Y. X is called a **parent** of Y, and Y is called a **child** of X.

Definition 2.3.3. Siblings [8]
Nodes are **siblings** if they have the same parent, that is, they are immediately dominated by the same node.

Definition 2.3.4. Lowest Common Ancestor
The **lowest common ancestor** of two nodes, X and Y, is the deepest node that dominates both X and Y.

In the ontology described by Figure 1, *vertebrate* is the parent of both *bird* and *mammal*, which are children of *vertebrate*. *Vertebrate* dominates *finch*, but is not its parent because *bird* is an intermediate node. *Finch*, *rosella*, and *sparrow* are siblings, and their parent is *bird*. *Vertebrate* is the lowest common ancestor of *finch* and *whale*.

3 Wikipedia

Wikipedia (`http://en.wikipedia.org`) is the largest existent knowledge repository [2,11]. As of this writing, Wikipedia contains over 5 million articles. In addition to its size, Wikipedia offers a host of explicitly defined semantics and a highly organized structure due to its links, portals, and categories [2,4,11]. Moreso, the concepts expressed in Wikipedia are intuitive and easily explained, as humans defined the concepts for others to understand.

Studies have also found the quality of the articles' content to be comparable to other highly regarded encyclopedias [3]. [2] notes that Wikipedia articles are almost noise-free and qualify as Standard Written English. These qualities make Wikipedia an attractive resource for NLP tasks, especially tasks concerning semantics, information retrieval, and information extraction.

3.1 Graph Properties of Wikipedia

The graph structure of Wikipedia comes from its links. The text to which a link is attached is called an **anchor**, as defined in [11]. There are three main types of links: infobox, categorical, and content [12]:

- **Infobox** links are located in the infobox sections of articles. These links often list a concept's attributes, such as a person's date and location of birth or a food's main ingredients.
- **Categorical** links reference category- and list-based articles. The linked-to articles often contain many links of their own, which can help one recognize inter-article relations such as hyponymy, hypernymy, and synonymy.
- **Content** links are those that appear in the article text itself.

4 Semantic Relatedness

In this section, we discuss some properties of an ontology that should be reflected in a semantic relatedness measure before moving on to some background on the most commonly used test sets. We then give some examples of the main categories of measures. Finally, three of the latest Wikipedia-based semantic relatedness measures are discussed.

4.1 Semantics Encoded in Ontologies

There are some hierarchical properties that a relatedness measure should capture. For instance, we expect that a node should be more similar to its child or its parent than to a node farther away. Additionally, sibling nodes should have a high similarity because they are separated by only one intermediate node (their parent). In general, there should be high relatedness among nodes on a given path from the root to a given node, with pairs of nodes closer together on the path having correspondingly higher similarities than pairs of nodes farther away on the path.

4.2 Evaluation of the Measures

Judgments of semantic relatedness are inherently innate and subjective [9]. This lack of an objective consensus means that human judgments should be defined as correct. Agreement among judges is also of great importance, since, as [2] notes, "it is this consensus that allows people to understand each other."

There are three main datasets used when evaluating the accuracy of a measure: Rubenstein-Goodenough [7], Miller-Charles [6], and WordSimilarity-353 [1]. They contain word pairs and human judgments on their semantic relatedness. Recreations of the original studies show that there is a high correlation ($r \geq 0.95$) among the perceived semantic similarities, even across groups and time periods [9]. This indicates that the innate similarity measure is steady. Accuracy of a semantic relatedness measure is given as the correlation between the measure's predicted relatedness scores and the human-determined relatedness scores.

4.3 Types of Semantic Relatedness Measures

There are many approaches to defining semantic relatedness measures. These include structure-based, information content-based, feature-based, and hybrid measures [9]. Structure-based measures do not consider the content of the concepts or documents. Rather, they take a purely hierarchical view in which semantic similarity is tied to the path length between two concepts, the concepts' locations in the hierarchy, and closest common ancestor of the concepts. Information content-based measures use the most informative ancestor of the two concepts being compared, as well as the amount of information shared by the concepts and the amount of information that differs between the concepts [4]. Feature-based measures make use of features present in the ontology's hierarchical structure, such as sets of shared ancestors. Hybrid measures combine these approaches. Listed below are some examples of structure-based and feature-based measures in order for one to get a sense of what characteristics are considered important when designing a measure.

4.3.1 Structure-based Measures

One of the simplest measures of semantic relatedness is the notion of shortest path. The similarity of concepts a and b is defined as:

$$sim(a, b) = 2 \cdot max(a, b) - min(a, b)$$

where $max(a, b)$ is the maximum path length between a and b, and $min(a, b)$ is the minimum path length between a and b [9]. There also exists a weighted version of the shortest-path measure.

The shortest-path measures do not account for the locations of the nodes in the hierarchy, just their relative positions. Other measures account for location, such as the Wu and Palmer measure, defined as:

$$sim(a, b) = \frac{2 \cdot N}{N_a + N_b + 2 \cdot N}$$

where N is the distance from the root to the lowest common ancestor of a and b, and N_a and N_b are the distances from a and b to their lowest common ancestor, respectively [4, 5, 9].

Measures can also be highly non-linear, as seen in the following measure defined by Li et al.:

$$sim(a,b) = e^{-\alpha \cdot min(a,b)} \cdot \frac{e^{\beta \cdot N} - e^{-\beta \cdot N}}{e^{\beta \cdot N} + e^{-\beta \cdot N}}$$

where α and β are empirically chosen, and $min(a,b)$ and N are defined as before [4, 9].

4.3.2 Feature-based Measures

A common feature used in measures is the set of ancestors of a node u, denoted as $A(u)$. A simple measure can then be defined as:

$$sim(a,b) = \frac{|A(a) \cap A(b)|}{|A(a) \cup A(b)|}$$

which is similar to the Jacard index [4].

A slightly different measure that has a tunable symmetry parameter is:

$$sim(a,b) = \alpha \frac{|A(a) \cup A(b)|}{|A(a)|} + (1-\alpha) \frac{|A(a) \cup A(b)|}{|A(b)|}$$

where $\alpha \in [0,1]$ [4].

4.4 Wikipedia-Based Algorithms for Determining Semantic Relatedness

In this section, we'll look at three algorithms that compute semantic relatedness between two words or documents through the use of Wikipedia, and discuss the performance, advantages, and disadvantages of each.

4.4.1 Semantic Relatedness via Concepts

One of the most successful approaches to semantic relatedness using Wikipedia is Explicit Semantic Analysis (ESA) [2]. ESA uses just the text within the given documents and the Wikipedia articles, and does not make use of any structure in the Wikipedia graph. Instead, it relies on the (generally valid) assumption that each Wikipedia article concerns only one topic.

ESA improves upon Latent Semantic Analysis (LSA) by ensuring that the concepts being manipulated are natural to humans. The dimensions found through the LSA dimensionality reduction technique are notoriously hard to interpret. ESA also is able to perform word sense disambiguation by looking to a word's neighbors for context, which LSA cannot do.

The key tool used in ESA is a "semantic interpreter" that maps words into the space of Wikipedia articles. For ease of computation, an inverted index is constructed that maps each word in a document to a list of concepts (articles) in which it appears. The inverted index is then pruned to remove weak connections. For a document D, its semantic interpretation vector V (of length N, the total number of Wikipedia articles) is calculated such that element v_i has a value of:

$$v_i = \sum_{w \in D} t_i \cdot k_j$$

where t_i is the weight of word w_i in the document's tf–idf vector, and k_j is the strength of the association of w_i with concept c_j as determined by the inverted index entry for w_i. Semantic relatedness between two documents is then computed as the cosine similarity of their semantic interpretation vectors.

ESA improves upon three main varieties of prior approaches: bag-of-words (BOW) models, techniques leveraging lexical resources, and LSA [2]. BOW treats each unique word as its own concept, which is inappropriate in many cases. This also means that BOW cannot handle synonymy. Techniques involving thesauri or other lexical resources, such as WordNet, are inherently limited in their scope because the resources are created manually [2, 11]. It is common for such resources to omit proper names and domain-specific terms. These resources often contain detailed information about specific words but not about world knowledge. Additionally, it is not possible to perform word sense disambiguation with WordNet. LSA is a dimensionality reduction technique that has seen much success. However, it is a purely mathematical technique that does not use any information about world knowledge or its structure, and the dimensions it projects words into are consequently difficult to interpret. ESA, in contrast, represents the meanings of documents through concepts that make sense to humans.

Performance: On the WordSimilarity-353 test set (word-word comparison), ESA has an accuracy of 0.75. On the Australian Broadcasting Corporation dataset (document-document comparison), ESA has an accuracy of 0.72 [2]. LSA previously had the best accuracy on those datasets, with scores of 0.56 and 0.60, respectively.

Advantages: There are many advantages of ESA. From an implementation side, input texts are simply plain text and do not require deep language understanding or pre-catalogued common-sense knowledge [2]. Furthermore, unlike some previous work, such as WikiRelate!, ESA can compare documents of any length [2]. There is no component in the algorithm that limits the length of the input text.

Disadvantages: ESA has disadvantages as well. It requires a large amount of pre-processing, although this cost is incurred only once. This pre-processing involves creating the semantic interpreter and inverted index. ESA also requires large amounts of textual data, as this approach relies solely on the textual content of Wikipedia articles. This is in excess of 13 Gb [11].

4.4.2 Semantic Relatedness via Content Links

One of the most glaring features of ESA is that it does not make use of the structure of Wikipedia. The Wikipedia Link-based Measure (WLM) [11] uses the content links in an article to tackle the problem of determining semantic relatedness. This approach makes note that most of the text in a Wikipedia article does not help to compare two concepts, and it is primarily the links that allow for such a comparison. That is to say, the manually defined semantics are more important than the textual content. Consideration of only the links in an article and not its content should make computation of semantic relatedness cheaper and more accurate than ESA [11].

WLM handles ambiguity and polysemy (concepts known by many names) through its use of anchors, defined earlier. The large number of links means there is a large number of anchors which capture ambiguity and polysemy [11].

The authors of [11] give two measures of similarity between two Wikipedia articles, one based on out-links and one based on in-links. The measure concerning out-links is nearly identical to a tf-idf vector. The difference is that link counts are weighted according to the probability of the link occurring, as opposed to word counts being weighted according to the probability of the word occurring. Only out-links of the two articles that are being compared are considered. The weight of a given link is defined as:

$$w(source \rightarrow target) = \begin{cases} -log(\frac{|T|}{|W|}) & source \in T \\ 0 & source \notin T \end{cases}$$

where T is the set of all articles linking to the target article, and W is the set of all Wikipedia articles [11]. Defined in this way, links to articles that have lots of in-links are given smaller weights and links to articles that have few in-links are given larger weights. The similarity of the two articles is calculated as the cosine similarity of the articles' tf-idf-like vectors.

The measure concerning in-links is based on the Normalized Google Distance, and for two articles a and b is defined as:

$$semrel(a, b) = \frac{log(max(|A|, |B|)) - log(|A \cap B|)}{log(|W|) - log(min(|A|, |B|))}$$

where A and B are the sets of all articles linking to a and b, respectively. This measure captures the original intent of the Normalized Google Distance, which is that documents containing both terms indicate relatedness and documents containing only one of the terms indicate lack of relatedness.

Performance: When evaluated on the WordSimilarity-353 dataset, the tf-idf measure has an accuracy of 0.66, and the Normalized Google Distance measure has an accuracy of 0.72. Both of these results are slightly worse than ESA's. However, by performing a combination of disambiguation techniques, the authors are able to achieve an accuracy of 0.78, which outperforms ESA.

Advantages: One of the most prominent advantages of WLM over ESA is its lower overhead. The preprocessing in ESA requires all of the text in Wikipedia articles, and the construction of the inverted index requires log-linear sorting. WLM requires less than 5% of the data (590 Mb vs. 13 Gb) that ESA needs, and there is no preprocessing required beyond linear-time extraction of anchor statistics from Wikipedia dumps. The authors of [11] see WLM as a competitive alternative to ESA, given its smaller resource demands.

Disadvantages: The authors of [11] found that anchors tend to link to instances of concepts rather than the concepts themselves. To avoid this issue, they reduce the set of articles to consider only those that receive a certain proportion of an anchor's links. Additionally, WLM performs worst when the terms being compared do not have appropriate Wikipedia articles. When the terms have suitable articles, WLM performance approaches that of ESA. Finally, WLM only uses content links and does not take full advantage of the types of links available in Wikipedia articles.

4.4.3 Semantic Relatedness via Random Walks

WLM uses only content links and does not consider infobox or categorical links. The intuition behind WikiWalk is that these links contribute significantly to the relatedness of the concepts discussed in Wikipedia articles [12].

WikiWalk uses a random walk algorithm to compare similarity. In a random walk, a hypothetical particle randomly walks along the edges of a graph, but it may "teleport" to a new node with some given probability. These probabilities are stored in a "teleport vector." The weight assigned to a node is equal to the steady-state probability of the particle occupying that node, in other words, how long the particle spends at that node. The nodes in the WikiWalk graph are articles, and the edges are links [5, 12].

To compute relatedness between two words, each word is assigned a teleport vector, their random walks are performed, and the resulting node weights are compared (as vectors) using cosine similarity. The best performance is seen when the teleport vectors are initialized according to the vectors returned by running ESA, with some additional normalization and pruning. Document relatedness is performed almost equivalently (all words are considered instead of just one word pair) [12].

The authors of [12] also use the concept of *generality*, which is calculated as the ratio of in-links of a target article to the in-links of a source article. More general articles should have more in-links than highly specific articles.

The authors make two main points in their findings: Wikipedia's textual content seems to be more useful than its link structure in determining semantic relatedness, and not all links offer the same relatedness information. When the vectors returned by WikiWalk are compared directly to those returned by ESA, it can be seen that using all the information available within a subset of the edges (links) yields gains over ESA. Pruning, categorizing by link type, and categorizing by generality are all important in the WikiWalk method [12].

Performance: On the Miller-Charles test set, WikiWalk has a maximum accuracy of 0.751. WikiWalk has a correlation of 0.634 with the WordSimilarity-353 test set. These figures are below those of ESA, and the authors of [12] attribute this to the inequality of link significance. Not all links contribute equally to semantic relatedness. Random walks can lower performance below that of ESA, and experiments show that categorical links are the most harmful and infobox links are the least harmful [12]. When accounting for links with only a given generality or specificity, WikiWalk is able to achieve a document-document accuracy of 0.772, which slightly outperforms ESA.

Advantages: WikiWalk's main advantage is that it uses all the links available in Wikipedia. When the teleport vectors are initialized according to ESA, the textual content is included also, so no information has been ignored. Using all of the information available in this way gives an accuracy that slightly outperforms the state of the art. WikiWalk is also easily adaptable so that links of certain types can be included or excluded from consideration.

Disadvantages: In most cases, WikiWalk performs worse than ESA, and a dictionary-based initialization procedure yields accuracies well below much prior work. Initializing the teleport vectors by running ESA brings with it the large pre-processing cost of ESA.

5 Conclusion

Ontologies are taxonomic structures that represent knowledge about some domain. Reasoning about relations such as *is-a* and *part-of* and ontological kinship properties such as parent, child, sibling, and lowest common ancestor give way to many semantic relatedness measures.

Semantic relatedness measures are evaluated on test sets of human-judged word pairs or document pairs, most notably the Rubenstein-Goodenough, Miller-Charles, and WordSimilarity-353 datasets. Human judgments are taken to be ground truth because judgments are inherently qualitative and also because the measures aim to emulate the relatednesses inferred during human conversation.

Measures come in many varieties, most notably structure-based, information content-based, feature-based, and hybrid. Structure-based measures concern the layout of the ontology, information content-based measures take inspiration from information theory, feature-based measures deal with characteristics such as ancestors of nodes and synonym sets, and hybrid measures combine these approaches.

Wikipedia is an attractive choice for studying semantic relatedness due to its size, manually defined semantics, link structure, and quality. Another helpful property is that each article concerns only one concept. The three main types of links that could be considered when computing semantic relatedness are: infobox, categorical, and content links, and each type contributes differently to the overall semantics of a Wikipedia article.

Explicit Semantic Analysis shows that the textual content of Wikipedia does a very good job of determining semantic relatedness. Wikipedia Link-based Measure and WikiWalk give slight increases in performance over ESA, indicating that the link structure of Wikipedia is useful as well. In fact, WLM does not use textual content at all and requires far less data than ESA. The best performance of WikiWalk requires running ESA, so the resource requirement remains large.

There is still much room for advancement in this field. There is little discussion of the algorithmic complexity of various semantic relatedness measures. Most comparisons are qualitative (one algorithm requires more pre-processing than another, etc.). Such information would help one decide which measure to use when considering the performance-cost tradeoff.

Another limitation of the Wikipedia-based measures discussed above is English-centricity. These measures concern the English Wikipedia only due to its size. There exist many smaller versions of Wikipedia in other languages. Nothing in the ESA, WLM, or WikiWalk algorithms is inherently English-specific, and it is only a lack of data that prevents these approaches from being applied to other languages. The creation of foreign-language test sets will help to evaluate the measures on foreign languages, and similar performance will corroborate the measures.

References

[1] Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM, 2001.

[2] Evgeniy Gabrilovich and Shaul Markovitch. Computing semantic relatedness using wikipedia-based explicit semantic analysis. In *IJcAI*, volume 7, pages 1606–1611, 2007.

[3] Jim Giles. Internet encyclopaedias go head to head. *Nature*, 438(7070):900–901, 2005.

[4] Sebastien Harispe, Sylvie Ranwez, Stefan Janaqi, and Jacky Montmain. Semantic similarity from natural language and ontology analysis. *Synthesis Lectures on Human Language Technologies*, 8(1):1–254, 2015.

[5] Rada Mihalcea and Dragomir Radev. *Graph-based natural language processing and information retrieval*. Cambridge University Press, 2011.

[6] George A Miller and Walter G Charles. Contextual correlates of semantic similarity. *Language and cognitive processes*, 6(1):1–28, 1991.

[7] Herbert Rubenstein and John B Goodenough. Contextual correlates of synonymy. *Communications of the ACM*, 8(10):627–633, 1965.

[8] Beatrice Santorini and Anthony Kroch. The syntax of natural language: An online introduction using the trees program. *Online textbook: http://www. ling. upenn. edu/beatrice/syntax-textbook/ch7. html*, 2000.

[9] Thabet Slimani. Description and evaluation of semantic similarity measures approaches. *arXiv preprint arXiv:1310.8059*, 2013.

[10] Daniel Sonntag. *Ontologies and Adaptivity in Dialogue for Question Answering*, volume 4. IOS Press, 2010.

[11] Ian Witten and David Milne. An effective, low-cost measure of semantic relatedness obtained from wikipedia links. In *Proceeding of AAAI Workshop on Wikipedia and Artificial Intelligence: an Evolving Synergy, AAAI Press, Chicago, USA*, pages 25–30, 2008.

[12] Eric Yeh, Daniel Ramage, Christopher D Manning, Eneko Agirre, and Aitor Soroa. Wikiwalk: random walks on wikipedia for semantic relatedness. In *Proceedings of the 2009 Workshop on Graph-based Methods for Natural Language Processing*, pages 41–49. Association for Computational Linguistics, 2009.

[13] Ontology relations, http://www.geneontology.org, 2015.