A Web Search Engine-Based Approach to Measure Semantic Similarity between Words

> *"The **Jaguar** is the largest **cat** in Western Hemisphere and can subdue larger prey than can the puma"*

Fig. 1. A snippet retrieved for the query **Jaguar** AND **cat**.

Snippets, a brief window of text extracted by a search engine around the query term in a document, provide useful information regarding the local context of the query term. Semantic similarity measures defined over snippets, have been used in query expansion [2], personal name disambiguation [3], and community mining [4]. Processing snippets is also efficient because it obviates the trouble of downloading webpages, which might be time consuming depending on the size of the pages. However, a widely acknowledged drawback of using snippets is that, because of the huge scale of the web and the large number of documents in the result set, only those snippets for the top-ranking results for a query can be processed efficiently. Ranking of search results, hence snippets, is determined by a complex combination of various factors unique to the underlying search engine. Therefore, no guarantee exists that all the information we need to measure semantic similarity between a given pair of words is contained in the top-ranking snippets.

We propose a method that considers both page counts and lexical syntactic patterns extracted from snippets that we show experimentally to overcome the above mentioned problems. For example, let us consider the snippet shown in Fig. 1 retrieved from Google for the query **Jaguar** AND **cat**.

Here, the phrase *is the largest* indicates a hypernymic relationship between Jaguar and cat. Phrases such as *also known as, is a, part of, is an example of* all indicate various semantic relations. Such indicative phrases have been applied to numerous tasks with good results, such as hypernym extraction [5] and fact extraction [6]. From the previous example, we form the pattern **X** is the largest **Y**, where we replace the two words *Jaguar* and *cat* by two variables **X** and **Y**.

Our contributions are summarized as follows:

- We present an automatically extracted lexical syntactic patterns-based approach to compute the semantic similarity between words or entities using text snippets retrieved from a web search engine. We propose a lexical pattern extraction algorithm that considers word subsequences in text snippets. Moreover, the extracted set of patterns are clustered to identify the different patterns that describe the same semantic relation.
- We integrate different web-based similarity measures using a machine learning approach. We extract synonymous word pairs from WordNet synsets as positive training instances and automatically generate negative training instances. We then train a two-class support vector machine (SVM) to classify synonymous and nonsynonymous word pairs. The integrated measure outperforms all existing web-based semantic similarity measures on a benchmark data set.

- We apply the proposed semantic similarity measure to identify relations between entities, in particular people, in a community extraction task. In this experiment, the proposed method outperforms the baselines with statistically significant precision and recall values. The results of the community mining task show the ability of the proposed method to measure the semantic similarity between not only words, but also between named entities, for which manually created lexical ontologies do not exist or incomplete.

2 RELATED WORK

Given a taxonomy of words, a straightforward method to calculate similarity between two words is to find the length of the shortest path connecting the two words in the taxonomy [7]. If a word is polysemous, then multiple paths might exist between the two words. In such cases, only the shortest path between any two senses of the words is considered for calculating similarity. A problem that is frequently acknowledged with this approach is that it relies on the notion that all links in the taxonomy represent a uniform distance.

Resnik [8] proposed a similarity measure using information content. He defined the similarity between two concepts C_1 and C_2 in the taxonomy as the maximum of the information content of all concepts C that subsume both C_1 and C_2. Then, the similarity between two words is defined as the maximum of the similarity between any concepts that the words belong to. He used WordNet as the taxonomy; information content is calculated using the Brown corpus.

Li et al. [9] combined structural semantic information from a lexical taxonomy and information content from a corpus in a nonlinear model. They proposed a similarity measure that uses shortest path length, depth, and local density in a taxonomy. Their experiments reported a high Pearson correlation coefficient of 0.8914 on the Miller and Charles [10] benchmark data set. They did not evaluate their method in terms of similarities among named entities. Lin [11] defined the similarity between two concepts as the information that is in common to both concepts and the information contained in each individual concept.

Cilibrasi and Vitanyi [12] proposed a distance metric between words using only page counts retrieved from a web search engine. The proposed metric is named *Normalized Google Distance* (NGD) and is given by

$$NGD(P, Q) = \frac{\max\{\log H(P), \log H(Q)\} - \log H(P, Q)}{\log N - \min\{\log H(P), \log H(Q)\}}.$$

Here, P and Q are the two words between which distance $NGD(P, Q)$ is to be computed, $H(P)$ denotes the page count for the word P, and $H(P, Q)$ is the page count for the query P AND Q. NGD is based on normalized information distance [13], which is defined using Kolmogorov complexity. Because NGD does not take into account the context in which the words co-occur, it suffers from the drawbacks described in the previous section that are characteristic to similarity measures that consider only page counts.

Sahami and Heilman [2] measured semantic similarity between two queries using snippets returned for those queries by a search engine. For each query, they collect snippets from a search engine and represent each snippet as a TF-IDF-weighted term vector. Each vector is L_2 normalized and the centroid of the set of vectors is computed. Semantic similarity between two queries is then defined as the inner product between the corresponding centroid vectors. They did not compare their similarity measure with taxonomy-based similarity measures.

Chen et al. [4] proposed a double-checking model using text snippets returned by a web search engine to compute semantic similarity between words. For two words P and Q, they collect snippets for each word from a web search engine. Then, they count the occurrences of word P in the snippets for word Q and the occurrences of word Q in the snippets for word P. These values are combined nonlinearly to compute the similarity between P and Q. The *Co-occurrence Double-Checking* (CODC) measure is defined as

$$\text{CODC}(P,Q) = \begin{cases} 0, & \text{if } f(P@Q) = 0, \\ \exp\left(\log\left[\frac{f(P@Q)}{H(P)} \times \frac{f(Q@P)}{H(Q)}\right]^{\alpha}\right), & \text{otherwise.} \end{cases}$$

Here, $f(P@Q)$ denotes the number of occurrences of P in the top-ranking snippets for the query Q in Google, $H(P)$ is the page count for query P, and α is a constant in this model, which is experimentally set to the value 0.15. This method depends heavily on the search engine's ranking algorithm. Although two words P and Q might be very similar, we cannot assume that one can find Q in the snippets for P, or vice versa, because a search engine considers many other factors besides semantic similarity, such as publication date (novelty) and link structure (authority) when ranking the result set for a query. This observation is confirmed by the experimental results in their paper which reports zero similarity scores for many pairs of words in the Miller and Charles [10] benchmark data set.

Semantic similarity measures have been used in various applications in natural language processing such as word sense disambiguation [14], language modeling [15], synonym extraction [16], and automatic thesauri extraction [17]. Semantic similarity measures are important in many web-related tasks. In query expansion [18], a user query is modified using synonymous words to improve the relevancy of the search. One method to find appropriate words to include in a query is to compare the previous user queries using semantic similarity measures. If there exists a previous query that is semantically related to the current query, then it can be either suggested to the user, or internally used by the search engine to modify the original query.

3 METHOD

3.1 Outline

Given two words P and Q, we model the problem of measuring the semantic similarity between P and Q, as a one of constructing a function $\text{sim}(P,Q)$ that returns a value in range $[0,1]$. If P and Q are highly similar (e.g., synonyms), we expect $\text{sim}(P,Q)$ to be closer to 1. On the
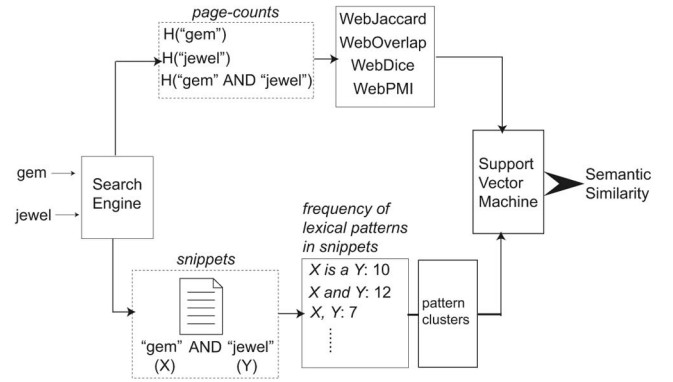
Fig. 2. Outline of the proposed method.

other hand, if P and Q are not semantically similar, then we expect $\text{sim}(P,Q)$ to be closer to 0. We define numerous features that express the similarity between P and Q using page counts and snippets retrieved from a web search engine for the two words. Using this feature representation of words, we train a two-class support vector machine to classify synonymous and nonsynonymous word pairs. The function $\text{sim}(P,Q)$ is then approximated by the confidence score of the trained SVM.

Fig. 2 illustrates an example of using the proposed method to compute the semantic similarity between two words, *gem* and *jewel*. First, we query a web search engine and retrieve page counts for the two words and for their conjunctive (i.e., "*gem*," "*jewel*," and "*gem AND jewel*"). In Section 3.2, we define four similarity scores using page counts. Page counts-based similarity scores consider the global co-occurrences of two words on the web. However, they do not consider the local context in which two words co-occur. On the other hand, snippets returned by a search engine represent the local context in which two words co-occur on the web. Consequently, we find the frequency of numerous lexical syntactic patterns in snippets returned for the conjunctive query of the two words. The lexical patterns we utilize are extracted automatically using the method described in Section 3.3. However, it is noteworthy that a semantic relation can be expressed using more than one lexical pattern. Grouping the different lexical patterns that convey the same semantic relation, enables us to represent a semantic relation between two words accurately. For this purpose, we propose a sequential pattern clustering algorithm in Section 3.4. Both page counts-based similarity scores and lexical pattern clusters are used to define various features that represent the relation between two words. Using this feature representation of word pairs, we train a two-class support vector machine [19] in Section 3.5.

3.2 Page Count-Based Co-Occurrence Measures

Page counts for the query P AND Q can be considered as an approximation of co-occurrence of two words (or multi-word phrases) P and Q on the web. However, page counts for the query P AND Q alone do not accurately express semantic similarity. For example, Google returns 11,300,000 as the page count for "*car*" AND "*automobile*," whereas the same is 49,000,000 for "*car*" AND "*apple*." Although,

automobile is more semantically similar to *car* than *apple* is, page counts for the query "*car*" AND "*apple*" are more than four times greater than those for the query "*car*" AND "*automobile*." One must consider the page counts not just for the query *P* AND *Q*, but also for the individual words *P* and *Q* to assess semantic similarity between *P* and *Q*.

We compute four popular co-occurrence measures; Jaccard, Overlap (Simpson), Dice, and Pointwise mutual information (PMI), to compute semantic similarity using page counts. For the remainder of this paper, we use the notation $H(P)$ to denote the page counts for the query P in a search engine. The *WebJaccard* coefficient between words (or multiword phrases) P and Q, WebJaccard(P, Q), is defined as

$$
\text{WebJaccard}(P, Q) = \begin{cases} 0, & \text{if } H(P \cap Q) \leq c, \\ \dfrac{H(P \cap Q)}{H(P) + H(Q) - H(P \cap Q)}, & \text{otherwise.} \end{cases}
$$

Therein, $P \cap Q$ denotes the conjunction query *P* AND *Q*. Given the scale and noise in web data, it is possible that two words may appear on some pages even though they are not related. In order to reduce the adverse effects attributable to such co-occurrences, we set the WebJaccard coefficient to zero if the page count for the query $P \cap Q$ is less than a threshold c.[2] Similarly, we define *WebOverlap*, WebOverlap(P, Q), as

$$
\text{WebOverlap}(P, Q) = \begin{cases} 0, & \text{if } H(P \cap Q) \leq c, \\ \dfrac{H(P \cap Q)}{\min(H(P), H(Q))}, & \text{otherwise.} \end{cases}
$$

WebOverlap is a natural modification to the Overlap (Simpson) coefficient. We define the *WebDice* coefficient as a variant of the Dice coefficient. WebDice(P, Q) is defined as

$$
\text{WebDice}(P, Q) = \begin{cases} 0, & \text{if } H(P \cap Q) \leq c, \\ \dfrac{2H(P \cap Q)}{H(P) + H(Q)}, & \text{otherwise.} \end{cases} \tag{3}
$$

Pointwise mutual information [20] is a measure that is motivated by information theory; it is intended to reflect the dependence between two probabilistic events. We define *WebPMI* as a variant form of pointwise mutual information using page counts as

$$
\text{WebPMI}(P, Q) = \begin{cases} 0, & \text{if } H(P \cap Q) \leq c, \\ \log_2 \left(\dfrac{\frac{H(P \cap Q)}{N}}{\frac{H(P)}{N} \frac{H(Q)}{N}} \right), & \text{otherwise.} \end{cases}
$$

Here, N is the number of documents indexed by the search engine. Probabilities in (4) are estimated according to the maximum likelihood principle. To calculate PMI accurately using (4), we must know N, the number of documents indexed by the search engine. Although estimating the number of documents indexed by a search engine [21] is an interesting task itself, it is beyond the scope of this work. In

2. We set $c = 5$ in our experiments.

Fig. 3. A snippet retrieved for the query "*cricket*" AND "*sport*."

the present work, we set $N = 10^{10}$ according to the number of indexed pages reported by Google. As previously discussed, page counts are mere approximations to actual word co-occurrences in the web. However, it has been shown empirically that there exists a high correlation between word counts obtained from a web search engine (e.g., Google and Altavista) and that from a corpus (e.g., British National corpus) [22]. Moreover, the approximated page counts have been successfully used to improve a variety of language modeling tasks [23].

3.3 Lexical Pattern Extraction

Page counts-based co-occurrence measures described in Section 3.2 do not consider the local context in which those words co-occur. This can be problematic if one or both words are polysemous, or when page counts are unreliable. On the other hand, the snippets returned by a search engine for the conjunctive query of two words provide useful clues related to the semantic relations that exist between two words. A snippet contains a window of text selected from a document that includes the queried words. Snippets are useful for search because, most of the time, a user can read the snippet and decide whether a particular search result is relevant, without even opening the url. Using snippets as contexts is also computationally efficient because it obviates the need to download the source documents from the web, which can be time consuming if a document is large. For example, consider the snippet in Fig. 3. Here, the phrase *is a* indicates a semantic relationship between **cricket** and **sport**. Many such phrases indicate semantic relationships. For example, *also known as, is a, part of, is an example of* all indicate semantic relations of different types. In the example given above, words indicating the semantic relation between *cricket* and *sport* appear between the query words. Replacing the query words by variables X and Y, we can form the pattern *X is a Y* from the example given above.

Despite the efficiency of using snippets, they pose two main challenges: first, a snippet can be a fragmented sentence, second, a search engine might produce a snippet by selecting multiple text fragments from different portions in a document. Because most syntactic or dependency parsers assume complete sentences as the input, deep parsing of snippets produces incorrect results. Consequently, we propose a shallow lexical pattern extraction algorithm using web snippets, to recognize the semantic relations that exist between two words. Lexical syntactic patterns have been used in various natural language processing tasks such as extracting hypernyms [5], [24], or meronyms [25], question answering [26], and paraphrase extraction [27]. Although a search engine might produce a snippet by selecting multiple text fragments from different portions in a document, a predefined delimiter is used to separate the different fragments. For example, in Google, the delimiter "…" is used to separate different fragments in

> **Ostrich, a large, flightless bird** that lives in the dry grasslands of Africa.

Fig. 4. A snippet retrieved for the query "$ostrich * * * * * bird$."

a snippet. We use such delimiters to split a snippet before we run the proposed lexical pattern extraction algorithm on each fragment.

Given two words P and Q, we query a web search engine using the wildcard query "$P * * * * * Q$" and download snippets. The "$*$" operator matches one word or none in a webpage. Therefore, our wildcard query retrieves snippets in which P and Q appear within a window of seven words. Because a search engine snippet contains ca. 20 words on average, and includes two fragments of texts selected from a document, we assume that the seven word window is sufficient to cover most relations between two words in snippets. In fact, over 95 percent of the lexical patterns extracted by the proposed method contain less than five words. We attempt to approximate the local context of two words using wildcard queries. For example, Fig. 4 shows a snippet retrieved for the query "$ostrich * * * * * bird$."

For a snippet δ, retrieved for a word pair (P, Q), first, we replace the two words P and Q, respectively, with two variables X and Y. We replace all numeric values by D, a marker for digits. Next, we generate all subsequences of words from δ that satisfy all of the following conditions:

1. A subsequence must contain exactly one occurrence of each X and Y.
2. The maximum length of a subsequence is L words.
3. A subsequence is allowed to skip one or more words. However, we do not skip more than g number of words consecutively. Moreover, the total number of words skipped in a subsequence should not exceed G.
4. We expand all negation contractions in a context. For example, *didn't* is expanded to *did not*. We do not skip the word *not* when generating subsequences. For example, this condition ensures that from the snippet X *is not a* Y, we do not produce the subsequence X *is a* Y.

Finally, we count the frequency of all generated subsequences and only use subsequences that occur more than T times as lexical patterns.

The parameters L, g, G, and T are set experimentally, as explained later in Section 3.6. It is noteworthy that the proposed pattern extraction algorithm considers all the words in a snippet, and is *not* limited to extracting patterns only from the mid fix (i.e., the portion of text in a snippet that appears between the queried words). Moreover, the consideration of gaps enables us to capture relations between distant words in a snippet. We use a modified version of the *prefixspan* algorithm [28] to generate subsequences from a text snippet. Specifically, we use the constraints (2-4) to prune the search space of candidate subsequences. For example, if a subsequence has reached the maximum length L, or the number of skipped words is G, then we will not extend it further. By pruning the search

space, we can speed up the pattern generation process. However, none of these modifications affect the accuracy of the proposed semantic similarity measure because the modified version of the prefixspan algorithm still generates the exact set of patterns that we would obtain if we used the original prefixspan algorithm (i.e., without pruning) and subsequently remove patterns that violate the above mentioned constraints. For example, some patterns extracted from the snippet shown in Fig. 4 are: **X**, *a large* **Y**, **X** *a flightless* **Y**, and **X**, *large* **Y** *lives*.

3.4 Lexical Pattern Clustering

Typically, a semantic relation can be expressed using more than one pattern. For example, consider the two distinct patterns, **X** *is a* **Y**, and **X** *is a large* **Y**. Both these patterns indicate that there exists an *is-a* relation between **X** and **Y**. Identifying the different patterns that express the same semantic relation enables us to represent the relation between two words accurately. According to the distributional hypothesis [29], words that occur in the same context have similar meanings. The distributional hypothesis has been used in various related tasks, such as identifying related words [16], and extracting paraphrases [27]. If we consider the word pairs that satisfy (i.e., co-occur with) a particular lexical pattern as the *context* of that lexical pair, then from the distributional hypothesis, it follows that the lexical patterns which are similarly distributed over word pairs must be semantically similar.

We represent a pattern a by a vector **a** of word-pair frequencies. We designate **a**, the *word-pair frequency* vector of pattern a. It is analogous to the *document frequency* vector of a word, as used in information retrieval. The value of the element corresponding to a word pair (P_i, Q_i) in **a**, is the frequency, $f(P_i, Q_i, a)$, that the pattern a occurs with the word pair (P_i, Q_i). As demonstrated later, the proposed pattern extraction algorithm typically extracts a large number of lexical patterns. Clustering algorithms based on pairwise comparisons among all patterns are prohibitively time consuming when the patterns are numerous. Next, we present a sequential clustering algorithm to efficiently cluster the extracted patterns.

Given a set Λ of patterns and a clustering similarity threshold θ, Algorithm 1 returns clusters (of patterns) that express similar semantic relations. First, in Algorithm 1, the function $SORT$ sorts the patterns into descending order of their total occurrences in all word pairs. The total occurrence $\mu(a)$ of a pattern a is the sum of frequencies over all word pairs, and is given by

$$\mu(a) = \sum_i f(P_i, Q_i, a). \tag{5}$$

After sorting, the most common patterns appear at the beginning in Λ, whereas rare patterns (i.e., patterns that occur with only few word pairs) get shifted to the end. Next, in line 2, we initialize the set of clusters, C, to the empty set. The outer for loop (starting at line 3), repeatedly takes a pattern $\mathbf{a_i}$ from the ordered set Λ, and in the inner for loop (starting at line 6), finds the cluster, c^* ($\in C$) that is most similar to $\mathbf{a_i}$. First, we represent a cluster by the centroid of all word-pair frequency vectors corresponding to the patterns in that cluster to compute the similarity

between a pattern and a cluster. Next, we compute the cosine similarity between the cluster centroid ($\mathbf{c_j}$), and the word-pair frequency vector of the pattern ($\mathbf{a_i}$). If the similarity between a pattern $\mathbf{a_i}$, and its most similar cluster, $\mathbf{c}^*$, is greater than the threshold θ, we append $\mathbf{a_i}$ to $\mathbf{c}^*$ (line 14). We use the operator $\oplus$ to denote the vector addition between $\mathbf{c}^*$ and $\mathbf{a_i}$. Then, we form a new cluster $\{\mathbf{a_i}\}$ and append it to the set of clusters, C, if $\mathbf{a_i}$ is not similar to any of the existing clusters beyond the threshold θ.

Algorithm 1. Sequential pattern clustering algorithm.
Input: patterns $\Lambda = \{\mathbf{a_1}, \ldots, \mathbf{a_n}\}$, threshold θ
Output: clusters C
1: SORT(Λ)
2: $C \leftarrow \{\}$
3: **for** pattern $\mathbf{a_i} \in \Lambda$ **do**
4:    $max \leftarrow -\infty$
5:    $\mathbf{c}^* \leftarrow null$
6:    **for** cluster $\mathbf{c_j} \in C$ **do**
7:       $sim \leftarrow \text{cosine}(\mathbf{a_i}, \mathbf{c_j})$
8:       **if** $sim > max$ **then**
9:          $max \leftarrow sim$
10:          $\mathbf{c}^* \leftarrow \mathbf{c_j}$
11:       **end if**
12:    **end for**
13:    **if** $max > \theta$ **then**
14:       $\mathbf{c}^* \leftarrow \mathbf{c}^* \oplus \mathbf{a_i}$
15:    **else**
16:       $C \leftarrow C \cup \{\mathbf{a_i}\}$
17:    **end if**
18: **end for**
19: **return** C

By sorting the lexical patterns in the descending order of their frequency and clustering the most frequent patterns first, we form clusters for more common relations first. This enables us to separate rare patterns which are likely to be outliers from attaching to otherwise clean clusters. The greedy sequential nature of the algorithm avoids pairwise comparisons between all lexical patterns. This is particularly important because when the number of lexical patterns is large as in our experiments (e.g., over 100,000), pairwise comparisons between all patterns are computationally prohibitive. The proposed clustering algorithm attempts to identify the lexical patterns that are similar to each other more than a given threshold value. By adjusting the threshold, we can obtain clusters with different granularity.

The only parameter in Algorithm 1, the similarity threshold, θ, ranges in $[0,1]$. It decides the *purity* of the formed clusters. Setting θ to a high value ensures that the patterns in each cluster are highly similar. However, high θ values also yield numerous clusters (increased model complexity). In Section 3.6, we investigate, experimentally, the effect of θ on the overall performance of the proposed relational similarity measure.

The initial sort operation in Algorithm 1 can be carried out in time complexity of $O(n\log n)$, where n is the number of patterns to be clustered. Next, the sequential assignment of lexical patterns to the clusters requires complexity of $O(n|C|)$, where $|C|$ is the number of clusters. Typically, n is much larger than $|C|$ (i.e., $n \gg |C|$). Therefore, the overall

time complexity of Algorithm 1 is dominated by the sort operation, hence $O(n\log n)$. The sequential nature of the algorithm avoids pairwise comparisons among all patterns. Moreover, sorting the patterns by their total word-pair frequency prior to clustering ensures that the final set of clusters contains the most common relations in the data set.

3.5 Measuring Semantic Similarity

In Section 3.2, we defined four co-occurrence measures using page counts. Moreover, in Sections 3.3 and 3.4, we showed how to extract clusters of lexical patterns from snippets to represent numerous semantic relations that exist between two words. In this section, we describe a machine learning approach to combine both page counts-based co-occurrence measures, and snippets-based lexical pattern clusters to construct a robust semantic similarity measure.

Given N clusters of lexical patterns, first, we represent a pair of words (P,Q) by an $(N+4)$-dimensional feature vector $\mathbf{f}_{PQ}$. The four page counts-based co-occurrence measures defined in Section 3.2 are used as four distinct features in $\mathbf{f}_{PQ}$. For completeness, let us assume that $(N+1)$st, $(N+2)$nd, $(N+3)$rd, and $(N+4)$th features are set, respectively, to WebJaccard, WebOverlap, WebDice, and WebPMI. Next, we compute a feature from each of the N clusters as follows: first, we assign a weight w_{ij} to a pattern a_i that is in a cluster c_j as follows:

$$w_{ij} = \frac{\mu(a_i)}{\sum_{t \in c_j} \mu(t)}. \qquad (6)$$

Here, $\mu(a)$ is the total frequency of a pattern a in all word pairs, and it is given by (5). Because we perform a hard clustering on patterns, a pattern can belong to only one cluster (i.e., $w_{ij} = 0$ for $a_i \notin c_j$). Finally, we compute the value of the jth feature in the feature vector for a word pair (P,Q) as follows:

$$\sum_{a_i \in c_j} w_{ij} f(P,Q,a_i). \qquad (7)$$

The value of the jth feature of the feature vector $\mathbf{f}_{PQ}$ representing a word pair (P,Q) can be seen as the weighted sum of all patterns in cluster c_j that co-occur with words P and Q. We assume all patterns in a cluster to represent a particular semantic relation. Consequently, the jth feature value given by (7) expresses the significance of the semantic relation represented by cluster j for word pair (P,Q). For example, if the weight w_{ij} is set to 1 for all patterns a_i in a cluster c_j, then the jth feature value is simply the sum of frequencies of all patterns in cluster c_j with words P and Q. However, assigning an equal weight to all patterns in a cluster is not desirable in practice because some patterns can contain misspellings and/or can be grammatically incorrect. Equation (6) assigns a weight to a pattern proportionate to its frequency in a cluster. If a pattern has a high frequency in a cluster, then it is likely to be a canonical form of the relation represented by all the patterns in that cluster. Consequently, the weighting scheme described by (6) prefers high frequent patterns in a cluster.

To train a two-class SVM to detect synonymous and nonsynonymous word pairs, we utilize a training data set $S = \{(P_k, Q_k, y_k)\}$ of word pairs. S consists of synonymous

word pairs (positive training instances) and nonsynonymous word pairs (negative training instances). Training data set S is generated automatically from WordNet synsets as described later in Section 3.6. Label $y_k \in \{-1, 1\}$ indicates whether the word pair (P_k, Q_k) is a synonymous word pair (i.e., $y_k = 1$) or a nonsynonymous word pair (i.e., $y_k = -1$). For each word pair in S, we create an $(N + 4)$-dimensional feature vector as described above. To simplify the notation, let us denote the feature vector of a word pair (P_k, Q_k) by $\mathbf{f_k}$. Finally, we train a two-class SVM using the labeled feature vectors.

Once we have trained an SVM using synonymous and nonsynonymous word pairs, we can use it to compute the semantic similarity between two given words. Following the same method, we used to generate feature vectors for training, we create an $(N + 4)$-dimensional feature vector $\mathbf{f}^*$ for a pair of words (P^*, Q^*), between which we must measure semantic similarity. We define the semantic similarity $\text{sim}(P^*, Q^*)$ between P^* and Q^* as the posterior probability, $p(y^* = 1|\mathbf{f}^*)$, that the feature vector $\mathbf{f}^*$ corresponding to the word pair (P^*, Q^*) belongs to the synonymous-words class (i.e., $y^* = 1$). $\text{sim}(P^*, Q^*)$ is given by

$$\text{sim}(P^*, Q^*) = p(y^* = 1|\mathbf{f}^*). \tag{8}$$

Because SVMs are large margin classifiers, the output of an SVM is the distance from the classification hyperplane. The distance $d(\mathbf{f}^*)$ to an instance $\mathbf{f}^*$ from the classification hyperplane is given by

$$d(\mathbf{f}^*) = h(\mathbf{f}^*) + b.$$

Here, b is the bias term and the hyperplane, $h(\mathbf{f}^*)$, is given by

$$h(\mathbf{f}^*) = \sum_i y_k \alpha_k K(\mathbf{f_k}, \mathbf{f}^*).$$

Here, α_k is the Lagrange multiplier corresponding to the support vector $\mathbf{f_k}$,[3] and $K(\mathbf{f_k}, \mathbf{f}^*)$ is the value of the kernel function for a training instance $\mathbf{f_k}$ and the instance to classify, $\mathbf{f}^*$. However, $d(\mathbf{f}^*)$ is not a calibrated posterior probability. Following Platt [30], we use sigmoid functions to convert this uncalibrated distance into a calibrated posterior probability. The probability, $p(y = 1|d(\mathbf{f}))$, is computed using a sigmoid function defined over $d(\mathbf{f})$ as follows:

$$p(y = 1|d(\mathbf{f})) = \frac{1}{1 + \exp(\lambda d(\mathbf{f}) + \mu)}.$$

Here, λ and μ are parameters which are determined by maximizing the likelihood of the training data. Log likelihood of the training data is given by

$$L(\lambda, \mu) = \sum_{k=1}^{N} \log p(y_k|\mathbf{f_k}; \lambda, \mu)$$

$$= \sum_{k=1}^{N} \{t_k \log(p_k) + (1 - t_k) \log(1 - p_k)\}.$$

Here, to simplify the notation, we have used $t_k = (y_k + 1)/2$ and $p_k = p(y_k = 1|\mathbf{f_k})$. The maximization in (9) with respect to parameters λ and μ is performed using model-trust minimization [31].

---

3. From K.K.T. conditions, it follows that the Lagrange multipliers corresponding to nonsupport vectors become zero.

TABLE 1
Number of Patterns Extracted for Training Data

word pairs	synonymous	non-synonymous
# word pairs	3000	3000
# extracted patterns	5365398	515848
# selected patterns	270762	38978

3.6 Training

To train the two-class SVM described in Section 3.5, we require both synonymous and nonsynonymous word pairs. We use WordNet, a manually created English dictionary, to generate the training data required by the proposed method. For each sense of a word, a set of synonymous words is listed in WordNet synsets. We randomly select 3,000 nouns from WordNet, and extract a pair of synonymous words from a synset of each selected noun. If a selected noun is polysemous, then we consider the synset for the dominant sense. Obtaining a set of nonsynonymous word pairs (negative training instances) is difficult, because there does not exist a large collection of manually created nonsynonymous word pairs. Consequently, to create a set of nonsynonymous word pairs, we adopt a random shuffling technique. Specifically, we first randomly select two synonymous word pairs from the set of synonymous word pairs created above, and exchange two words between word pairs to create two new word pairs. For example, from two synonymous word pairs (A, B) and (C, D), we generate two new pairs (A, C) and (B, D). If the newly created word pairs do not appear in any of the word net synsets, we select them as nonsynonymous word pairs. We repeat this process until we create 3,000 nonsynonymous word pairs. Our final training data set contains 6,000 word pairs (i.e., 3,000 synonymous word pairs and 3,000 nonsynonymous word pairs).

Next, we use the lexical pattern extraction algorithm described in Section 3.3 to extract numerous lexical patterns for the word pairs in our training data set. We experimentally set the parameters in the pattern extraction algorithm to $L = 5$, $g = 2$, $G = 4$, and $T = 5$. Table 1 shows the number of patterns extracted for synonymous and nonsynonymous word pairs in the training data set. As can be seen from Table 1, the proposed pattern extraction algorithm typically extracts a large number of lexical patterns. Figs. 5 and 6, respectively, show the distribution of patterns extracted for synonymous and nonsynonymous word pairs. Because of the noise in web snippets such as, ill-formed snippets and misspells, most patterns occur only a few times in the list of extracted patterns. Consequently, we ignore any patterns that occur less than five times. Finally, we deduplicate the patterns that appear for both synonymous and nonsynonymous word pairs to create a final set of 3,02,286 lexical patterns. The remainder of the experiments described in the paper use this set of lexical patterns.

We determine the clustering threshold θ as follows: first, we run Algorithm 1 for different θ values, and with each set of clusters, we compute feature vectors for synonymous word pairs as described in Section 3.5. Let W denote the set of synonymous word pairs (i.e., $W = \{(P_i, Q_i)|(P_i, Q_i, y_i) \in S, y_i = 1\}$). Moreover, let $\mathbf{f_W}$ be the centroid vector of all
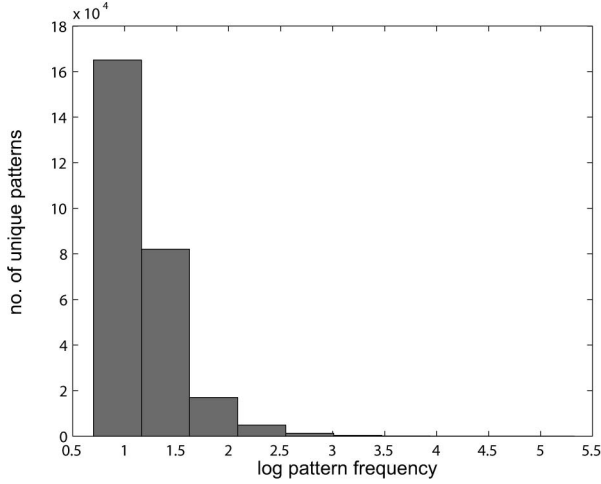
Fig. 5. Distribution of patterns extracted from synonymous word pairs.

feature vectors representing synonymous word pairs, which is given by

$$\mathbf{f_W} = \frac{1}{|W|} \sum_{(P,Q) \in W} \mathbf{f_{PQ}}. \qquad (10)$$

Next, we compute the average Mahalanobis distance, $D(\theta)$, between $\mathbf{f_W}$ and feature vectors that represent synonymous as follows:

$$D(\theta) = \frac{1}{|W|} \sum_{(P,Q) \in W} \text{Mahala}(\mathbf{f_W}, \mathbf{f_{PQ}}). \qquad (11)$$

Here, $|W|$ is the number of word pairs in W, and mahala$(\mathbf{f_w}, \mathbf{f_{PQ}})$ is the Mahalanobis distance defined by

$$\text{mahala}(\mathbf{f_w}, \mathbf{f_{PQ}}) = (\mathbf{f_w} - \mathbf{f_{PQ}})^{\mathrm{T}} C^{-1} (\mathbf{f_w} - \mathbf{f_{PQ}}). \qquad (12)$$

Here, C^{-1} is the inverse of the intercluster correlation matrix, C, where the (i, j) element of C is defined to be the inner product between the vectors $\mathbf{c_i}$, $\mathbf{c_j}$ corresponding to clusters c_i and c_j. Finally, we set the optimum value of
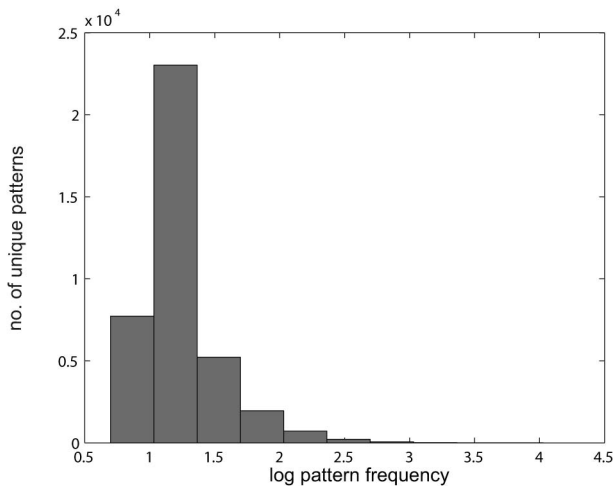
Fig. 6. Distribution of patterns extracted from nonsynonymous word pairs.
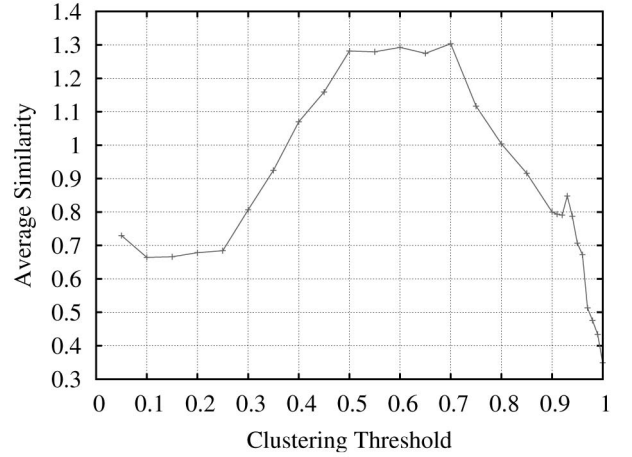
Fig. 7. Average similarity versus clustering threshold θ.

clustering threshold, $\hat{\theta}$, to the value of θ that minimizes the average Mahalanobis distance as follows:

$$\hat{\theta} = \underset{\theta \in [0,1]}{\arg\min} \, D(\theta).$$

Alternatively, we can define the reciprocal of $D(\theta)$ as average similarity, and minimize this quantity. Note that the average in (11) is taken over a large number of synonymous word pairs (3,000 word pairs in W), which enables us to determine θ robustly. Moreover, we consider Mahalanobis distance instead of euclidean distances, because a set of pattern clusters might not necessarily be independent. For example, we would expect a certain level of correlation between the two clusters that represent an *is-a* relation and a *has-a* relation. Mahalanobis distance considers the correlation between clusters when computing distance. Note that if we take the identity matrix as C in (12), then we get the euclidean distance.

Fig. 7 plots average similarity between centroid feature vector and all synonymous word pairs for different values of θ. From Fig. 7, we see that initially average similarity increases when θ is increased. This is because clustering of semantically related patterns reduces the sparseness in feature vectors. Average similarity is stable within a range of θ values between 0.5 and 0.7. However, increasing θ beyond 0.7 results in a rapid drop of average similarity. To explain this behavior, consider Fig. 8 where we plot the sparsity of the set of clusters (i.e., the ratio between singletons to total clusters) against threshold θ. As seen from Fig. 8, high θ values result in a high percentage of singletons because only highly similar patterns will form clusters. Consequently, feature vectors for different word pairs do not have many features in common. The maximum average similarity score of 1.31 is obtained with $\theta = 0.7$, corresponding to 32,207 total clusters out of which 23,836 are singletons with exactly one pattern (sparsity = 0.74). For the remainder of the experiments in this paper, we set θ to this optimal value and use the corresponding set of clusters.

We train an SVM with a radial basis function (RBF) kernel. Kernel parameter γ and soft-margin trade-off C are, respectively, set to 0.0078125 and 1,31,072 using fivefold

Fig. 8. Sparsity versus clustering threshold θ.

cross validation on training data. We used LibSVM[4] as the SVM implementation. Remainder of the experiments in the paper use this trained SVM model.

4 EXPERIMENTS

4.1 Benchmark Data Sets

Following the previous work, we evaluate the proposed semantic similarity measure by comparing it with human ratings in three benchmark data sets: Miller-Charles (MC) [10], Rubenstein-Goodenough (RG) [32], and WordSimilarity-353 (WS) [33]. Each data set contains a list of word pairs rated by multiple human annotators (MC: 28 pairs, 38 annotators, RG: 65 pairs, 36 annotators, and WS: 353 pairs, 13 annotators). A semantic similarity measure is evaluated using the correlation between the similarity scores produced by it for the word pairs in a benchmark data set and the human ratings. Both Pearson correlation coefficient and Spearman correlation coefficient have been used as evaluation measures in previous work on semantic similarity. It is noteworthy that Pearson correlation coefficient can get severely affected by nonlinearities in ratings. Contrastingly, Spearman correlation coefficient first assigns ranks to each list of scores, and then computes correlation between the two lists of ranks. Therefore, Spearman correlation is more appropriate for evaluating semantic similarity measures, which might not be necessarily linear. In fact, as we shall see later, most semantic similarity measures are nonlinear. Previous work that used RG and WS data sets in their evaluations has chosen Spearman correlation coefficient as the evaluation measure. However, for the MC data set, which contains only 28 word pairs, Pearson correlation coefficient has been widely used. To be able to compare our results with previous work, we use both Pearson and Spearman correlation coefficients for experiments conducted on MC data set, and Spearman correlation coefficient for experiments on RG and WS data sets. It is noteworthy that we exclude all words that appear in the benchmark data set from the training data created from WordNet as described in Section 3.6. Benchmark data sets are reserved

4. http://www.csie.ntu.edu.tw/cjlin/libsvm/.

for evaluation purposes only and we do not train or tune any of the parameters using benchmark data sets.

4.2 Semantic Similarity

Table 2 shows the experimental results on MC data set for the proposed method (**Proposed**); previously proposed web-based semantic similarity measures: Sahami and Heilman [2] (**SH**), Co-occurrence double-checking model [4], and Normalized Google Distance [12]; and the four page counts-based co-occurrence measures in Section 3.2. **No Clust** baseline , which resembles our previously published work [34], is identical to the proposed method in all aspects except for that it does not use cluster information. It can be understood as the proposed method with each extracted lexical pattern in its own cluster. No Clust is expected to show the effect of clustering on the performance of the proposed method. All similarity scores in Table 2 are normalized into $[0, 1]$ range for the ease of comparison, and Fisher's confidence intervals are computed for Spearman and Pearson correlation coefficients. NGD is a distance measure and was converted to a similarity score by taking the inverse. Original papers that proposed NGD and SH measures did not present their results on MC data set. Therefore, we reimplemented those systems following the original papers. The proposed method achieves the highest Pearson and Spearman coefficients in Table 2 and outperforms all other web-based semantic similarity measures. WebPMI reports the highest correlation among all page counts-based co-occurrence measures. However, methods that use snippets such as SH and CODC, have better correlation scores. MC data set contains polysemous words such as *father* (priest versus parent), *oracle* (priest versus database system), and *crane* (machine versus bird), which are problematic for page counts-based measures that do not consider the local context of a word. The No Clust baseline which combines both page counts and snippets outperforms the CODC measure by a wide margin of 0.2 points. Moreover, by clustering the lexical patterns, we can further improve the No Clust baseline.

Table 3 summarizes the experimental results on RG and WS data sets. Likewise on the MC data set, the proposed method outperforms all other methods on RG and WS data sets. In contrast to MC data set, the proposed method outperforms the No Clust baseline by a wide margin in RG and WS data sets. Unlike the MC data set which contains only 28 word pairs, RG and WS data sets contain a large number of word pairs. Therefore, more reliable statistics can be computed on RG and WS data sets. Fig. 9 shows the similarity scores produced by six methods against human ratings in the WS data set. We see that all methods deviate from the $y = x$ line, and are not linear. We believe this justifies the use of Spearman correlation instead of Pearson correlation by previous work on semantic similarity as the preferred evaluation measure.

Tables 4, 5, and 6, respectively, compare the proposed method against previously proposed semantic similarity measures. Despite the fact that the proposed method does not require manually created resources such as WordNet, Wikipedia, or fixed corpora, the performance of the proposed method is comparable with methods that use such resources. The nondependence on dictionaries is

TABLE 2
Semantic Similarity Scores on MC Data Set

word pair	MC	WebJaccard	WebDice	WebOverlap	WebPMI	CODC [4]	SH [2]	NGD [12]	No Clust	Proposed
automobile-car	1.00	0.65	0.66	0.83	0.43	0.69	1.00	0.15	0.98	0.92
journey-voyage	0.98	0.41	0.42	0.16	0.47	0.42	0.52	0.39	1.00	1.00
gem-jewel	0.98	0.29	0.30	0.07	0.69	1.00	0.21	0.42	0.69	0.82
boy-lad	0.96	0.18	0.19	0.59	0.63	0.00	0.47	0.12	0.97	0.96
coast-shore	0.94	0.78	0.79	0.51	0.56	0.52	0.38	0.52	0.95	0.97
asylum-madhouse	0.92	0.01	0.01	0.08	0.81	0.00	0.21	1.00	0.77	0.79
magician-wizard	0.89	0.29	0.30	0.37	0.86	0.67	0.23	0.44	1.00	1.00
midday-noon	0.87	0.10	0.10	0.12	0.59	0.86	0.29	0.74	0.82	0.99
furnace-stove	0.79	0.39	0.41	0.10	1.00	0.93	0.31	0.61	0.89	0.88
food-fruit	0.78	0.75	0.76	1.00	0.45	0.34	0.18	0.55	1.00	0.94
bird-cock	0.77	0.14	0.15	0.14	0.43	0.50	0.06	0.41	0.59	0.87
bird-crane	0.75	0.23	0.24	0.21	0.52	0.00	0.22	0.41	0.88	0.85
implement-tool	0.75	1.00	1.00	0.51	0.30	0.42	0.42	0.91	0.68	0.50
brother-monk	0.71	0.25	0.27	0.33	0.62	0.55	0.27	0.23	0.38	0.27
crane-implement	0.42	0.06	0.06	0.10	0.19	0.00	0.15	0.40	0.13	0.06
brother-lad	0.41	0.18	0.19	0.36	0.64	0.38	0.24	0.26	0.34	0.13
car-journey	0.28	0.44	0.45	0.36	0.20	0.29	0.19	0.00	0.29	0.17
monk-oracle	0.27	0.00	0.00	0.00	0.00	0.00	0.05	0.45	0.33	0.80
food-rooster	0.21	0.00	0.00	0.41	0.21	0.00	0.08	0.42	0.06	0.02
coast-hill	0.21	0.96	0.97	0.26	0.35	0.00	0.29	0.70	0.87	0.36
forest-graveyard	0.20	0.06	0.06	0.23	0.49	0.00	0.00	0.54	0.55	0.44
monk-slave	0.12	0.17	0.18	0.05	0.61	0.00	0.10	0.77	0.38	0.24
coast-forest	0.09	0.86	0.87	0.29	0.42	0.00	0.25	0.36	0.41	0.15
lad-wizard	0.09	0.06	0.07	0.05	0.43	0.00	0.15	0.66	0.22	0.23
cord-smile	0.01	0.09	0.10	0.02	0.21	0.00	0.09	0.13	0.00	0.01
glass-magician	0.01	0.11	0.11	0.40	0.60	0.00	0.14	0.21	0.18	0.05
rooster-voyage	0.00	0.00	0.00	0.00	0.23	0.00	0.20	0.21	0.02	0.05
noon-string	0.00	0.12	0.12	0.04	0.10	0.00	0.08	0.21	0.02	0.00
Spearman	1.00	0.39	0.39	0.40	0.52	0.69	0.62	0.13	0.83	0.85
Lower	1.00	0.02	0.02	0.04	0.18	0.42	0.33	−0.25	0.66	0.69
Upper	1.00	0.67	0.67	0.68	0.75	0.84	0.81	0.48	0.92	0.93
Pearson	1.00	0.26	0.27	0.38	0.55	0.69	0.58	0.21	0.83	0.87
Lower	1.00	−0.13	−0.12	0.01	0.22	0.42	0.26	−0.18	0.67	0.73
Upper	1.00	0.58	0.58	0.66	0.77	0.85	0.78	0.54	0.92	0.94

particularly attractive when measuring the similarity between named entities which are not well covered by dictionaries such as WordNet. We further evaluate the ability of the proposed method to compute the semantic similarity between named entities in Section 4.3.

In Table 7, we analyze the effect of clustering. We compare No Clust (i.e., does not use any clustering information in feature vector creation), singletons excluded (remove all clusters with only one pattern), and singletons included (considering all clusters). From Table 7, we see in all three data sets, we obtain the best results by considering all clusters (singletons incl.). If we remove all singletons, then the performance drops below No Clust. Note that out of the 32,207 clusters used by the proposed method, 23,836 are singletons (sparsity = 0.74). Therefore, if we remove all

singletons, we cannot represent some word pairs adequately, resulting in poor performance.

Table 8 shows the contribution of page counts-based similarity measures, and lexical patterns extracted from snippets, on the overall performance of the proposed method. To evaluate the effect of page counts-based co-occurrence measures on the proposed method, we generate feature vectors only using the four page counts-based co-occurrence measures, to train an SVM. Similarly, to evaluate the effect of snippets, we generate feature vectors only using lexical pattern clusters. From Table 8, we see that on all three data sets, snippets have a greater impact on the performance of the proposed method than page counts. By considering both page counts as well as snippets, we can further improve the performance reported by individual methods. The improvement in performance when we use snippets only is statistically significant over that when we use page counts only in RG and WS data sets. However, the performance gain in the combination is not statistically significant. We believe that this is because most of the words in the benchmark data sets are common nouns that co-occur a lot in web snippets. On the other hand, having page counts in the model is particularly useful when two words do not appear in numerous lexical patterns.

4.3 Community Mining

Measuring the semantic similarity between named entities is vital in many applications such as query expansion [2], entity disambiguation (e.g., namesake disambiguation), and

TABLE 3
Correlation with RG and WS Data Sets

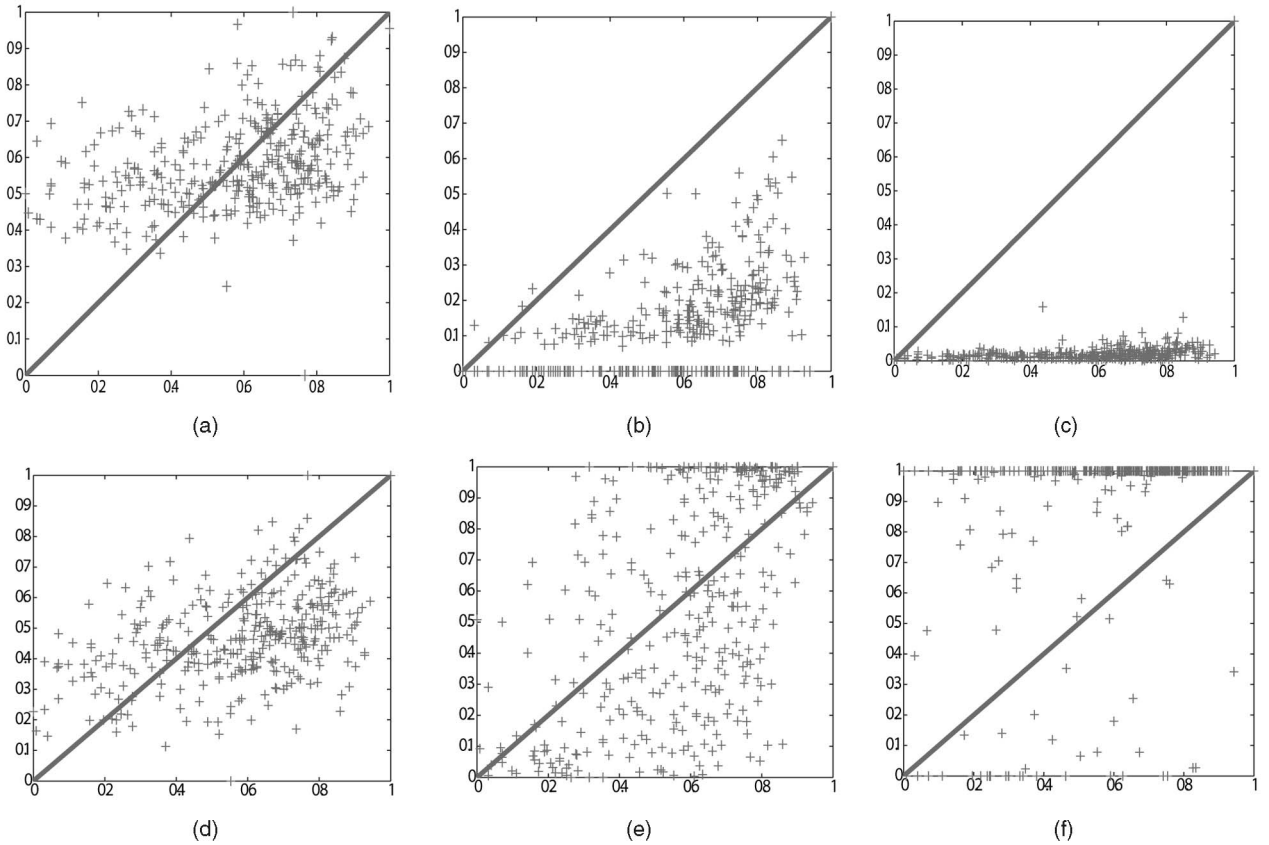| Method	WS	RG
WebJaccard	0.26 [0.16, 0.35]	0.51 [0.30, 0.67]
WebDice	0.26 [0.16, 0.35]	0.51 [0.30, 0.67]
WebOverlap	0.27 [0.17, 0.36]	0.54 [0.34, 0.69]
WebPMI	0.36 [0.26, 0.45]	0.49 [0.28, 0.66]
CODC [4]	0.55 [0.48, 0.62]	0.65 [0.49, 0.77]
SH [2]	0.36 [0.26, 0.45]	0.31 [0.07, 0.52]
NGD [12]	0.40 [0.31, 0.48]	0.56 [0.37, 0.71]
No Clust	0.53 [0.45, 0.60]	0.73 [0.60, 0.83]
Proposed	0.74 [0.69, 0.78]	0.86 [0.78, 0.91]

Fig. 9. Spearman rank correlation between various similarity measures (y-axis) and human ratings (x-axis) on the WS data set.

community mining [46]. Because most named entities are not covered by WordNet, similarity measures that are based on WordNet cannot be used directly in these tasks. Unlike common English words, named entities are being created constantly. Manually maintaining an up-to-date taxonomy of named entities is costly, if not impossible. The proposed semantic similarity measure is appealing for these applications because it does not require precompiled taxonomies.

In order to evaluate the performance of the proposed measure in capturing the semantic similarity between named entities, we set up a community mining task. We select 50 personal names from five communities: *tennis players*, *golfers*, *actors*, *politicians*, and *scientists*,[5] (10 names from each community) from the open directory project (DMOZ).[6] For each pair of names in our data set, we measure their similarity using the proposed method and baselines. We use group-average agglomerative hierarchical clustering (GAAC) to cluster the names in our data set into five clusters.

Initially, each name is assigned to a separate cluster. In subsequent iterations, group-average agglomerative clustering process, merges the two clusters with highest correlation. Correlation, $\text{Corr}(\Gamma)$ between two clusters A and B is defined as the following:

$$\text{Corr}(\Gamma) = \frac{1}{2} \frac{1}{|\Gamma|(|\Gamma| - 1)} \sum_{(u,v) \in \Gamma} \text{sim}(u, v).$$

Here, Γ is the merger of the two clusters A and B. $|\Gamma|$ denotes the number of elements (persons) in Γ and $\text{sim}(u, v)$ is the semantic similarity between two persons u and v in Γ. We terminate GAAC process when exactly five clusters are formed. We adopt this clustering method with different semantic similarity measures $\text{sim}(u, v)$ to compare their accuracy in clustering people who belong to the same community.

We employed the *B-CUBED* metric [47] to evaluate the clustering results. The B-CUBED evaluation metric was originally proposed for evaluating cross-document coreference chains. It does not require the clusters to be labeled.

TABLE 4
Comparison with Previous Work on MC Data Set

Method	Source	Pearson
Wikirelate! [35]	Wikipedia	0.46
Sahami & Heilman [2]	Web Snippets	0.58
Gledson [36]	Page Counts	0.55
Wu & Palmer [37]	WordNet	0.78
Resnik [8]	WordNet	0.74
Leacock [38]	WordNet	0.82
Lin [11]	WordNet	0.82
Jiang & Conrath [39]	WordNet	0.84
Jarmasz [40]	Roget's	0.87
Li et al. [9]	WordNet	0.89
Schickel-Zuber [41]	WordNet	0.91
Agirre et al [42]	WordNet+Corpus	0.93
Proposed	WebSnippets+Page Counts	0.87

5. www.miv.t.u-tokyo.ac.jp/danushka/data/people.tgz.
6. http://dmoz.org.

TABLE 5
Comparison with Previous Work on RG Data Set

Method	Source	Spearman
Wikirelate! [35]	Wikipedia	0.56
Gledson [36]	Page Counts	0.55
Jiang & Conrath [39]	WordNet	0.73
Hirst & St. Onge [43]	WordNet	0.73
Resnik [8]	WordNet	0.80
Lin [11]	WordNet	0.83
Leacock [38]	WordNet	0.85
Proposed	WebSnippets+Page Counts	0.86

TABLE 6
Comparison with Previous Work on WS Data Set

Method	Source	Spearman
Jarmasz [40]	WordNet	0.35
Wikirelate! [35]	Wikipedia	0.48
Jarmasz [40]	Roget's	0.55
Hughes & Ramage [44]	WordNet	0.55
Finkelstein et al. [33]	Corpus+WordNet	0.56
Gabrilovich [45]	ODP	0.65
Gabrilovich [45]	Wikipedia	0.75
Proposed	WebSnippets+Page Counts	0.74

TABLE 7
Effect of Pattern Clustering (Spearman)

Method	MC	RG	WS
No Clust	0.83	0.73	0.53
[upper, lower]	[0.66, 0.92]	[0.60, 0.83]	[0.45, 0.60]
singletons excl.	0.59	0.68	0.49
[upper, lower]	[0.29, 0.79]	[0.52, 0.79]	[0.40, 0.56]
singletons incl.	0.85	0.86	0.74
[upper, lower]	[0.69, 0.93]	[0.78, 0.91]	[0.69, 0.78]

We compute precision, recall, and F-score for each name in the data set and average the results over the data set. For each person p in our data set, let us denote the cluster that p belongs to by $C(p)$. Moreover, we use $A(p)$ to denote the affiliation of person p, e.g., $A(\text{"Tiger Woods"}) = \text{"Tennis Player."}$ Then, we calculate precision and recall for person p as

$$\text{Precision}(p) = \frac{\text{No. of people in C(p) with affiliation A(p)}}{\text{No. of people in C(p)}},$$

$$\text{Recall}(p) = \frac{\text{No. of people in C(p) with affiliation A(p)}}{\text{Total No. of people with affiliation A(p)}}.$$

Since, we selected 10 people from each of the five categories, the total number of people with a particular affiliation is 10 for all the names p. Then, the F-score of person p is defined as

$$\text{F}(p) = \frac{2 \times \text{Precision }(p) \times \text{Recall}(p)}{\text{Precision}(p) + \text{Recall}(p)}.$$

Overall precision, recall, and F-score are computed by taking the averaged sum over all the names in the data set.

$$\text{Precision} = \frac{1}{N} \sum_{p \in Data\ Set} \text{Precision}(p)$$

TABLE 8
Page Counts versus Snippets (Spearman)

Method	MC	RG	WS
Page counts only	0.57	0.57	0.37
[upper, lower]	[0.25, 0.77]	[0.39, 0.72]	[0.29, 0.46]
Snippets only	0.82	0.85	0.72
[upper, lower]	[0.65, 0.91]	[0.77, 0.90]	[0.71, 0.79]
Both	0.85	0.86	0.74
[upper, lower]	[0.69, 0.93]	[0.78, 0.91]	[0.69, 0.78]

TABLE 9
Results for Community Mining

Method	Precision	Recall	F Measure
WebJaccard	0.59	0.71	0.61
WebOverlap	0.59	0.68	0.59
WebDice	0.58	0.71	0.61
WebPMI	0.26	0.42	0.29
Sahami [2]	0.63	0.66	0.64
Chen [4]	0.47	0.62	0.49
No Clust	0.79	0.80	0.78
Proposed	0.85	0.87	0.86

$$\text{Recall} = \frac{1}{N} \sum_{p \in Data\ Set} \text{Recall}(p)$$

$$F\text{-Score} = \frac{1}{N} \sum_{p \in Data\ Set} \text{F}(p).$$

Here, $Data\ Set$ is the set of 50 names selected from the open directory project. Therefore, $N = 50$ in our evaluations.

Experimental results are shown in Table 9. The proposed method shows the highest entity clustering accuracy in Table 9 with a statistically significant ($p \leq 0.01$ Tukey HSD) F-score of 0.86. Sahami and Heilman's [2] snippet-based similarity measure, WebJaccard, WebDice, and WebOverlap measures yield similar clustering accuracies. By clustering semantically related lexical patterns, we see that both precision as well as recall can be improved in a community mining task.

5 CONCLUSION

We proposed a semantic similarity measure using both page counts and snippets retrieved from a web search engine for two words. Four word co-occurrence measures were computed using page counts. We proposed a lexical pattern extraction algorithm to extract numerous semantic relations that exist between two words. Moreover, a sequential pattern clustering algorithm was proposed to identify different lexical patterns that describe the same semantic relation. Both page counts-based co-occurrence measures and lexical pattern clusters were used to define features for a word pair. A two-class SVM was trained using those features extracted for synonymous and nonsynonymous word pairs selected from WordNet synsets. Experimental results on three benchmark data sets showed that the proposed method outperforms various baselines as well as previously proposed web-based semantic similarity measures, achieving a high correlation with human ratings. Moreover, the proposed method improved the F-score in a community mining

task, thereby underlining its usefulness in real-world tasks, that include named entities not adequately covered by manually created resources.

REFERENCES

[1] A. Kilgarriff, "Googleology Is Bad Science," *Computational Linguistics,* vol. 33, pp. 147-151, 2007.

[2] M. Sahami and T. Heilman, "A Web-Based Kernel Function for Measuring the Similarity of Short Text Snippets," *Proc. 15th Int'l World Wide Web Conf.,* 2006.

[3] D. Bollegala, Y. Matsuo, and M. Ishizuka, "Disambiguating Personal Names on the Web Using Automatically Extracted Key Phrases," *Proc. 17th European Conf. Artificial Intelligence,* pp. 553-557, 2006.

[4] H. Chen, M. Lin, and Y. Wei, "Novel Association Measures Using Web Search with Double Checking," *Proc. 21st Int'l Conf. Computational Linguistics and 44th Ann. Meeting of the Assoc. for Computational Linguistics (COLING/ACL '06),* pp. 1009-1016, 2006.

[5] M. Hearst, "Automatic Acquisition of Hyponyms from Large Text Corpora," *Proc. 14th Conf. Computational Linguistics (COLING),* pp. 539-545, 1992.

[6] M. Pasca, D. Lin, J. Bigham, A. Lifchits, and A. Jain, "Organizing and Searching the World Wide Web of Facts - Step One: The One-Million Fact Extraction Challenge," *Proc. Nat'l Conf. Artificial Intelligence (AAAI '06),* 2006.

[7] R. Rada, H. Mili, E. Bichnell, and M. Blettner, "Development and Application of a Metric on Semantic Nets," *IEEE Trans. Systems, Man and Cybernetics,* vol. 19, no. 1, pp. 17-30, Jan./Feb. 1989.

[8] P. Resnik, "Using Information Content to Evaluate Semantic Similarity in a Taxonomy," *Proc. 14th Int'l Joint Conf. Aritificial Intelligence,* 1995.

[9] D. Mclean, Y. Li, and Z.A. Bandar, "An Approach for Measuring Semantic Similarity between Words Using Multiple Information Sources," *IEEE Trans. Knowledge and Data Eng.,* vol. 15, no. 4, pp. 871-882, July/Aug. 2003.

[10] G. Miller and W. Charles, "Contextual Correlates of Semantic Similarity," *Language and Cognitive Processes,* vol. 6, no. 1, pp. 1-28, 1998.

[11] D. Lin, "An Information-Theoretic Definition of Similarity," *Proc. 15th Int'l Conf. Machine Learning (ICML),* pp. 296-304, 1998.

[12] R. Cilibrasi and P. Vitanyi, "The Google Similarity Distance," *IEEE Trans. Knowledge and Data Eng.,* vol. 19, no. 3, pp. 370-383, Mar. 2007.

[13] M. Li, X. Chen, X. Li, B. Ma, and P. Vitanyi, "The Similarity Metric," *IEEE Trans. Information Theory,* vol. 50, no. 12, pp. 3250-3264, Dec. 2004.

[14] P. Resnik, "Semantic Similarity in a Taxonomy: An Information Based Measure and Its Application to Problems of Ambiguity in Natural Language," *J. Artificial Intelligence Research,* vol. 11, pp. 95-130, 1999.

[15] R. Rosenfield, "A Maximum Entropy Approach to Adaptive Statistical Modelling," *Computer Speech and Language,* vol. 10, pp. 187-228, 1996.

[16] D. Lin, "Automatic Retrieval and Clustering of Similar Words," *Proc. 17th Int'l Conf. Computational Linguistics (COLING),* pp. 768-774, 1998.

[17] J. Curran, "Ensemble Methods for Automatic Thesaurus Extraction," *Proc. ACL-02 Conf. Empirical Methods in Natural Language Processing (EMNLP),* 2002.

[18] C. Buckley, G. Salton, J. Allan, and A. Singhal, "Automatic Query Expansion Using Smart: Trec 3," *Proc. Third Text REtreival Conf.,* pp. 69-80, 1994.

[19] V. Vapnik, *Statistical Learning Theory.* Wiley, 1998.

[20] K. Church and P. Hanks, "Word Association Norms, Mutual Information and Lexicography," *Computational Linguistics,* vol. 16, pp. 22-29, 1991.

[21] Z. Bar-Yossef and M. Gurevich, "Random Sampling from a Search Engine's Index," *Proc. 15th Int'l World Wide Web Conf.,* 2006.

[22] F. Keller and M. Lapata, "Using the Web to Obtain Frequencies for Unseen Bigrams," *Computational Linguistics,* vol. 29, no. 3, pp. 459-484, 2003.

[23] M. Lapata and F. Keller, "Web-Based Models for Natural Language Processing," *ACM Trans. Speech and Language Processing,* vol. 2, no. 1, pp. 1-31, 2005.

[24] R. Snow, D. Jurafsky, and A. Ng, "Learning Syntactic Patterns for Automatic Hypernym Discovery," *Proc. Advances in Neural Information Processing Systems (NIPS),* pp. 1297-1304, 2005.

[25] M. Berland and E. Charniak, "Finding Parts in Very Large Corpora," *Proc. Ann. Meeting of the Assoc. for Computational Linguistics on Computational Linguistics (ACL '99),* pp. 57-64, 1999.

[26] D. Ravichandran and E. Hovy, "Learning Surface Text Patterns for a Question Answering System," *Proc. Ann. Meeting on Assoc. for Computational Linguistics (ACL '02),* pp. 41-47, 2001.

[27] R. Bhagat and D. Ravichandran, "Large Scale Acquisition of Paraphrases for Learning Surface Patterns," *Proc. Assoc. for Computational Linguistics: Human Language Technologies (ACL '08: HLT),* pp. 674-682, 2008.

[28] J. Pei, J. Han, B. Mortazavi-Asi, J. Wang, H. Pinto, Q. Chen, U. Dayal, and M. Hsu, "Mining Sequential Patterns by Pattern-Growth: The Prefixspan Approach," *IEEE Trans. Knowledge and Data Eng.,* vol. 16, no. 11, pp. 1424-1440, Nov. 2004.

[29] Z. Harris, "Distributional Structure," *Word,* vol. 10, pp. 146-162, 1954.

[30] J. Platt, "Probabilistic Outputs for Support Vector Machines and Comparison to Regularized Likelihood Methods," *Advances in Large Margin Classifiers,* pp. 61-74, MIT Press, 2000.

[31] P. Gill, W. Murray, and M. Wright, *Practical Optimization.* Academic Press, 1981.

[32] H. Rubenstein and J. Goodenough, "Contextual Correlates of Synonymy," *Comm. ACM,* vol. 8, pp. 627-633, 1965.

[33] L. Finkelstein, E. Gabrilovich, Y. Matias, E. Rivlin, Z. Solan, G. Wolfman, and E. Ruppin, "Placing Search in Context: The Concept Revisited," *ACM Trans. Information Systems,* vol. 20, pp. 116-131, 2002.

[34] D. Bollegala, Y. Matsuo, and M. Ishizuka, "Measuring Semantic Similarity between Words Using Web Search Engines," *Proc. Int'l Conf. World Wide Web (WWW '07),* pp. 757-766, 2007.

[35] M. Strube and S.P. Ponzetto, "Wikirelate! Computing Semantic Relatedness Using Wikipedia," *Proc. Nat'l Conf. Artificial Intelligence (AAAI '06),* pp. 1419-1424, 2006.

[36] A. Gledson and J. Keane, "Using Web-Search Results to Measure Word-Group Similarity," *Proc. Int'l Conf. Computational Linguistics (COLING '08),* pp. 281-288, 2008.

[37] Z. Wu and M. Palmer, "Verb Semantics and Lexical Selection," *Proc. Ann. Meeting on Assoc. for Computational Linguistics (ACL '94),* pp. 133-138, 1994.

[38] C. Leacock and M. Chodorow, "Combining Local Context and Wordnet Similarity for Word Sense Disambiguation," *WordNet: An Electronic Lexical Database,* vol. 49, pp. 265-283, MIT Press, 1998.

[39] J. Jiang and D. Conrath, "Semantic Similarity Based on Corpus Statistics and Lexical Taxonomy," *Proc. Int'l Conf. Research in Computational Linguistics (ROCLING X),* 1997.

[40] M. Jarmasz, "Roget's Thesaurus as a Lexical Resource for Natural Language Processing," technical report, Univ. of Ottawa, 2003.

[41] V. Schickel-Zuber and B. Faltings, "OSS: A Semantic Similarity Function Based on Hierarchical Ontologies," *Proc. Int'l Joint Conf. Artificial Intelligence (IJCAI '07),* pp. 551-556, 2007.

[42] E. Agirre, E. Alfonseca, K. Hall, J. Kravalova, M. Pasca, and A. Soroa, "A Study on Similarity and Relatedness Using Distributional and Wordnet-Based Approaches," *Proc. Human Language Technologies: The 2009 Ann. Conf. North Am. Chapter of the Assoc. for Computational Linguistics (NAACL-HLT '09),* 2009.

[43] G. Hirst and D. St-Onge, "Lexical Chains as Representations of Context for the Detection and Correction of Malapropisms," *WordNet: An Electronic Lexical Database,* pp. 305-332, MIT Press, 1998.

[44] T. Hughes and D. Ramage, "Lexical Semantic Relatedness with Random Graph Walks," *Proc. Joint Conf. Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL '07),* pp. 581-589, 2007.

[45] E. Gabrilovich and S. Markovitch, "Computing Semantic Relatedness Using Wikipedia-Based Explicit Semantic Analysis," *Proc. Int'l Joint Conf. Artificial Intelligence (IJCAI '07),* pp. 1606-1611, 2007.

[46] Y. Matsuo, J. Mori, M. Hamasaki, K. Ishida, T. Nishimura, H. Takeda, K. Hasida, and M. Ishizuka, "Polyphonet: An Advanced Social Network Extraction System," *Proc. 15th Int'l World Wide Web,* 2006.

[47] A. Bagga and B. Baldwin, "Entity-Based Cross Document Coreferencing Using the Vector Space Model," *Proc. 36th Ann. Meeting of the Assoc. for Computational Linguistics and 17th Int'l Conf. Computational Linguistics (COLING-ACL),* pp. 79-85, 1998.

Danushka Bollegala received the BS, MS, and PhD degrees from the University of Tokyo, Japan, in 2005, 2007, and 2009, respectively. He is currently a research fellow of the Japanese society for the promotion of science (JSPS). His research interests are natural language processing, web mining, and artificial intelligence.

Mitsuru Ishizuka (M'79) received the BS and PhD degrees in electronic engineering from the University of Tokyo in 1971 and 1976. He is a professor at the Graduate School of Information Science and Technology, the University of Tokyo, Japan. His research interests include artificial intelligence, web intelligence, and life-like agents. He is a member of the IEEE.

Yutaka Matsuo received the BS, MS, and PhD degrees from the University of Tokyo, in 1997, 1999, and 2002, respectively. He is an associate professor at the Institute of Engineering Innovation, the University of Tokyo, Japan. He joined the National Institute of Advanced Industrial Science and Technology (AIST) from 2002 to 2007. He is interested in social network mining, text processing, and semantic web in the context of artificial intelligence research.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.