

## Semantic Clustering in Dutch

Automatically inducing semantic classes from large-scale corpora

TIM VAN DE CRUYS

*CLCG, University of Groningen*

### Abstract

Handcrafting semantic classes is a difficult and time-consuming job, and depends on human interpretation. Unsupervised machine learning techniques might be much faster, and they do not rely on interpretation, because they stick to the data. The goal of this research is to present some clustering techniques that make it possible to automatically achieve Dutch word classes. More particularly, vector space measures are used to compute the semantic similarity of nouns according to the adjectives those nouns collocate with. Such semantic similarity measures provide a thorough basis to cluster nouns into semantic classes. Partitional clustering algorithms, that produce stand-alone clusters, as well as agglomerative clustering algorithms, that produce hierarchical trees, are investigated. For the evaluation of the clusters, evaluation frameworks will be used that compare the clusters to the hand-crafted Dutch EuroWordNet and the Interlingual Wordnet synsets. Additionally, the clustering of adjectives according to the collocating nouns has been investigated.

## 2.1 Introduction

Automatically acquiring semantics from text is a subject that has gathered a lot of attention for quite some time now. As Manning and Schütze (2000) point out, most work on acquiring semantic properties of words has focused on *semantic similarity*. ‘Automatically acquiring a relative measure of how similar a word is to known words (...) is much easier than determining what the actual meaning is.’ (Manning and Schütze 2000, 295)

Most work on semantic similarity relies on the Distributional Hypothesis (Harris 1985). This hypothesis states that words that occur in similar contexts tend to be similar. Take for example the invented word *sneup*, used in a number of contexts:

- verse ‘fresh’ sneup
- gezouten ‘salty’ sneup
- lekkere ‘tasty’ sneup
- zoete ‘sweet’ sneup
- taaie ‘tough’ sneup

A speaker of Dutch who is not familiar with the word *sneup* can easily infer from the context that it is some kind of food. In the same way, a computer might be able to extract similar words from similar contexts, and group them into clusters. There are, however, some problems with such an automatic approach. Ambiguity is the most important problem. Take the examples:

- (1) een oneven nummer  
a odd number  
‘an odd number’
- (2) een steengoed nummer  
a great number  
‘a great song’

The word *nummer* does not have the same meaning in these examples. In example 1, *nummer* is used in the sense of ‘designator of quantity’. In example 2, it is used in the sense of ‘musical performance’. Accordingly, we would like the word *nummer* to end up in two different clusters, the first cluster consisting of words like *getal* ‘number’, *cijfer* ‘digit’ and the second cluster containing words like *liedje* ‘song’, *song* ‘song’.

While it is relatively easy for a human language user to distinguish between the two senses, this is a difficult task for a computer. Moreover, the results get blurred because the attributes of both senses (in this example *oneven* and *steengoed*) are grouped together. In 2.2.3, an approach is touched upon that might be able to resolve this kind of ambiguity. But in this research, an active disambiguation of words has not been pursued.

## 2.2 Concepts and Methodology

### 2.2.1 Vector Space Measures

The actual semantic similarity of words is determined by means of vector space measures. The two words, for which the semantic similarity is to be calculated, are represented as vectors in a multi-dimensional space. With regard to quantitative data, there are two possible vector space representations: binary vector spaces and real-valued vector spaces. **Binary vectors** only have one bit of information on each dimension, to indicate presence or absence of a feature. For linguistic objects, the **real-valued vector space** is more appropriate, as this makes it possible to encode the frequency of the attribute.

In this research, the vector space consists of the adjectives (modifiers) of the nouns. Figure 3 gives an example of four nouns represented as vectors in *modifier space*.

	rood	lekker	snel	tweedehands
appel	2	1	0	0
wijn	2	2	0	0
auto	1	0	1	2
vrachtwagen	1	0	1	1

FIGURE 3 A noun-by-adjective matrix

The matrix shows that the modifier *rood* collocates with all four nouns, while *lekker* only collocates with *appel* and *wijn*. On the other hand, *snel* and *tweedehands* only collocate with *auto* and *vrachtwagen*.

This example shows how it might be possible to make a judgement about the similarity of nouns according to the collocating adjectives. However, in order to make this approach really useful, an appropriate similarity measure is needed. Such a measure is discussed below.

### 2.2.2 Similarity measure

Several similarity measures are available to calculate the similarity among various patterns. A few possibilities are *Dice coefficient*, *Jaccard coefficient* and *Overlap coefficient*. An overview of various similarity measures for lexical distributional similarity is given in Weeds, Weir and McCarthy (2004). van der Plas and Bouma (2005) provide an evaluation of distributional similarity measures applied to Dutch syntactic relations.

In these experiments, the *cosine measure* has been used. The cosine measure penalizes less in cases where the number of non-zero entries is very different. This seems appropriate in the context of distributional similarity, since the amount of data available for certain words might be quite different, and we do not want to qualify words as dissimilar because of this property.

For the general case of two n-dimensional vectors  $\vec{x}$  and  $\vec{y}$  in a real-valued space, the cosine measure can be calculated as follows:

$$\cos(\vec{x}, \vec{y}) = \frac{\vec{x} \cdot \vec{y}}{|\vec{x}| |\vec{y}|} = \frac{\sum_{i=1}^n x_i y_i}{\sqrt{\sum_{i=1}^n x_i^2} \sqrt{\sum_{i=1}^n y_i^2}}$$

This formula yields a number between 0 and 1, where 0 means no similarity at all and 1 means identical vectors. When applying the cosine similarity measure to the vectors in figure 3, we get:

- $\cos(\text{appel}, \text{wijn}) = \frac{6}{\sqrt{40}} \cong 0.94$
- $\cos(\text{auto}, \text{vrachtwagen}) = \frac{4}{\sqrt{18}} \cong 0.94$
- $\cos(\text{appel}, \text{vrachtwagen}) = \frac{2}{\sqrt{15}} \cong 0.51$

These simple examples show how semantically similar words are found: semantically similar words get high cosine values due to equal contexts of collocating adjectives, while semantically dissimilar words get a lower cosine value because of differing collocating adjectives.

### 2.2.3 Clustering

There are various clustering methods. An extensive overview of clustering is given in Jain, Murty and Flynn (1999). In general, a distinction can be made between:

- partitional clustering algorithms: algorithms that produce ‘stand-alone’ clusters which are not embedded in a structure;
- agglomerative (hierarchical) clustering algorithms: algorithms that assign a complete branching structure to the various clusters, up to the root node.

Both algorithms are worth exploring in the framework of semantic clustering. Partitional clustering is interesting to check whether similar words get grouped together. Hierarchical clustering is relevant for testing whether this kind of clustering is able to produce a sensible wordnet, that is comparable to hand-crafted wordnets. Both approaches have been explored in this paper.

#### Partitional clustering

As a partitional algorithm, K-means (MacQueen 1967) has been used. The procedure of K-means is as follows:

1. Choose  $k$  cluster centers, which are usually  $k$  randomly-chosen patterns or  $k$  randomly defined points inside the vector space;
2. assign each pattern to the closest cluster center;
3. recompute the cluster centers using the current cluster memberships;

4. if a convergence criterion is not met, go to step 2. Otherwise, stop the algorithm. The convergence criterium might be: no (or minimal) reassignment of patterns to new cluster centers, or a minimal decrease in squared error (a measure to check whether there are still many differences in cluster assignment in each iterative step).

### **Hierarchical clustering**

There are three well-known algorithms for hierarchical clustering: single-link, complete-link and group-average agglomerative clustering. These algorithms only differ in the way they characterize the similarity between a pair of clusters. In the single-link method (Sneath and Sokal 1973), the distance between two clusters is the minimum of the distances between all pairs of patterns drawn from the two clusters (one pattern from the first cluster, the other from the second). In the complete-link algorithm (King 1967), the distance between two clusters is the maximum of all pairwise distances between patterns in the two clusters. In group-average agglomerative clustering (Han and Kamber 2001), the distance between two clusters is the average distance between patterns in the two clusters. In the three algorithms, two clusters are merged to form a larger cluster based on minimum distance criteria.

This research has opted for the group-average agglomerative clustering algorithm. Group-average agglomerative clustering stands midway between single-link, which quickly merges clusters together, and complete-link, which tends to be conservative in cluster merging. The procedure of this algorithm is elaborated below:

1. The algorithm starts by taking each individual pattern in the pattern set to form a cluster;
2. next, the two clusters which are most similar are grouped together. Most similar means: the two clusters with the smallest distance between the averages of the clusters;
3. step two is repeated until there is only one cluster left. When the algorithm terminates, all clusters are hierarchically connected to the root node.

### **Hard and soft clustering**

An extra distinction needs to be made between hard clustering and soft clustering algorithms. In hard clustering algorithms, each element is assigned to exactly one cluster. In soft clustering algorithms, an element may be assigned to several clusters. Usually, soft clustering algorithms yield a probability distribution for each pattern, in which some patterns are more likely to belong to certain clusters than to others. This might seem a tempting approach for natural language processing, because ambiguity requires some words to be assigned to several clusters. However, soft clustering, as it is generally understood, is not the most appropriate approach to cope with disambiguous words. Since all the attributes of ambiguous words are taken into account (attributes that belong to different senses of the word), the vector that is constructed cannot represent both senses of the word, but it will present some kind of average, in which the most dominant sense will have

the upper hand. Therefore, only hard clustering approaches have been pursued (equally running the risk of wrong cluster assignment with ambiguous words, and missing out on less frequent senses of an ambiguous word).

There is, however, another soft clustering approach. In this approach, an ambiguous word is first assigned to a (dominant) sense of the word (found with all the attributes of the word). Once assigned to a certain cluster, the attributes that belong to this cluster are removed from the word vector, so that other, less common senses of the word might be revealed. Such algorithms are called *disjunctive clustering models*. It might be interesting to develop such an algorithm for Dutch. The algorithm discussed by Pantel and Lin (2002) would be a good algorithm to start from. This algorithm indeed tries to find less common word senses by stripping the values of more common senses off the feature vector. However, finding less common senses of a word most likely requires much more data.

## 2.2.4 Experimental Design

All adjective-noun collocations have been extracted from the Twente Nieuws Corpus (TwNC). The corpus was tagged with Mbt (Daelemans, Zavrel et al. 1996), a memory-based tagger, and lemmatized with Mblem (van den Bosch and Daelemans 1999), a memory-based lemmatizer. The frequency of the adjectives has been logarithmically smoothed ( $f(x) = 1 + \ln(x)$  for each  $x > 0$ ), in order to normalize the occurrence of many instances of one single adjective.

Various parameters have been used for clustering. A combination of the 5,000 most frequent nouns (adjectives) together with the 20,000 most frequent adjectives (nouns) was experimentally found to be functioning best (clustering more nouns yields much worse cluster quality, clustering the same nouns with more adjectives does not yield any significant improvement). This boils down to a frequency cut-off of 200 individual noun-adjective collocations per noun, while each adjective occurs at least 5 times.

## 2.3 Results

### 2.3.1 Partitional clustering

#### Noun Clustering

Below are some of the clusters that have been found by the K-means algorithm, clustering the 5,000 nouns into 700 clusters<sup>2</sup>:

- april januari november februari oktober maart mei juni augustus december september juli
- sprinter schaatser coureur speelster middenvelder vedette wielrenner aan-

---

<sup>2</sup>Note that the corpus has been lemmatized, but due to errors of the lemmatizer, tokens might sometimes end up in the clusters.

voerder keeper landgenoot speler atlete aanvaller renner verdediger atleet  
kopman bokser voetballer zwemmer spits tennisser doelman

- dollar ton euro meter kilometer kilo pond gulden centimeter
- hoofdredacteur chef commandant secretaris-generaal bestuursvoorzitter  
bevelhebber hoofdofficier directeur
- maand zomer winter winters week eeuw herfst
- stijger stijgers daler kanshebber winnaars boosdoener verliezer troef ver-  
liezers
- bisschop priesters Kerk predikant priester kerk dominee gelovigen kerken  
bisschoppen

### Adjective Clustering

The clustering has also been done the other way around: the 5,000 most frequent adjectives have been clustered according to the 20,000 most frequent collocating nouns. Such kind of clustering produces results of the kind below:

- geel paars zwart groen blauw grijs oranje bruin roze wit rood
- Duits Amerikaans Zweeds Russisch buitenlands Brits Belgisch Nederlands  
Frans Engels Japans Zwitsers Italiaans Spaans Chinees
- rk roomskatholieke russisch-orthodoxe servisch-orthodoxe oud-katholiek  
anglicaans r.k. koptisch Koptisch grieks-orthodoxe
- zonovergoten herfstig winters druilerig zomers regenachtig zonnig
- deplorabel abominabel erbarmelijk penibel mensonterend mensonwaardig  
benard miserabel

### 2.3.2 Agglomerative Clustering

Agglomerative clustering also yields some remarkable results. What is remarkable is that the upper nodes present broad semantic categories, such as persons, objects, abstract entities, . . . Figure 4 shows part of an example of an agglomerative tree, grouping together nouns that designate a time entity. Note that the edges in the tree should not be interpreted as actual ‘is a’-relations (as is the case in Wordnet), but rather as relations of semantic relatedness.

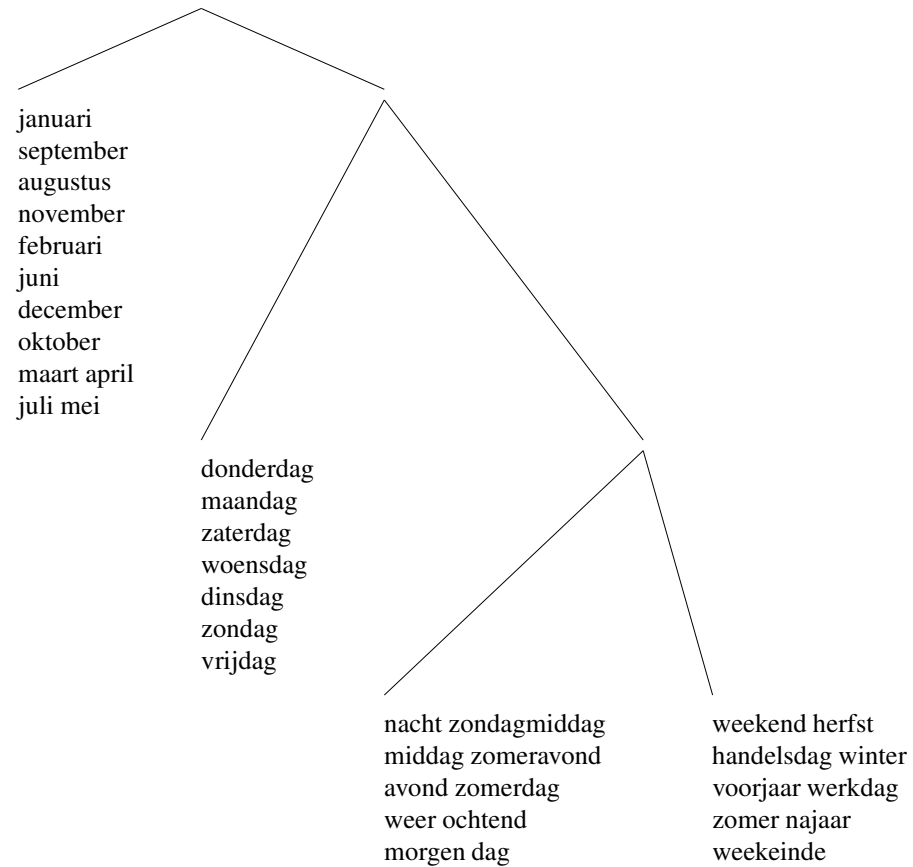


FIGURE 4 Example of agglomerative clustering: days, months, seasons

## 2.4 Evaluation

### 2.4.1 Automatic Evaluation with EuroWordNet

#### Methodology

For the evaluation of the clusters, precision and recall has been calculated according to the relations that exist in the Dutch version of EuroWordNet. The procedure of the evaluation is as follows:

- The wordnet relations that are used for the evaluation are:
  - synonyms
  - hyponyms
  - hypernyms



- co-hyponyms (hyponyms of the hypernyms)
- For each cluster, it is checked in EuroWordNet which word from the cluster has most relations in EuroWordNet with the other words from the cluster. This word is taken to be the most central word of the cluster.
- For this word, the synonyms, hyponyms, hypernyms and co-hyponyms are drawn from EuroWordNet.
- To calculate precision, it is checked how many words from the cluster are actually appearing in the EuroWordNet-relations.
- To calculate recall, it is checked how many of the EuroWordNet-relations are not appearing in the found cluster.

A few remarks are to be made with regard to this evaluation framework. Recall will be a low number, because different kind of relations are considered in the evaluation framework: hyponyms, hypernyms, synonyms and co-hyponyms are considered all together. The real recall (as acknowledged by human judges) is probably a lot higher. To a human judge, a cluster that contains the 7 days of the week seems quite complete, but in this evaluation framework, it gets a recall of 9.21%. Therefore, recall is not such a good measure in this evaluation framework; precision will therefore be considered the most important value.

## Results

Figure 5 presents the results of the **partitional noun clustering** evaluation. The precision and recall values are plotted against the number of clusters used.

The figures show a precision that is lower with few, large clusters, rising towards an optimal number of clusters, and then declining again when the number of clusters gets too small. Recall is faintly showing the opposite tendency, being larger with fewer but large clusters, and declining when the number of clusters get larger. But as has been explained before, recall is not such a good measure in this case.

With the optimal number of clusters, the clustering algorithm is able to reach a precision of **42.50%** (and a recall of about **8%**). Taking into account that EuroWordNet itself is incomplete (a large part of the clustered words is not known by EuroWordNet), that the evaluation algorithm is only looking one level up and down in the wordnet hierarchy, and that there is an error margin due to mistakes of the lemmatizer, the results obtained by the clustering algorithm are quite good, also given that the random baseline (results for clusters that have been randomly compiled using a hash table) is about **5.5%** for precision and about **3%** for recall.

It is interesting to have a look at the share of each relationship in the precision measure. This gives an indication of the relationships that are found by the clustering algorithm, and to what extent they are found. Figure 6 presents the share of each relationship graphically.

The majority of words found by the clustering algorithm are clearly co-hyponyms. This result was to be expected, as the horizontal relationship (which is mainly the co-hyponym




FIGURE 5 Evaluation of 5,000 nouns clustering with EuroWordNet

relationship) is the relationship in which most similarity is to be found. Synonyms (which is the second horizontal relationship) are also found by the algorithm, but the share of synonyms is of the same order as the hyponym and hypernym relationships. Most likely, this is because synonyms are less numerous than co-hyponyms. What is remarkable, is that co-hyponyms and synonyms seem to be following the same pattern: starting at a lower precision with large (fewer) clusters, rising to an optimum with middle-size clusters and declining again when the clusters get too small (too many clusters). This is not the case with the other relationships: the hypernym precision stays roughly at the same level (and is even rising a bit), while the hyponym precision is declining with the number of clusters. These results seem to indicate that hyponyms tend to get clustered more easily than hypernyms, when the margin is large enough (fewer but larger clusters). Of course, these tendencies are not significant enough to draw conclusions.

#### 2.4.2 Evaluation with Wu & Palmer's measure

The algorithm discussed in 2.4.1 tries to evaluate the cluster quality by comparing the clusters to a fixed set of words extracted from EuroWordNet. This makes it possible to calculate precision and recall values. Another kind of evaluation relies on measures of **semantic similarity** according to hierarchical wordnets.

A number of such similarity measures have been developed. Among these measures, the most important are Wu & Palmer's (Wu and Palmer 1994), Resnik's (Resnik 1995) and Lin's (Lin 1998).




FIGURE 6 Share of each relationship in total precision

In this evaluation, Wu & Palmer's measure will be adopted. Wu and Palmer (1994) have proposed a measure that calculates the similarity between two words according to their position in a hierarchical wordnet. The similarity is calculated according to the formula given below, in which  $N_1$  and  $N_2$  are the number of *is-a* links from  $A$  and  $B$  to their most specific common superclass  $C$ ;  $N_3$  is the number of *is-a* links from  $C$  to the root of the taxonomy.

$$sim_{Wu\&Palmer}(A, B) = \frac{2N_3}{N_1 + N_2 + 2N_3}$$

For example, the most common superclass of *hond* en *zalm* is *dier* (as can be seen on the extract from Dutch EuroWordNet in figure 7). Consequently,  $N_1 = 2$ ,  $N_2 = 2$ ,  $N_3 = 4$  and  $sim_{Wu\&Palmer}(hond, zalm) = 0.67$ .

The results have not been calculated by using the Dutch EuroWordNet directly, as was the case with the former evaluation framework. Instead, the words have been converted to Interlingual WordNet offsets.<sup>3</sup> This way, it was possible to make use of a perl module that implements the computation of Wu & Palmer's measure in the English WordNet (Pedersen, Patwardhan and Michelizzi 2004).

<sup>3</sup>Interlingual WordNet offsets are identification codes connected with a particular WordNet synset, that have been designed to make translations among various languages possible. They are available in EuroWordNet, and are basically the same as the ones used in the English WordNet (there are differences among different versions of WordNet, but conversion procedures exist).




FIGURE 7 Extract from the Dutch EuroWordNet hierarchy

The average cluster similarity has been calculated as follows:

- For each cluster, the most central word has been taken (which is the word that was the closest to the cluster's centroid);
- the similarity between this most central word and every other word in the cluster has been calculated;
- If a word is ambiguous (i.e. has more than one synset), similarity has been calculated for all synsets, and the highest value has been retained;
- the average of these similarities has been taken, and every cluster average has been added up;
- all cluster averages have been divided by the total number of clusters, to get the total average;
- words not known by WordNet have been ignored.

Figure 8 shows the results of the evaluation with Wu&Palmer's measure. There's an average of 60% similarity within the clusters, while randomly compiled clusters have an average of 29% similarity. These results seem to confirm the results found by the former evaluation framework.

## 2.5 Conclusion & Further Work

In this research, clustering techniques have been explored that make it possible to automatically acquire semantic classes in Dutch. More particularly, vector space measures have been used to calculate semantic similarity. These semantic similarity measures are then used to cluster nouns into classes. Partitional K-means clustering has been used as a




FIGURE 8 Evaluation the clustering of nouns (Wu&Palmer similarity measure)

partitional clustering algorithm and group-average agglomerative clustering has been used as a hierarchical clustering algorithm.

The results and the evaluation of the clusters have shown that using the syntactic context of words (more particularly, using modifiers to cluster nouns) is indeed a useful approach for extracting semantic classes. The evaluation of the clusters shows significant similarities with Wordnet-relations, in my own evaluation framework evaluating direct relationships, as well as with Wu & Palmer's similarity measure. The clustering of modifiers (adjectives) according to their heads (nouns) also seems to yield quite good results, although this kind of clustering has not yet been evaluated.

There are, however, some issues that make the automatic clustering of nouns less straightforward. Ambiguity is one of the problems that is difficult to tackle for a computer. Ambiguity blurs the results, because both senses of the word (with their accompanying values) get grouped together into one sense. Disambiguating these ambiguous words into different clusters is one of the main goals to be solved, in order to reach a semantic clustering that is able to compete with hand-crafted semantic classes.

Another issue that requires more research is the automatic extraction of complete wordnets, instead of stand-alone clusters. Instead of focusing on the horizontal semantic relationships (creating clusters of similar words) it would be interesting to explore algorithms that automatically acquire the vertical relationships. This way, it might be possible to automatically construct a complete wordnet, similar to hand-crafted wordnets available. The

agglomerative clustering algorithm is a first step towards this direction, though still far from perfect.

Also, the field of verb clustering remains to be explored. Subject-verb and verb-object relations are quite different from adjective-noun relations. How these relations might be used in order to cluster verbs, is subject to further research. It also remains to be investigated how these relations might help in improving the clustering of nouns.

A final interesting subject for future research is the application of dimensionality reduction techniques (LSA, PLSA) to counter data sparseness and noise. The application of these techniques will be explored in order to bring about a better clustering.

### Acknowledgements

This research has been carried out during a three month internship at the *Centrum voor Nederlandse Taal en Spraak* at the University of Antwerp. The work has been supervised by Walter Daelemans and Bart Decadt, and has substantially benefited from their remarks and contributions.

### References

- Daelemans, W., Zavrel, J. et al.(1996), Mbt: A memory-based part of speech tagger-generator, in E. Ejerhed and I. Dagan (eds), *Proceedings of the Fourth Workshop on Very Large Corpora*, Copenhagen, Denmark, pp. 14–27.
- EuroWordNet, Building a multilingual database with wordnets for several European languages*(n.d.), website. <http://www.illc.uva.nl/EuroWordNet/>.
- Han, J. and Kamber, M.(2001), *Data Mining – Concepts and Techniques*, Morgan Kaufmann, Boston.
- Harris, Z.(1985), Distributional structure, in J. J. Katz (ed.), *The Philosophy of Linguistics*, Oxford University Press, pp. 26–47.
- Jain, A. K., Murty, M. N. and Flynn, P.(1999), Data clustering: a review, *ACM Computing Surveys* **31**(3), 264–323.
- King, B.(1967), Step-wise clustering procedures, *Journal of the American Statistical Association* **69**, 86–101.
- Lin, D.(1998), An information-theoretic definition of similarity, *Proceedings of the 15th International Conference on Machine Learning*, Morgan Kaufmann, San Francisco, pp. 296–304.
- MacQueen, J.(1967), Some methods for classification and analysis of multivariate observations, in L. M. Le Cam and J. Neyman (eds), *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, University of California Press, Berkeley, California, pp. 281–297.
- Manning, C. and Schütze, H.(2000), *Foundations of Statistical Natural Language Processing*, MIT Press, Cambridge, Massachusetts.
- Pantel, P. and Lin, D.(2002), Discovering word senses from text, *Proceedings of ACM*

- Conference on Knowledge Discovery and Data Mining*, Edmonton, Canada, pp. 613–619.
- Pedersen, T., Patwardhan, S. and Michelizzi, J.(2004), Wordnet::similarity - measuring the relatedness of concepts. <http://citeseer.ist.psu.edu/665035.html>.
- Rennie, J.(2000), Wordnet::querydata: a Perl module for accessing the WordNet database, website. <http://people.csail.mit.edu/~jrennie/WordNet>.
- Resnik, P.(1995), Using information content to evaluate semantic similarity in a taxonomy, *IJCAI*, pp. 448–453.
- Sneath, P. H. A. and Sokal, R. R.(1973), *Numerical taxonomy: the principles and practice of numerical classification*, Freeman, San Francisco.
- Twente News Corpus*(n.d.), website. <http://www.vf.utwente.nl/~druid/TwNC/TwNC-main.html>.
- van den Bosch, A. and Daelemans, W.(1999), Memory-based morphological analysis, *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, University of Maryland, USA, pp. 285–292.
- van der Plas, L. and Bouma, G.(2005), Syntactic contexts for finding semantically similar words, in T. van der Wouden et al. (eds), *Computational Linguistics in the Netherlands 2004. Selected Papers from the Fifteenth CLIN Meeting*, LOT, Utrecht, pp. 173–184.
- Weeds, J., Weir, D. and McCarthy, D.(2004), Characterising measures of lexical distributional similarity, *Proceedings of the 20th International Conference of Computational Linguistics*.
- WordNet, a lexical database for the English language*(n.d.), website. <http://wordnet.princeton.edu/>.
- Wu, Z. and Palmer, M.(1994), Verb semantics and lexical selection, *32nd. Annual Meeting of the Association for Computational Linguistics*, New Mexico State University, Las Cruces, New Mexico, pp. 133–138.

