

Class-Based Statistical Models for Lexical Knowledge Acquisition


Stephen Clark

CSRP 539

September, 2001

ISSN 1350-3162

UNIVERSITY OF



Cognitive Science
Research Papers

1000 - C 1000

2	Uses of lexical sense preferences	
2	Using probabilities to represent preferences	
2	Estimating probabilities of senses	
4	Overview of chapters	4
2	Lexical knowledge acquisition	
2	Acquisition of selectional preferences	
2	2 Distributional similarity	
2	Cooccurrence extraction	
2	4 Smoothing for probability estimation	2
2.2	Applications	24
2.2	Structural disambiguation PP attachment	24
2.2.2	Parse selection	2

C

	Problem specification	
	The semantic hierarchy	
2	Class based probability estimation	

2 Tf m 2 a t2 io Estimating the relevant probabilities in Class based probab e 4 i b e 2 a 4 se 4 e 4 d 2 v 2 p ar 4

2	Example thesaurus and cut	2
2.2	Example Bayesian network	
2	Example dependencies and base NPs	
	Part of the WordNet hierarchy	4
2	Proof of proposition	
	Proof of proposition	
4	An algorithm for determining $\text{top}(\bar{c}, r)$	4
	An example generalisation determining $\text{top}(\langle \text{soup} \rangle, s, \bar{c}, \text{obj})$	42
4	Calculation of $A(\overline{\langle \text{beverage} \rangle}, r, \bar{c}, \text{obj})$	
4.2	Estimates for calculating $A(C, r)$ for a set of concepts C	2
4	Calculation of $A(\overline{\langle \text{twelve} \rangle}, \bar{c}, o, \text{obj})$	4
4.4	Calculation of $A(\overline{\langle \text{food} \rangle}, \bar{c}, \text{obj})$	
	Example dependency structure	
2	Sequence of decisions generating a dependency structure	2
	Part generation of a dependency structure	
4	The complete dependency structure	
	The grammatical relations used in the implementation	
	Example parser output	
	Example dependency structure	
	Dependency probabilities by relation	4 4 4 2 Tf 2 2 4 Tm ica

2	Contingency table for the bigram	2
	Contingency table for the children of $\langle \text{canine} \rangle$ in the subject position of r_{un}	
2	Contingency table for the children of $\langle \text{liquid} \rangle$ in the object position of r_{un}	4
	Example levels of generalisation for different values of α	4
4	Example levels of generalisation for different sample sizes	44
	The extent of generalisation for different values of α and sample sizes	4
	The extent of generalisation for G^2 and X^2	4
4	Examples of re estimated frequencies	
4 2	How the distribution of the count changes	
4	Disambiguation results across a range of α values	
	Frequency of each type of relation in the test suite	4
2	Lower and upper bounds for the F score	
	Results for the LR parser	
4	Results for the dependency model	
	PP results for the complete test set of , test cases	2
2	PP results when $\langle \text{root} \rangle$ is selected for neither attachment point	2
	PP results when $\langle \text{root} \rangle$ is selected for at most one of the attachment points	
4	PP results for G^2 and X^2 on complete test set of , cases	
	PP results for G^2 and X^2 when $\langle \text{root} \rangle$ is selected for at most one attachment point	4
	Results for the pseudo disambiguation task	
	Results for the pseudo disambiguation task with th training data	
	Disambiguation results for G^2 and X^2	

2 Coreference

This basic approach can be applied to other problems such as anaphora resolution and word sense disambiguation. Consider the problem of determining the referent of *it* in the following sentence taken from Wilks:

I bought the wine, sat on a rock and drank it.

To determine the correct referent, we can use the fact that the correct sense of *it* is more likely to be an object of *drink* than the correct sense of *rock*.

Resnik (2001) argues that the constraints a predicate places on its arguments are not Boolean constraints as in the classical account of selectional restrictions (Katz and Fodor, 1963) but that the constraints are satisfied to a certain degree. Resnik cites McCawley (1998) and Fodor (1998) as earlier critics of Katz and Fodor's theory. We follow Resnik in modelling the constraints as graded preferences and in line with other recent work in this area (Ribas, 2005; Li and Abe, 2006; McCarthy, 2006; Wagner, 2006) probabilities are used to encode the preferences. An important question is whether the preference measure should define a probability distribution over the possible arguments of a predicate.

Resnik's measure of selectional preference which he calls selectional association is defined in terms of probabilities but the measure does not define a probability distribution over the possible arguments of a predicate; the values for selectional association need not lie between zero and one and do not sum to one over the possible arguments. This is also true of a number of related measures in the literature such as the chi squared statistic (Kilgarriff, 1991), likelihood ratio statistics (Dunning, 1994) and mutual information (Church and Hanks, 1990). Aside from the question of whether these measures are appropriate for use in corpus based linguistics (Dunning, 1994) they all suffer from a limitation.

The limitation arises when determining the semantic plausibility of a complex linguistic event such as a parse tree. In order to do parse selection one can measure the overall extent

This chapter is divided into two sections; one section describes work from those areas of lexical acquisition that are of particular relevance to this thesis and the other section describes previous approaches to structural disambiguation and parse selection. These areas of application are considered because the problems of structural disambiguation and parse selection are dealt with in Chapters 3 and 4.

The knowledge acquisition section focuses on selectional preferences, describing in detail those approaches that have used WordNet and showing how they relate to the class based estimation method described in Chapter 3. We also describe some approaches to automatic clustering which is an important alternative to using a man made hierarchy for generalisation and also collocation extraction which has used statistics that are used in Chapters 3 and 4. Finally a number of smoothing techniques for probability estimation are described; this work is relevant because the class based estimation method described in Chapter 3 can be thought of as performing a kind of smoothing.

The applications section focuses on those approaches to structural disambiguation and parse selection that have used knowledge similar to lexical sense preferences; this includes much of the recent work on resolving PP attachment ambiguities and statistical parsing where there has been a move towards probability models based on lexical dependencies.

The role of the lexicon has taken on increasing importance in recent years both from a theoretical and a computational perspective. One of the main reasons for this is the development of the computational lexicon. The development of the computational lexicon has been a major factor in the development of the computational lexicon. The development of the computational lexicon has been a major factor in the development of the computational lexicon.

arguments but rather has a *pr* *rr* kind of argument. However, Wilks distanced himself from a probabilistic treatment of preferences: it is still the case that an individual preference is either satisfied or it is not, as with selectional restrictions. The difference is that an interpretation of a sentence can be preferred even if individual preferences are violated, as long as there is no alternative interpretation with less violations.

Resnik [\[10\]](#) took the notion of preference one step further by suggesting that preference should be measured on a continuous scale. Resnik uses the following list of examples, which originally appeared in Drange [\[11\]](#), to demonstrate that the preferences of $\bar{1}$ $\bar{0}$

sn² s o o s *quæstion* pr r n²

The parts of Resnik's work a b a b that are most relevant for this thesis are his solutions to the following questions

- How can a probability distribution over the WordNet hierarchy be defined?
- 2 How can we measure the extent to which an argument satisfies the preferences of a predicate

Each question will be dealt with in turn

Resnik defines his probability model in terms of classes where *quæstion* has the interpretation given above. Let $C = \{c_1, c_2, \dots, c_n\}$ be the set of classes in WordNet where n is the number of concepts so that each concept has a corresponding class. Resnik places the following constraints on any probability distribution over C

$$\begin{aligned} \text{if } c_j \text{ is a kind of } c_i \text{ then } p(c_i) &\geq p(c_j) & 2 \\ \sum_{c_i \in C} p(c_i) &= 1 & 2.4 \end{aligned}$$

Equation 2 agrees with the intuition that the probability of a class increases with the level of abstraction. Although note that the probability corresponding to a node in the hierarchy is not defined in terms of the sum of the probabilities of the children. Equation 2.4 is required by Resnik because he defines a random variable ranging over all the classes and defines information theoretic functions of that random variable such as entropy.


Resnik's aim is to model the fact that some verbs select more strongly for their arguments than others. For example *selects* more strongly for its direct object than *is*. Resnik's approach is based on the fact that for strongly selecting verbs the probability of a class conditional on the verb $p(c_i | v)$ is likely to differ largely from the unconditional probability $p(c_i)$. From an information theoretic perspective a strongly selecting verb provides more information about the

C p r _ r ~~g~~us or

A difficulty with using selectional association in an application is that the arguments are likely to be nouns rather than classes and so an appropriate class has to be chosen for the noun. This problem has two dimensions since a noun can have more than one sense but can also be repre

BEVERAGE FOOD LIQUID FLUID ... ENTITY Each of these classes would receive a count of $1/2$ for each instance of x_i in the data. Note that this method of class estimation is unusual among the work in this area and is motivated by the desire to define a probability distribution over the set of all classes. The other work described here does not

2 C p r _ r ϕ_{us} or



relative to the entire data size and the number of words it generalizes them into a class. When the differences are especially noticeable relative to the entire data size and the number of the words, on the other hand, it stops generalization at that level.

As we shall see, a similar approach to generalization is taken in this thesis, but not using MDL. One of the problems with this generalization approach is that it is based on frequencies, which

considering The rst modi cation is based on the following observation that removing parts of the hierarchy based on the nouns that occur in the data can result in large parts being excised For example if *n* appeared in the data a large proportion of the complete hierarchy would be removed namely that part of the hierarchy dominated by `<entity>` McCarthy s alternative solution is to create new leaf nodes for each internal node in the hierarchy; for example the synset for the concept `<entity>` would be represented at a new leaf node having the internal `<entity>` node as a parent This modi cation results in all the nouns in the hierarchy being represented at leaf nodes Counts for nouns are distributed initially at leaf nodes and then passed up to internal nodes representing the classes

McCarthy s response to the DAG problem is to leave the hierarchy as a DAG and argue that since only around % of the nodes in WordNet have more than one parent the resulting tree cut models are unlikely to differ much from the tree case McCarthy also notes that the majority of cases of multiple inheritance occur low down in the hierarchy



each HMM remains the same but the values of the probabilities vary

To give an example consider how the noun *no* is generated for the object position of *question*. In fact since *no* has more than one sense in WordNet there are numerous paths through WordNet that generate the noun but let us assume that the noun is generated via the food sense. The hypernyms of the food sense of *no* are as follows $\langle \text{bread} \rangle$ $\langle \text{baked_good} \rangle$ $\langle \text{foodstuff} \rangle$ $\langle \text{food} \rangle$ $\langle \text{substance} \rangle$ $\langle \text{object} \rangle$ $\langle \text{entity} \rangle$. First a child of the root of the hierarchy is chosen in this case the $\langle \text{entity} \rangle$ node according to the transition probabilities associated with the root. Then the concept $\langle \text{object} \rangle$ is chosen according to the transition probabilities associated with $\langle \text{entity} \rangle$.

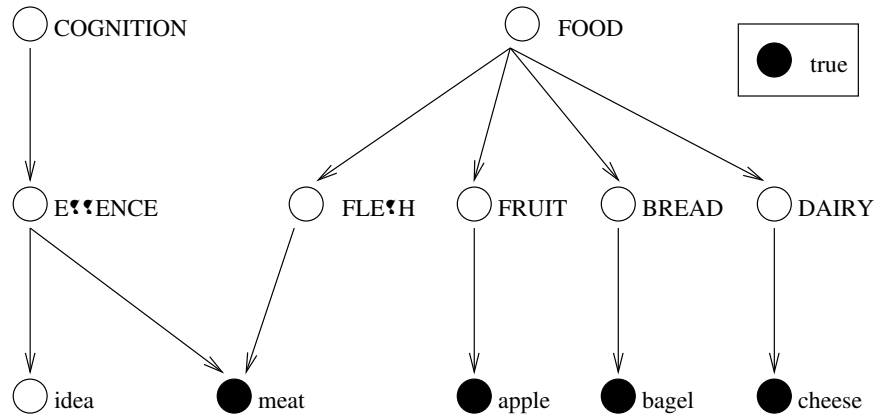


Figure 2.2 Example Bayesian network

variable which can be in one of two states r_u or s . A synset node has the value r_u if the concept represented by the synset is selected for by the verb and a word node has the value r_u if the word can appear as an argument of the verb.

Each variable A with parents B, \dots, B_n has associated with it a conditional probability table (CPT) which stores the probabilities $p(A|B, \dots, B_n)$. Ciaramita and Johnson call these probabilities the *probors* and they are defined according to the following principles. First, it is understood that a verb selects for a concept *probors* and a word can appear as an argument of the verb.

Cluster-based

Cluster

Pereira, Tishby, and Lee acquire clusters of nouns for the direct object position of verbs. The clustering is soft in that each word belongs to a cluster according to a cluster membership probability, and it is also hierarchical in that the clustering algorithm works in a top-down iterative fashion, splitting existing clusters at each iteration. The decision to keep two nouns in the same cluster is based on the difference between their conditional verb distributions $p_n(\cdot)$, which is measured using the KL divergence.

In contrast, Brown, Della Pietra, de Rouza, Lai, and Mercer [2] adopt a bottom-up iterative approach in which initially the clusters are the individual words themselves, and the decision to merge two classes is based on the minimal loss of mutual information. The clustering is hard in that a noun either belongs to a cluster or it does not, and there is no notion of degrees of membership. The clustering model was used to try and improve a language model, although no improvements in perplexity were gained by using a cluster b



no

equation

2 C p r _ r \mathcal{X} or

u u n or \mathcal{X}

The mutual information between two words x and y in some cooccurrence relation is defined as follows

$$I(x, y) = \log_2 \frac{p(x, y)}{p(x)p(y)} \quad 2$$

The mutual information described here is often referred to as *pointwise mutual information* to distinguish it from the notion used in information theory. Pointwise mutual information is derived from the information theoretic notion but the information theoretic version is defined as an average over random variables. Also the pointwise version has less of a theoretical basis; Jelinek warns that interpreting $I(x, y)$ as the mutual information between x and y gives only an intuitive interpretation. [p 4](#)

Pointwise mutual information compares the joint probability of observing x and y together p

(w_1, w_2)	$(\neg w_1, w_2)$
$(w_1, \neg w_2)$	$(\neg w_1, \neg w_2)$

Table 2 Contingency table for the bigram $w_1 w_2$

(w_1, w_2) is the number of times w_2 follows w_1 in the data and $(\neg w_1, w_2)$ is the number of times w_2 follows a word other than w_1 in the data. The other frequencies in the table are defined analogously. The null hypothesis corresponding to the table is that w_1 and w_2 appear independently of each other and a statistic such as chi squared can be used to determine how likely the null hypothesis is to be true. If the chi squared statistic has a high value then this gives strong evidence that the null hypothesis is false and that w_1 and w_2 are highly associated. Thus bigrams with high chi squared scores should correspond to highly associated pairs of words or collocations.

The chi squared statistic that is usually encountered in text books is the *person* chi squared statistic. However the problem with this statistic as Dunning demonstrates is that it can over estimate the significance of rare events. This means that the bigrams producing the highest scores are often based on very low counts which makes the test unreliable. Most of the top ranked bigrams in Dunning's experiments occurred only once in the data and among the highest ranked bigrams were cases like *pr* *r* *n* *his* *n* and *s* *nn* *r* *s* which are hardly highly associated pairs of words. As a remedy to this problem Dunning considers the log likelihood ratio statistic denoted G^2 which does not over estimate the significance of rare events in the same way. The top ranking bigrams produced according to this statistic were much more intuitive.

Dunning's analysis of his results is based on the following claim that the sampling distribution of G^2 approaches chi squared quicker than the sampling distribution of X^2 . However this part of Dunning's analysis is debatable since Agresti makes exactly the opposite claim.

The sampling distributions of X^2 and G^2 get closer to chi squared as the sample size n increases ... The convergence is quicker for X^2 than G^2 p 4

Given Agresti's comments a more likely explanation lies in the conservative nature of G^2 which means that X^2 is more likely to return a significant result for a table based on small counts. This would explain Dunning's results in which pairs of words occurring infrequently in the corpus obtain high scores according to X^2 but not G^2 . These issues will be discussed further in Chapter where a chi squared test is used as part of a procedure for selecting a suitable level of abstraction in WordNet.

Pedersen suggests using Fisher's exact test Agresti for bigram discovery rather than a chi squared statistic. The advantage of Fisher's exact test is that it can be applied to any contingency table regardless of the size of the counts and the result will be reliable. However the test is computationally expensive since it involves computing every contingency table that could have led to the marginal totals observed in the sampled table. The marginal totals are not shown in Table 2 but are simply the totals obtained by summing the scores in each row and column. In addition the results obtained by Pedersen for the exact test did not differ greatly from those obtained for the log likelihood statistic and so it is not clear that the benefits of using the test outweigh the additional computational burden.

Corpus Smoothing

Many of the smoothing techniques used in corpus based NLP were developed for language modelling and so to demonstrate some of the most widely used techniques we consider the problem of estimating an n gram model. More specifically the problem is to estimate the probability of a word conditional on the previous $n - 1$ words $p(w_n | w_{n-1}, \dots, w_1)$. A maximum likelihood

As an example consider using 2.22 to estimate $p(\langle \text{fox} \rangle | r_{un}, \text{subj})$ and $p(\langle \text{carpet} \rangle | r_{un}, \text{subj})$ assuming that neither $\langle \text{fox} \rangle$ nor $\langle \text{carpet} \rangle$ appear with r_{un} in the data. Unlike additive smoothing the two unseen senses are unlikely to receive the same estimate since the estimates based on less context are unlikely to be the same for the two senses. However $\langle \text{fox} \rangle$ will not necessarily receive a higher estimate than $\langle \text{carpet} \rangle$; the problem is that the estimates based on less context ignore the verb. In contrast the estimation method presented in Chapter 3 is able to make use of the verb by determining whether semantically similar senses to $\langle \text{fox} \rangle$ and $\langle \text{carpet} \rangle$ appear as subjects of r_{un} .

Good Turing

Another widely used technique is the Good Turing method (Good, 1952) which states that an n -gram that has occurred r times in the data should have an adjusted frequency r^* where

$$r^* = (r + 1) \frac{E(r + 1)}{E(r)} \quad (r \geq 1) \tag{2.2}$$

$E(r)$ is the expected number of n -grams that occur r times in the data. Relative frequencies based on the r^* values can be used to estimate the probabilities. Note that 2.2 only applies to values of r greater than zero; a further result of Good (1952) is that the total probability mass assigned to unseen objects is $E(0) / E(1)$ where $E(1)$ is the total number of n -grams in the data.

In practice the actual number of n -grams that occur r times in the data, n_r , can be used to approximate the expected values if the actual values are suitably smoothed themselves. To see

This section describes previous work on structural disambiguation which is a problem considered later in the thesis. The section describes work on PP attachment and then work that has considered the more general problem of parse selection. Not all previous approaches are considered since the literature in both cases is very large and we describe only those approaches that are most relevant to the work in this thesis.

The type of structural ambiguity that has been most covered in the literature is PP attachment ambiguity. This is a pervasive form of ambiguity and a potentially damaging one in that increasing the number of PPs in a sentence can lead to a combinatorial explosion in the number of possible analyses (Church and Patil, 2002). A number of early studies in the psycholinguistics domain suggested possible strategies for resolving attachment ambiguities. Two of the most cited studies are those of Kimball (1973) who suggested that a constituent tends to attach to another constituent immediately to its right (right association) and Frazier (1986) who suggested that there is a preference for attachments that lead to the parse tree with the fewest nodes (minimal attachment). However, later work (Whittemore, Ferrara, and Brunner, 1998; Taraban and McClelland, 1999) demonstrated that lexical information is a better predictor of attachments and most of the recent corpus based approaches to structural disambiguation including PP attachment have been based on lexical associations.

The PP problem that is usually addressed only considers sequences of the following form: $n_1 \text{ PP } n_2$ or $n_1 \text{ PP } n_2 \text{ PP } n_3$. Moreover, only the heads of the noun phrases are usually considered. The problem can then be characterised as taking a four tuple (n_1, n_2, p, n_3) and deciding whether the PP attaches to n_1 or n_2 as in the much used example: *She saw the man with a telescope*. Note that this is an easier problem than the most general form of PP attachment since only two possible attachment sites are being considered. In the general case there may be more than two sites. Consider this example from Hindle and Rooth:

- 2.24 NBC was so afraid of hostile advocacy groups and unnerving advertisers that it shot its dramatization of the landmark court case that legalised abo

-

$$p(A|, n, pr, n_2) = \begin{cases} \text{if } A \text{ is noun attach} \\ \text{if } A \text{ is verb attach} \end{cases}$$

An interesting result of the paper is that the optimum value for ϕ_{us} was found to be zero at all stages. This means that even if a context occurs only once in the training data it is better to use an estimate based on that context rather than back off to another level. We present a related result in Chapter 4 regarding the use of low count events in the training data. We find that

simply compares probabilities corresponding to the possible attachment sites. An advantage of our approach is that these probabilities can be easily integrated into a model for parse selection.

2 C p r _ r ~~h~~us or

The problem of parse selection is to select the correct parse for a sentence from a number of alternatives. As Collins (1999) points out, this can be an astonishingly severe problem in broad domains such as the Wall Street Journal (WSJ). Collins cites a number of factors that are responsible for the severity of the problem: the need for a large grammar to obtain broad coverage; long sentences being typical in a broad domain; and many common sources of syntactic ambiguity such as PP attachment leading to exponential growth in the number of analyses relative to sentence length. There are many examples in the literature of ordinary looking sentences having hundreds, sometimes thousands, of different analyses according to some grammar. The parser of

C p r _ r ⁺ ~~o~~us or



the probability model

Briscoe and Carroll (1999) define a probability model based on the moves of an LR parser (see also Briscoe and Carroll (1999), Carroll and Briscoe (1999), Carroll, Minnen, and Briscoe (1999)). The grammar underlying the parser is a hand-written phrase structure grammar. The probability model is structural and does not account for the probabilities of lexical dependencies. However, more context is taken into account than a PCFG, since the history that is considered at each parsing decision is conditional on the LR state, which can encode information in addition to the non-terminal being expanded. A dependency-based evaluation in Carroll, Minnen, and Briscoe (1999) shows that the latest version of the parsing system can identify some grammatical relations, such as subject and direct object, with high accuracy, but is less successful with other relations, such as the second object in a ditransitive construction and indirect object. The accurate identification of some relations, such as those corresponding to PP attachment, is likely to require a more lexicalised probability model.

A current version of the Briscoe and Carroll parser is used throughout this thesis. The parser is highly robust and has been used to provide large amounts of training data for the experiments reported in Chapters 3 and 4. It was also used for the parse selection experiments in Chapter 5 in order to provide the possible parses for a set of test sentences. A feature of the latest version is that the output is in the form of head dependency relations, which were used to create a dependency structure for each possible parse. In addition, the performance of the parser provided a useful benchmark against which to measure the performance of the dependency model.

Hektoen (1999) defines a probability model over logical forms, rather than syntactic structures, arguing that semantic relations are the key to accurate parse selection. A hand-written grammar was developed especially for this work, so that the requisite logical forms could be derived. A further novel aspect of the approach is that Bayesian estimation is used to estimate the parameters. Hektoen did attempt a direct comparison with PATTERN and Collins' conditional model, although the use of a hand-written grammar meant that only a subset of sentences from the Penn Treebank could be parsed. Also, Hektoen argues that the Parseval measures are not very suitable for his system, since they measure the ability of the system to parse a sentence, rather than the ability to select the correct parse.

... **C** ...

... **C** ...

The problem addressed in this chapter is how to estimate $p(\bar{s} | r)$ where \bar{s} is a sense in a semantic hierarchy \bar{s} is a predicate and r is an argument position. The term 'predicate' is used loosely here in that the predicate does not have to be a semantic object. bu

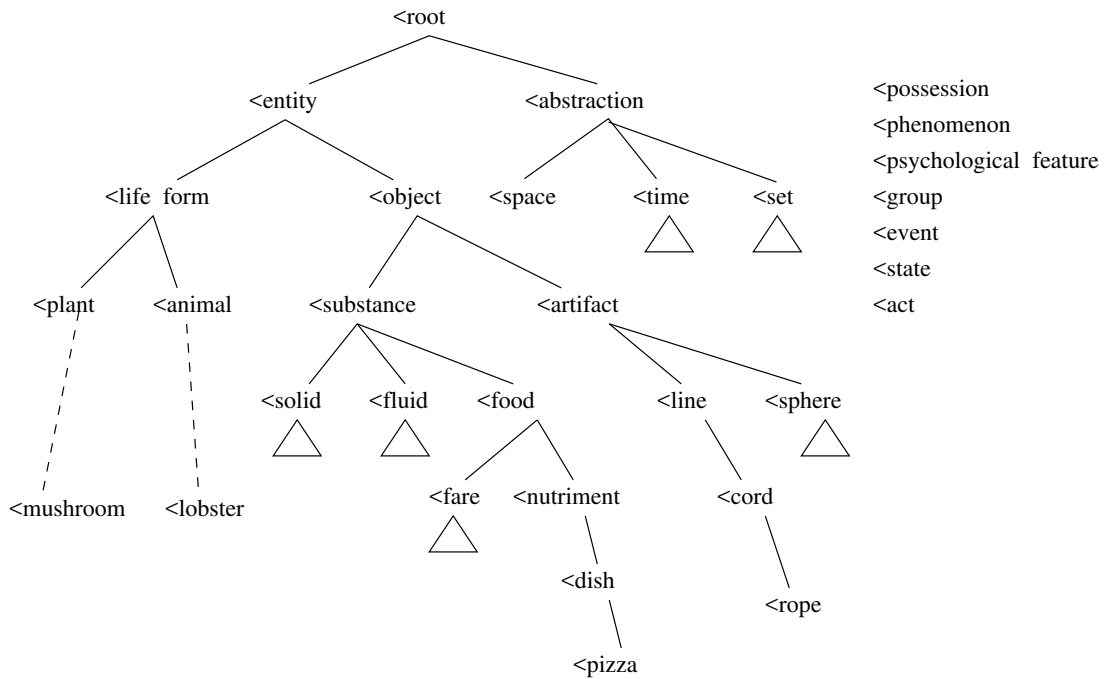


Figure Part of the WordNet hierarchy

concept \bar{c} and $cn(n) = \{ \bar{c} \mid n \in \text{syn}(\bar{c}) \}$ to denote the set of concepts that can be denoted by the noun n

The hierarchy has the structure of a directed acyclic graph although only around one percent of the nodes in the graph have more than one parent where the edges of the graph constitute what we call the direct-isa relation. Let $\text{isa} \subseteq \mathbf{C} \times \mathbf{C}$ be the transitive reflexive closure of direct-isa then $\bar{c} \text{ isa } \bar{c}$ implies \bar{c} is a kind of \bar{c} . If \bar{c} isa \bar{c} then \bar{c} is a *pon* of \bar{c} and \bar{c} is a *pon* of \bar{c} . In fact the hierarchy is not a single hierarchy but consists of nine separate sub hierarchies. The sub hierarchies are headed by the most general kind of concept and the roots of the sub hierarchies are shown in Figure which shows part of the WordNet hierarchy. The seven roots in addition

-- C ss s pro $\bar{1}\bar{1}$ s $\bar{1}$ $\bar{0}n$

non verbal predicates such as adjectives as well as verbs

$$\begin{aligned}
 p(\bar{c}, r) &= p(\bar{c} | r) \frac{p(r)}{p(\bar{c} | r)} & 4 \\
 &= \frac{p(r)}{p(\bar{c} | r)} \sum_{c' \in \bar{c}} p(c' | r) \\
 &= \frac{p(r)}{p(\bar{c} | r)} \sum_{c' \in \bar{c}} p(c' | r) \frac{p(c' | r)}{p(r)} \\
 &= \frac{1}{p(\bar{c} | r)} \sum_{c' \in \bar{c}} p(c' | r) \\
 &= \frac{1}{p(\bar{c} | r)} \sum_{c' \in \bar{c}} p(c' | r) \\
 &=
 \end{aligned}$$

Figure 2 Proof of proposition

compare the probabilities $p(\bar{c}, r)$ only. The proof of proposition is given in Figure and is explained in detail below.

The first line 2 applies Bayes theorem to the probability $p(\bar{c}, r)$. Line rewrites the probability $p(\bar{c} | r)$ as the sum of the probabilities of the sets dominated by the daughters of \bar{c} , $\sum_{c'} p(\bar{c}_{c'} | r)$, plus the probability of \bar{c} itself, $p(\bar{c} | r)$. This equality holds because the probability of a set of concepts $p(\bar{c} | r)$ has been defined in as the sum of the probabilities of the concepts in the set. However, note that the equality only holds in the tree case, and this is where the proofs in Figures 2 and differ. For a DAG, the probability of a set of concepts dominated by \bar{c} cannot be obtained by summing the probabilities of the sets dominated by the daughters of \bar{c} plus the probability of \bar{c} itself. The reason is that, in the sum $\sum_{c'} p(\bar{c}_{c'} | r)$, the probabilities of

$$\begin{aligned}
 p(\bar{\sigma}, r) &= p(\bar{\sigma} | r) \frac{p(r)}{p(\bar{\sigma} | r)} && 2 \\
 &= \frac{p(r)}{p(\bar{\sigma} | r)} \left(\sum_{\sigma} p(\bar{\sigma}_{\sigma} | r) + p(\sigma | r) \right) \\
 &= \frac{p(r)}{p(\bar{\sigma} | r)} \left(\sum_{\sigma} p(\bar{\sigma}_{\sigma} | r) \frac{p(\bar{\sigma}_{\sigma} | r)}{p(r)} + p(\sigma | r) \frac{p(\sigma | r)}{p(r)} \right) && 4 \\
 &= \frac{1}{p(\bar{\sigma} | r)} \left(\sum_{\sigma} p(\bar{\sigma}_{\sigma} | r) + p(\sigma | r) \right) \\
 &= \frac{1}{p(\bar{\sigma} | r)} \left(\sum_{\sigma} p(\bar{\sigma}_{\sigma} | r) + p(\sigma | r) \right) \\
 &=
 \end{aligned}$$

Figure Proof of proposition

su 4 o n 2 I IM true W H BPC ID f 2 2 2 2 Tm c 4 TJ 4 2 T

$\bar{\sigma}$	$(\bar{\sigma}; run, subj)$	$(\bar{\sigma}; subj) - (\bar{\sigma}; run, subj)$	$(\bar{\sigma}; subj) = \sum_{v \in V} (\bar{\sigma}; v, subj)$
$\langle bitch \rangle$		2 2	2
$\langle dog \rangle$	2	2 4 22	2
$\langle wolf \rangle$		4	
$\langle jackal \rangle$		2	2
$\langle wild_dog \rangle$			
$\langle hyena \rangle$	2		
$\langle fox \rangle$	2	2	2
	4		4

Table Contingency table for the children of $\langle canine \rangle$ in the subject position of run

senses but the data consist of nouns For now a simple approach is taken which is to estimate $(\bar{\sigma}, r)$ by distributing the count for each noun n in $\text{syn}(\bar{\sigma})$


```

    C top( $\bar{c}$ , r)
top ← c
sig_result ← false
CC parent  $\bar{c}_i$  gives lowest  $G^2$  value  $G^2_{\bar{c}_i}$ 
    not sig_result top  $\neq$  (root)
     $G^2_{\bar{c}_i} \leftarrow \infty$ 
    parents of top
    calculate  $G^2$  for sets dominated by children of parent
     $G^2 < G^2_{\bar{c}_i}$ 
     $G^2_{\bar{c}_i} \leftarrow G^2$ 
    parent  $\bar{c}_i \leftarrow$  parent
    chi squared test for parent  $\bar{c}_i$  is significant
    sig_result ← true
    move up to next node top ← parent  $\bar{c}_i$ 
return top

```

Figure 4 An algorithm for determining $\text{top}(\bar{c}, r)$

$\text{top}(\bar{c}, r)$

Figure gives an example of the procedure at work. Here $\text{top}(\langle \text{soup} \rangle, s, \text{obj})$ is being determined. The example is based on data from a subset of the BNC which had cases of an argument in the object position of s . The G^2 statistic is used together with an α value of . Initially top is set to $\langle \text{soup} \rangle$ and the probabilities corresponding to the children of $\langle \text{dish} \rangle$ are compared $p(s | \langle \text{soup} \rangle, \text{obj})$, $p(s | \langle \text{lasagne} \rangle, \text{obj})$, $p(s | \langle \text{haggis} \rangle, \text{obj})$ and so on for the rest of the children. The chi squared test results in a G^2 value of 4. compared to a critical value of . Since G^2 is less than the critical value the procedure moves up to the next node. This continues until a significant result is obtained which first occurs at $\langle \text{substance} \rangle$ when comparing the children of $\langle \text{object} \rangle$. Thus $\langle \text{substance} \rangle$ is the chosen level of generalisation.

Before giving some example levels of generalisation it is worth making some comparisons with the other WordNet approaches. First note that we have not made a uniform distribution as assumption as Li and Abe do (equation 2). Furthermore the problem described in Section 2 stemming from the fact that Li and Abe compare frequencies in order to generalise does not arise. This problem is avoided because we compare probabilities conditioned on sets of concepts rather than the frequencies of senses. And finally the generalisation procedure is able to return a suitable class for arguments that are negatively associated with some predicate (Section 2) explained how such arguments cause a problem for Resnik's approach. To see why consider applying the generalisation procedure to $\langle \text{location} \rangle$ in the object position of ; the procedure is unlikely to get as high as $\langle \text{entity} \rangle$ as we argued Resnik's approach is likely to do since the probabilities corresponding to the daughters of \langle

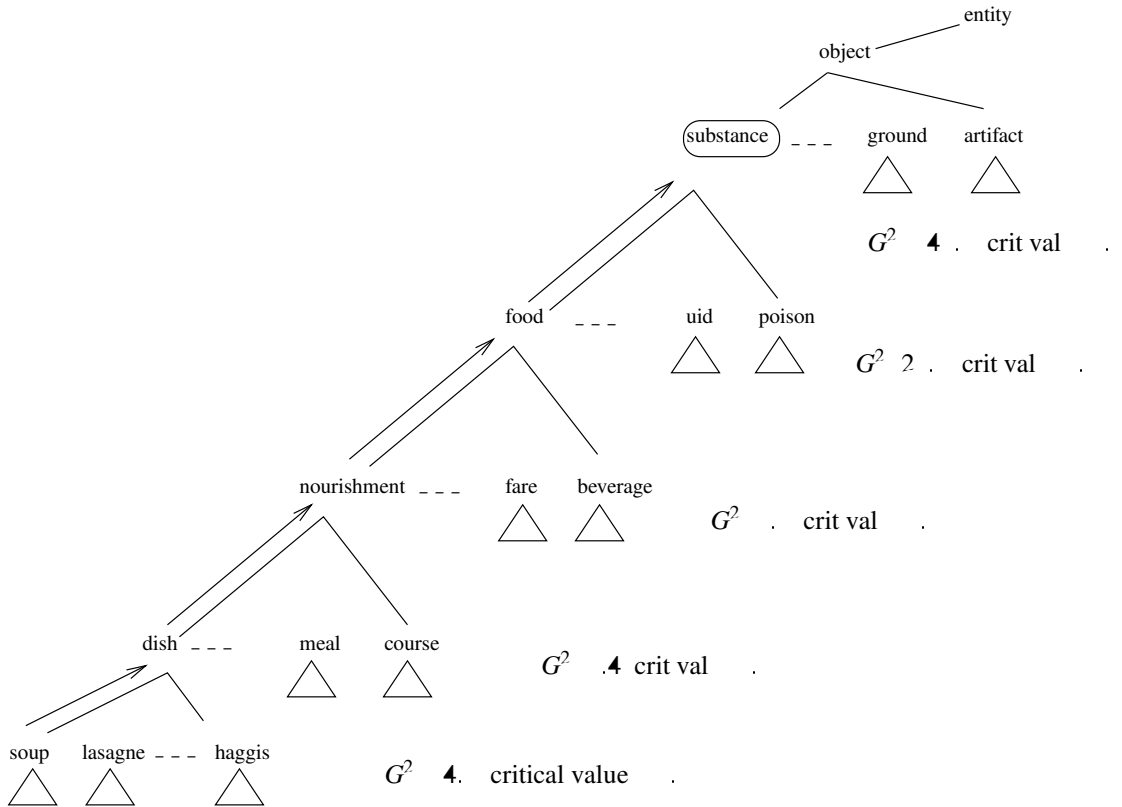


Figure An example generalisation determining $\text{top}(\langle\langle \text{soup} \rangle, s \rangle, \text{obj})$

1 **C** $\leftarrow \rightarrow \rightarrow \rightarrow \rightarrow \rightarrow$

In this section we show how the level of generalisation varies with the value for α and how

(\bar{s}, r)	α	
$\langle \text{coffee} \rangle, r_{\text{obj}}$.	$\langle \text{coffee} \rangle \langle \text{BEVERAGE} \rangle \langle \text{food} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$
$r_{\text{obj}} = 4$.	$\langle \text{coffee} \rangle \langle \text{BEVERAGE} \rangle \langle \text{food} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$
	.	$\langle \text{coffee} \rangle \langle \text{BEVERAGE} \rangle \langle \text{food} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$
	.	$\langle \text{coffee} \rangle \langle \text{BEVERAGE} \rangle \langle \text{food} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$
$\langle \text{hotdog} \rangle, \text{obj}$.	$\langle \text{hotdog} \rangle \langle \text{sandwich} \rangle \langle \text{snack_food} \rangle \langle \text{DISH} \rangle \dots \langle \text{RTfTmTj} \rangle \text{RTfTmTjRTfaTmTJRTcRTfTfR}$

\bar{s}, r	s, r	%	
$\langle \text{coffee} \rangle, r$	$\langle \text{coffee} \rangle, \text{obj}$	4	$\langle \text{coffee} \rangle \langle \text{BEVERAGE} \rangle \langle \text{liquid} \rangle \langle \text{fluid} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$ $\langle \text{coffee} \rangle \langle \text{BEVERAGE} \rangle \langle \text{liquid} \rangle \langle \text{fluid} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$ $\langle \text{coffee} \rangle \langle \text{beverage} \rangle \langle \text{liquid} \rangle \langle \text{FLUID} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle$ $\langle \text{coffee} \rangle \langle \text{beverage} \rangle \langle \text{liquid} \rangle \langle \text{fluid} \rangle \dots \langle \text{object} \rangle \langle \text{entity} \rangle \langle \text{ROOT} \rangle$
$\langle \text{hotdog} \rangle, r$	$\langle \text{hotdog} \rangle, \text{obj}$	4	$\langle \text{hotdog} \rangle \dots \langle \text{DISH} \rangle \langle \text{nourishment} \rangle \langle \text{food} \rangle \dots \langle \text{entity} \rangle$ $\langle \text{hotdog} \rangle \dots \langle \text{DISH} \rangle \langle \text{nourishment} \rangle \langle \text{food} \rangle \dots \langle \text{entity} \rangle$ $\langle \text{hotdog} \rangle \dots \langle \text{dish} \rangle \langle \text{NOURISHMENT} \rangle \langle \text{food} \rangle \dots \langle \text{entity} \rangle$ $\langle \text{hotdog} \rangle \dots \langle \text{dish} \rangle \langle \text{nourishment} \rangle \langle \text{food} \rangle \dots \langle \text{entity} \rangle \langle \text{ROOT} \rangle$
$\langle \text{Socrates} \rangle, r$	$\langle \text{Socrates} \rangle, \text{obj}$	4	$\langle \text{Socrates} \rangle \dots \langle \text{life_form} \rangle \langle \text{CAUSAL_AGENT} \rangle \langle \text{entity} \rangle$ $\langle \text{Socrates} \rangle \dots \langle \text{life_form} \rangle \langle \text{CAUSAL_AGENT} \rangle \langle \text{entity} \rangle$ $\langle \text{Socrates} \rangle \dots \langle \text{life_form} \rangle \langle \text{causal_agent} \rangle \langle \text{ENTITY} \rangle$ $\langle \text{Socrates} \rangle \dots \langle \text{life_form} \rangle \langle \text{causal_agent} \rangle \langle \text{entity} \rangle \langle \text{ROOT} \rangle$
$\langle \text{dream} \rangle, r$	$\langle \text{dream} \rangle, \text{obj}$	2	$\langle \text{dream} \rangle \dots \langle \text{preoccupation} \rangle \langle \text{cognitive_state} \rangle \langle \text{STATE} \rangle$ $\langle \text{dream} \rangle \dots \langle \text{preoccupation} \rangle \langle \text{cognitive_state} \rangle \langle \text{STATE} \rangle$ $\langle \text{dream} \rangle \dots \langle \text{preoccupation} \rangle \langle \text{cognitive_state} \rangle \langle \text{state} \rangle \langle \text{ROOT} \rangle$ $\langle \text{dream} \rangle \dots \langle \text{preoccupation} \rangle \langle \text{cognitive_state} \rangle \langle \text{state} \rangle \langle \text{ROOT} \rangle$
$\langle \text{man} \rangle, s$	$\langle \text{man} \rangle, \text{obj}$	4	$\langle \text{man} \rangle \dots \langle \text{mammal} \rangle \dots \langle \text{animal} \rangle \langle \text{LIFE_FORM} \rangle \langle \text{entity} \rangle$ $\langle \text{man} \rangle \dots \langle \text{mammal} \rangle \dots \langle \text{animal} \rangle \langle \text{LIFE_FORM} \rangle \langle \text{entity} \rangle$ $\langle \text{man} \rangle \dots \langle \text{mammal} \rangle \dots \langle \text{animal} \rangle \langle \text{LIFE_FORM} \rangle \langle \text{entity} \rangle$ $\langle \text{man} \rangle \dots \langle \text{mammal} \rangle \dots \langle \text{animal} \rangle \langle \text{life_form} \rangle \langle \text{entity} \rangle \langle \text{ROOT} \rangle$
$\langle \text{belief} \rangle, r$	$\langle \text{belief} \rangle, \text{obj}$	1	$\langle \text{belief} \rangle \dots \langle \text{cognition} \rangle$

α	%	%	%	%
.
.	2.	.	4.	.
.	2.	2.	4.	.4
.	.2	.	2.	.

Table The extent of generalisation for different values of α and sample sizes

α	G^2	X^2
.	.	.
.	2.	2.
.	2.	.
.	.2	.2

G^2 statistic The advantage of this test is that it can be applied to any contingency table irrespective of the size of the counts The main disadvantage is that it is computationally expensive especially for large contingency tables

What we have found in practice is that applying the chi squared test to tables with low counts tends to produce an insignificant result and the null hypothesis is not rejected This is especially true for the more conservative G^2 statistic The consequences of this for the generalisation procedure are that low count tables tend to result in the procedure moving up to the next node in the hierarchy This behaviour is clearly demonstrated in Tables 4 and But given that the purpose of the generalisation is to overcome the sparse data problem this behaviour is desirable and therefore we do not modify the test for tables with low counts

The next issue to consider is which statistic to use Dunning argues that G^2 is more suitable for corpus based linguistics than X^2 and Chapter 2 described Dunning's experiment comparing the use of X^2 and G^2 to identify highly associated bigrams Dunning's claim is that for small samples the sampling distribution of G^2 is a better approximation to the chi squared distribution than the sampling distribution of X^2 However in Chapter 2 we presented a quotation from Agresti which contradicts this claim A more likely explanation lies in the conservative nature of G^2 which means that X^2 is more likely to return a significant result for a table based on small counts This would explain Dunning's bigram results in which pairs of words occurring infrequently in the corpus obtain high scores according to X^2 but not G^2

Note that for some applications it may make little difference to the performance whether G^2 or X^2 is used The results for a PP attachment task described in Chapter are very similar for both statistics In fact the use of X^2 may even lead to better results for some applications The results of a pseudo disambiguation task also described in Chapter

plenty of counts; and since the point of this work is to overcome the sparse data problem the second consideration should override the first. The chi squared test has this overriding effect built in automatically particularly when using the conservative G^2 statistic since it measures the

This may appear to be a crude solution to the problem of ambiguous data but in practice it works surprisingly well. The reason is that counts for sets of concepts tend to accumulate in the right places. To see why consider this example adapted from Resnik. Resnik notes that a similar point is made by Yarowsky. Consider estimating probabilities for the object position of the verb $r_{\vec{h}}$ and suppose that $r_{\vec{h}}$ \vec{h} and $r_{\vec{h}}$ r occur as part of the data. The word r is a member of seven senses in WordNet and \vec{h} is a member of two senses. Thus for these data items splitting the count equally leads to each sense of r receiving .4 counts and each sense of \vec{h} . counts. But note that with regard to s s of concepts only those sets containing senses of both \vec{h} and r such as $\langle \text{beverage} \rangle$ will accumulate counts. The counts for the incorrect senses will be randomly dispersed throughout the hierarchy as noise and areas where counts would be expected to accumulate such as under $\langle \text{beverage} \rangle$ in this example will receive the majority of the overall count. As will be shown later this accumulation effect means that performance in applications can be good even if this simple estimation technique is used.

However there is an obvious problem with this approach although counts for sets tend to accumulate in the right places counts can be greatly underestimated. In the previous example $(\langle \text{beverage} \rangle, r_{\vec{h}}, \text{obj})$ is incremented by only .4 counts from the two data instances rather than the correct value of 2. In addition as Resnik himself notes the accumulation process has less effect on sets of concepts low down in the hierarchy since here the counts have had less chance to accumulate. The example Resnik gives is for o nos . In this case counts would be expected to be higher for the set dominated by the bodily sense of nos rather than the aircraft sense. However since both senses are low down in the hierarchy splitting counts equally is likely to lead to a similar count for each set. For the same reason counts for individual concepts as opposed to sets of concepts are likely to be inaccurate.

In response to this we note that the accumulation of counts leads to an obvious strategy use the fact that correct senses are likely to be members of sets where counts have accumulated as a way of re distributing the count. Continuing with the $r_{\vec{h}}$ \vec{h} example \vec{h} has a beverage sense and a colour sense in WordNet. If the above strategy is used equal counts will be given to each sense on the first iteration but on subsequent iterations more of the count will be given to the beverage sense. This is because counts would accumulate under $\langle \text{beverage} \rangle$ for the object position of $r_{\vec{h}}$ and not under $\langle \text{colour} \rangle$.

One issue to consider is how to determine a representative set for a concept. We have been assuming that $\langle \text{beverage} \rangle$ and $\langle \text{colour} \rangle$ are suitable for the two senses of \vec{h} but a procedure is needed which determines this automatically. The procedure needs to find a hypernym for each alternative sense such that the hypernym is high enough for counts to have accumulated in the set dominated by the hypernym; however it should not be so high that the alternative senses cannot be distinguished. An example of a hypernym that is too high is $\langle \text{root} \rangle$ the notional root of the hierarchy since if $\langle \text{root} \rangle$ were chosen for both senses of wine there would be no way to distinguish between the senses. Another reason not to go too high is that the sets need to be in some sense representative of the senses. Suppose \vec{u} \vec{h} occurs in the data and the food sense of \vec{u} \vec{h} and the electronic sense need to be distinguished. It would not be appropriate to represent the electronic sense using $\langle \text{entity} \rangle$ since this would not capture the rto 2 t 2 2 h a 2

$$A(C, r) = \frac{p(C|r)}{p(C)}$$

$$p(C|r) = \frac{(C, r)}{(r)}$$

$$p(C) = \frac{\sum_{v \in V} (C, v)}{\sum_{v \in V} (v)}$$

$$(C, r) = \sum_{i \in C} (i, r)$$

Figure 4.2 Estimates for calculating $A(C, r)$ for a set of concepts C ; V is the set of verbs in the data

entity is not homogeneous with respect to the object position of drink some entities are drunk some are not In contrast the set abstraction is fairly homogeneous in that on the whole kinds of abstraction are rarely drunk

The set beverage is also homogeneous which is a suitable representative for the beverage sense Note that the two sets abstraction and beverage are also maximally homogeneous in that the sets dominated by the parents of beverage and abstraction liquid and root respectively are not themselves homogeneous This motivates the idea that we should be looking for maximally homogeneous sets maximal because we want to allow counts to accumulate and noise to be dispersed The problem with using colour as a representative of wine is that colour is not high enough for this dispersal to have occurred

One way to recognise that liquid is not homogeneous is to note that the sets dominated by the daughters of liquid are associated to differing degrees with drink some liquids are drunk such as beverages liquor and water but some are not such as ammonia antifreeze and sheep

the verb. Thus it appears that the procedure can be applied directly to the problem of determining $[\bar{c}, r]$

However there are some differences between the problems being addressed in this and the previous chapter. In the previous chapter the problem was to find a generalisation level that would lead to a reasonable probability estimate. In this chapter the problem is to find a level where counts have accumulated and the noise dispersed sufficiently. A solution to both problems lies in finding homogeneous sets; the difference lies in the r of homogeneity that is likely to be optimal in each case. For the probability estimation problem it may be that the difference in association norms needs to be relatively small for a class based probability estimate to be a useful estimate. Results presented in Chapter 3 suggest that for some disambiguation tasks this is indeed the case. Another way to think of this is that for some tasks the optimal level of generalisation is quite low in the hierarchy on the whole. In contrast the re-estimation problem is likely to favour a level of generalisation that is quite high on the whole since it is here that counts have accumulated and noise dispersed.

Despite these differences the procedure can be adapted to both problems. The degree of homogeneity required can be controlled by the parameter α the level of significance of the chi squared test. The value of α controls the overall level of generalisation: a high value for α results in a low level of generalisation on the whole and a low value for α results in a high level of generalisation. Results from the previous chapter clearly demonstrate this. One way to set a value for α would be to estimate counts using a range of α values and use a held out test set to choose those counts that give the best performance on the task in hand.

Another useful feature of the procedure within the context of the re-estimation problem is that it employs a significance test to find homogeneous sets. This implies that the procedure automatically finds areas where counts have accumulated since it is only here that there will be enough data to return a significant result for the chi squared test. This point is especially true when the more conservative G^2 statistic is used and a low value for α .

As a final comment a point of clarification is needed. The previous chapter showed that the chosen level of generalisation is dependent on the size of the data sample as well as on the value of α . Thus the notion of homogeneity being used here is not an absolute notion but a relative one relative to the sample. If the procedure determines a maximally homogeneous set that does not accord with intuition this should not be automatically considered a failure. A comment in Clark and Weir states that $\langle \text{food} \rangle$ is heterogeneous with respect to the object position of

$$\begin{aligned}
 p(\overline{\langle \text{food} \rangle} | \text{ , obj}) &= \frac{(\overline{\langle \text{food} \rangle}, \text{ , obj})}{(\text{ , obj})} \\
 &= \frac{2}{4} \\
 &= .5
 \end{aligned}$$

$$\begin{aligned}
 p(\overline{\langle \text{food} \rangle} | \text{obj}) &= \frac{(\overline{\langle \text{food} \rangle}, \text{obj})}{(\text{obj})} \\
 &= \frac{2}{2} \\
 &= 1
 \end{aligned}$$

$$\begin{aligned}
 A(\overline{\langle \text{food} \rangle}, \text{ , obj}) &= .5 / 1 \\
 &= .5
 \end{aligned}$$

Figure 4.4 Calculation of $A(\overline{\langle \text{food} \rangle}, \text{ , obj})$

high value for $(\overline{\langle \text{entity} \rangle}, r)$ and so $p(\text{cajole} | \overline{\langle \text{entity} \rangle}, \text{obj})$ is not over estimated

The conclusion is that if the association norm is to be applied appropriately it should be applied to frequent verbs or to sets for which (C, r) is reasonably high; however since the re-estimation procedure relies on using sets where plenty of counts have accumulated this should not be a problem

1. Introduction

There are two evaluations in this section⁴ The first shows how the estimated counts change

1

C

C

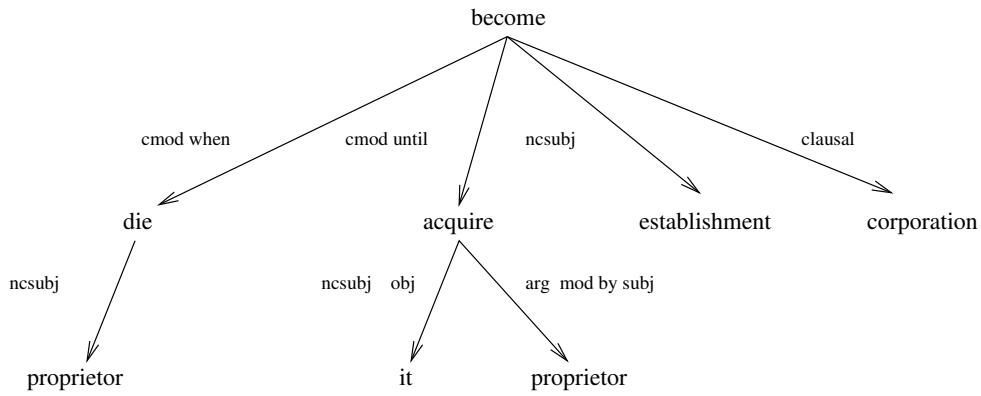


Figure Example dependency structure for the sentence *n proprietor or \bar{h} s s*
 \bar{s} n s ou \bar{o} \bar{c} orpor \bar{h} n un \bar{h} \bar{s} \bar{q} u \bar{h} no r propr \bar{h} or

\bar{o} un \bar{h} \bar{q} u \bar{h} the establishment should \bar{o} a corporation un \bar{h} it is \bar{q} u \bar{h} by another proprietor Here \bar{o} is the head in both cases and \bar{h} and \bar{q} u \bar{h} are dependents The prepositions \bar{h} and un \bar{h} introduce the dependents

- ncsubj denotes a non clausal subject The ncsubj examples simply encode a head and dependent except that the passive \bar{h} \bar{s} \bar{q} u \bar{h} is recognized as such by the symbol obj This appears in the triple labelling the edge \bar{q} u \bar{h} \bar{h} and indicates that \bar{h} is an underlying object of \bar{q} u \bar{h}
- arg_mod

2 C p r _ n r h Es ! on 6 n q u s h o rs 6 on s

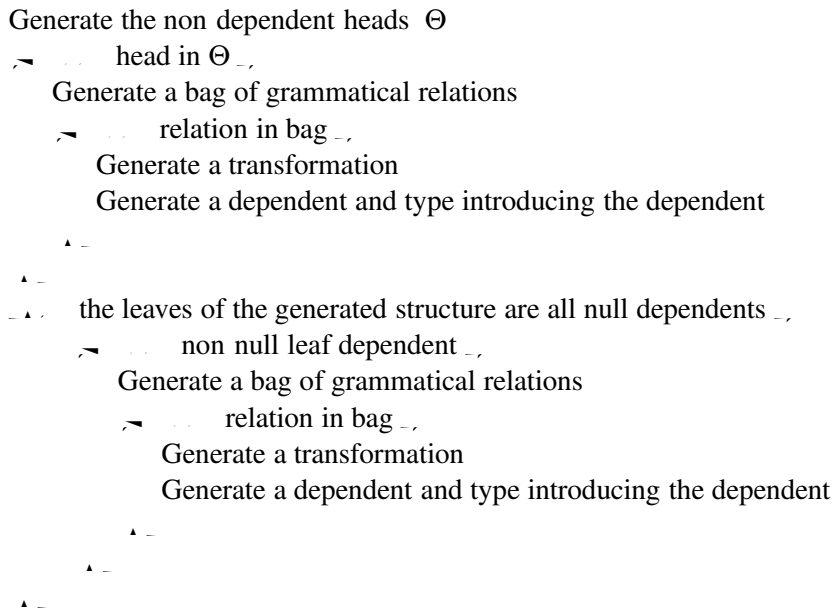


Figure 2 Sequence of decisions generating a dependency structure

The dependency structure with the highest probability is chosen as the correct structure together with the corresponding parse if necessary. The conditioning context ... is known as the history and is equivalent to the structure built up to that point. In order that the model have a manageable number of parameters a function Φ

C p r $_$ n r \bar{c} Es \bar{c} \bar{c} n \bar{c} u s \bar{c} o rs \bar{c} \bar{c} n s

$p(\bar{c}, |, r)$ where \bar{c} is a nominal dependent

The probabilities corresponding to the above examples are

- $p(\bar{c} = p \text{ ron} | p \bar{c}, \text{iobj})$
- $p(\bar{c} = r \bar{c} \text{ or } | r \bar{c} \bar{c}, \text{ncmod})$
- $p(\bar{c} = \text{on } \bar{c} | \bar{c}, \text{ncmod})$

Again the sense of \bar{c} is chosen which maximises the probability estimate and $p(\bar{c}, |, r)$ is used as a proxy for $p(\bar{c}, |, r)$ where \bar{c} is determined as follows

$$\bar{c} = \arg \max_{\bar{c} \in \text{cn}(\bar{c})} p_{s\bar{c}}(\bar{c}, |, r)$$

The class based approach can be used to obtain $p_{s\bar{c}}(\bar{c}, |, r)$ by first applying Bayes theorem and then conditioning on an appropriate set of concepts as before. The only difference is that the conditional probability of \bar{c} is now joint with

$$p(\bar{c}, |, r) = p(\bar{c}, | \bar{c}, r) \frac{p(\bar{c} | r)}{p(\bar{c} | r)}$$

$$\approx p(\bar{c}, | \bar{c}, r)$$

The set \bar{h} is obtained by applying the procedure described in Chapter and the probability $p(\bar{h}, r)$ is estimated using relative frequencies. If the head does not appear in WordNet, an estimate of $p(\langle \text{root} \rangle, r)$ is used, unless the head is a pronoun or proper name. If the head is a pronoun \bar{h}

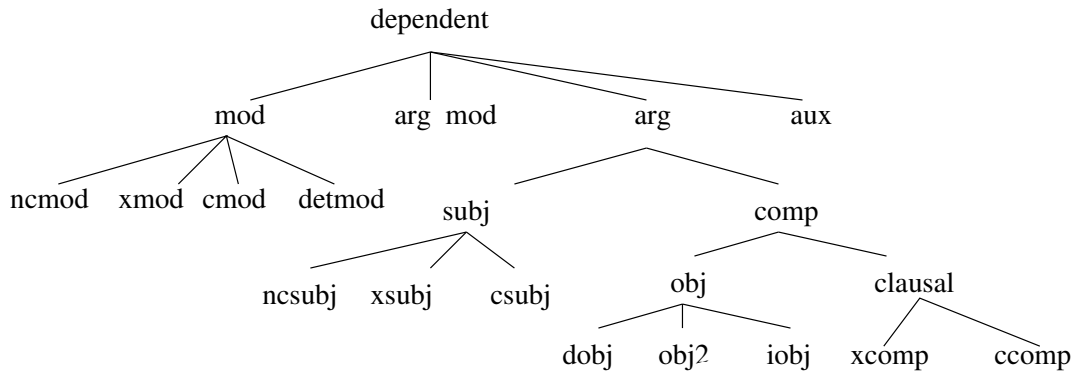


Figure The grammatical relations used in the implementation

1 C C

1 C C

The parser used for the evaluation is a more developed version of that described in Carroll and Briscoe. This version is able to produce output in the form of grammatical relations, which is the main reason the parser was chosen. The parser produces a set of parses for a sentence together with the corresponding sets of grammatical relations. Thus we were able to create a dependency structure for each parse and choose the parse with the most probable structure. A further advantage in using this parser is that there exists a manually created test suite which uses the same grammatical relation scheme as used by the parser (Carroll et al., 2000); this test suite was used for the evaluation.

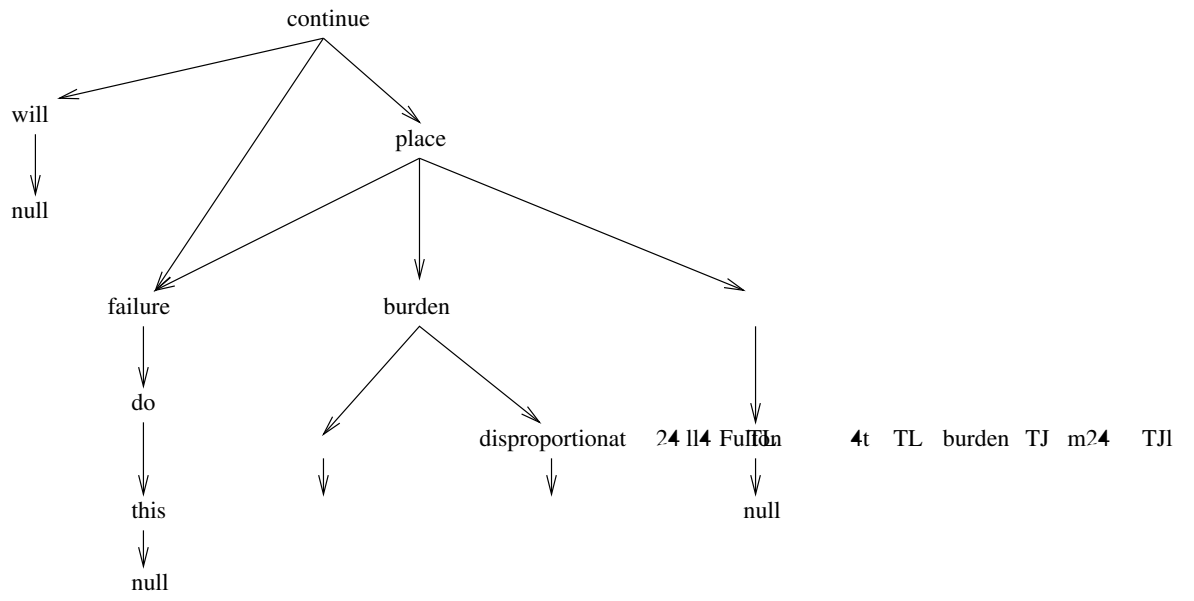
The relations used by the parser can be arranged in a hierarchy as shown in Figure 1. If the parser is unable to determine the precise nature of the relation and thus cannot return a relation at a leaf node, a more generic relation can be returned. Each relation is described in detail in Appendix A based on the descriptions given in Carroll et al. (2000a) and Carroll et al. (2000b). A brief description of each relation is given below.

-

C p r _ n r h Es n n q u s h o rs n s

```
(|ncsubj| |continue:6_VV0| |failure:1_NN1| _ )  
(|clausal| _ |continue:6_VV0| |place:8_VV0|)  
(|ncsubj| |place:8_VV0| |failure:1_NN1| _ )  
(|dobj| |place:8_VV0| |burden:11_NN1| _ )  
(|iobj| |on:12_II| |place:8_VV0| |tax-payer:14_NN2|)  
(|dobj| |do: _VD0| |this:4_DD1| _ )  
(|xcomp| |to:2_T0| |failure:1_NN1| |do: _VD0|)  
(|ncmod| _ |burden:11_NN1| |disproportionate:10_JJ|)  
(|ncmod| _ |tax-payer:14_NN2| |Fulton:1 _NP1|)  
(|detmod| _ |burden:11_NN1| |a:15_AT1|)  
(|aux| _ |continue:6_VV0| |will:16_VM|)
```

Figure Example parser output for the sentence



obtained from John Carroll who ran the parser over around million words of the BNC from around , sentences The parser output was in the same form as that given in Figure and the output was processed in the following way the formulaic expressions such as sums of money were found using simple regular expressions

- 4 digit numbers beginning or 2 were replaced with the word on
Numerical expressions were replaced with n q n
Monetary expressions not in WordNet were replaced with s o - on
Expressions denoting people not in WordNet such as Dr were replaced with so on
Expressions denoting companies not in WordNet such as Ltd were replaced with n p n
- Verbs and prepositions were reduced to lower case
- All words were lemmatized

The formulaic expressions were replaced with these particular words because each word has only one sense in WordNet and belongs to a relevant synset

Some parts of the data are much more accurate than others Table in the next section

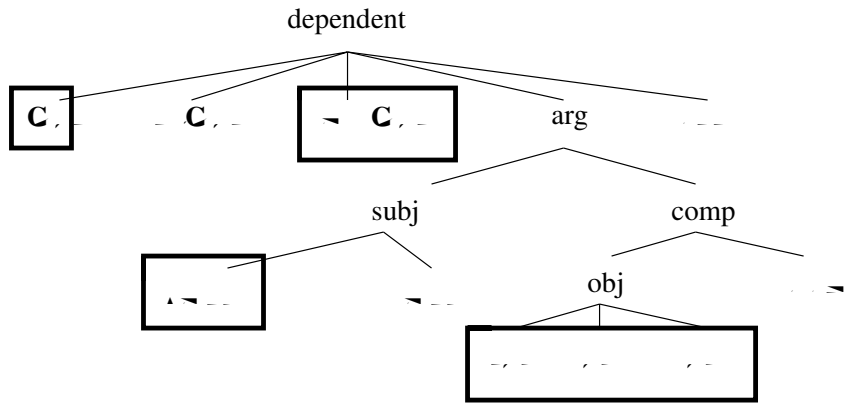


Figure Dependency probabilities by relation that can be estimated using WordNet

is half covered by a box because not all of the mod cases can be estimated using WordNet. For the test suite used for the evaluation, approximately 50% of the grammatical relations correspond to parameters that can be estimated using WordNet. The parameters corresponding to the remaining relations were estimated using the linear interpolation method.

1.2.2. Test Suite

The test suite consists of 1000 sentences taken from the *rusanne* corpus, covering a number of written genres and manually annotated with grammatical relation information.

4 C p r _ n r n Es n n q u s n o r s n s

Relation	occurrences	% occurrences
dependent		.
mod		.
ncmod	24	.2
xmod	2	2.
cmod	2	.2
detmod	24	.2
arg_mod	4	.
arg	2	.2
subj	4	.
ncsubj		.
xsubj		.
csubj		.
comp		.
obj		.
doobj	4	.
obj2		.
iobj		2.4
clausal	4	.2
xcomp	2	4.
ccomp		.2
aux		.
conj	4	2.

Table Frequency of each type of relation in the test suite

structures The model is likely to prefer incomplete structures with a small number of relations because in these cases less probabilities are multiplied together to get a total probability for the dependency structure

The dependency structures were processed in similar ways to the data in that each word was lemmatized and formulaic expressions were replaced with words in WordNet as described in Section 2 Because there is only a small amount of data in the test set we did not use any of it as held out data and the various parameters were selected by hand The parameters δ and ϵ described in Section 2.2 were set to , and respectively and the level of significance for the chi squared test α was set to . The results appear at t H BPC ID EI Qq 24f signi c

(t)3r0(h)5(e)993(i)3.0011(n)-36(t)2.t 10986(s)4(e)-227347.986(s)(s)901.0986(n)(t)7(e)3205.239250(s)a2(2240)9(d)16(s)198

C p r _ n r ĩ Es ĩ on ĩ n ĩ u s ĩ o rs ĩ on s

Relation	Precision %	Recall %	F score	GRs
dependent	.2	.	.	
mod	.	.2	.2	
ncmod	.	.4	.	2
xmod	.	.24	.4	4
cmod	.	.2	.	
detmod	.	.2	.2	2
arg_mod				
arg	.	.	.2	2
subj	.	.2	.	2
ncsubj	.	.	.4	
xsubj	.	.2	.	
csubj				2
comp	.	.	.4	
obj	.	.	.	
dobj	.	.22	.2	42
obj2	.2	.	.4	4
iobj	.2	.4	.42	
clausal	.2	.2	.	
xcomp	.4	.2	.2	Tf

22 Tm

42 4 4

The treatment of word sense ambiguity is another area that could be improved. Currently a rather cavalier approach is taken which is to select the sense that maximises the relevant probability estimate. One promising approach is to try and integrate the word sense disambiguation into the parsing model and perform the two simultaneously as Bikel [2] has attempted to do.

A tentative conclusion of this chapter is that the use of lexical sense preferences or selectional preferences alone is unlikely to lead to a highly accurate parse selection system. Even the successful statistical parsing models such as those of Collins and Charniak [2] which rely heavily on lexical information also make use of the structural properties of a parse. One way to extend this work would be to try and combine the dependency model with the structural model of Briscoe and Carroll.

As an evaluation of the class based estimation technique the results are inconclusive since the parse selection problem may not be a good way to isolate the performance of the WordNet estimation techniques. In order to have a more focused evaluation the method of estimation is applied to two disambiguation tasks that can be tackled using only parameters relating to lexical sense preferences; moreover the parameters can be estimated using reliable data. These tasks are presented in the next chapter.

C p r A u s o u n o p r s s s s n n s

For these examples it is hard to see that there is an ambiguity at all but the attachment problem assumes that any *r np pr p np* sequence results in an ambiguity In it is assumed that *o p n* could attach to ; in *4 n o n or* could attach to *u s*; and in *sp u n* could attach to *n s s*

Another reason why the telescope and stick examples are misleading is that they imply the PP attachment problem as we have defined it is harder than it really is For these two examples either attachment results in a plausible semantic reading and the correct reading depends on the wider context In a commonly cited paper Altmann and Steedman argue that the resolution of attachment ambiguities requires a model where the relevant entities are represented and reasoned about This argument led Hindle and Rooth to comment that if this is typical of PP attachment ambiguities then there is little hope of building computational models to solve the problem at least in the near future

Clearly some account of context is required for the resolution of some cases of attachment ambiguity However this may only apply to a small subset of cases The three treebank examples can be resolved without resorting to the wider context; in fact they can be resolved without even considering *n₂ ui*

The estimates $p_{s\bar{\sigma}}(\bar{\sigma}, pr|)$ and $p_{s\bar{\sigma}}(\bar{\sigma}_n, pr|n)$ are obtained using the method described in Chapter First Bayes rule is applied and then probabilities are conditioned on a set of concepts where appropriate The formulae are given for $p(\bar{\sigma}, pr|_{\mathbb{N}}$

6 - SO 2h 6 n 211111S



4 C p r A u s o n o s s s i n n s

α value	% correct G^2	% correct X^2
.	(4 cases	(cases
.	(, 2 cases	

$$\max_{\vec{s} \in \text{cn}(n)} p_{s\tau}(\vec{s}), \text{obj}$$

Tf

2 2

Time

ax

TJR

Tf 4 2

2 Td

Generalisation technique	% correct	av gen	sd gen
Similarity class			
$\alpha = .$.	.	2.
$\alpha = .$.4	2.	.
$\alpha = .$.	2.4	.
$\alpha = .$.	.	.
$\alpha = .$.	.2	.2
Low class	.	.	.
MDL	.	4.	.
Assoc	.	4.2	2.

Table 1 Results for the pseudo disambiguation task

it as a noun-noun sense pair. For example, the two instances of *wo* in the synsets $\{wo\}$ and $\{wo-h, wo-h, wo, sno, C\}$ are treated as separate nouns. We use $sep(n)$ to denote the set of separate instances of n in WordNet.

Adopting the MDL approach, the disambiguation decision was made as follows: p is used to denote an estimate using the MDL approach:

$$\max_{n' \in sep(n)} p(n'),$$

α value	% correct G^2	% correct X^2
.	(.)	4. (.)
.	.4 (2.)	. (2.)
.	(2.4)	4. (2.2)
.	(.)	4. (.)
.	(.2)	. (.2)

Table Disambiguation results for G^2 and X^2

important feature of these results is that the α values corresponding to the lowest scores lead to a significant amount of generalisation. This explains why the α

This Chapter considers each of the problems that have been addressed in this thesis outlining the proposed solution for each problem together with the original contribution. The ways in which the work could be extended are also considered. The discussion is organised by chapter.

Resnik [10] considered the problem of how to estimate the probability of a noun sense given a predicate and argument position. The proposed solution answers two questions: one, how to use a class from WordNet to estimate the probability of a noun sense, thereby overcoming the sparse data problem; and two, how to select a suitable class to represent a sense. The second question can be thought of as how to select a suitable level of generalisation in WordNet. The proposed generalisation procedure employs a chi squared test and the level of significance of the test α is treated as a parameter to be set empirically. Results were given showing how the chosen level of generalisation depends on both the sample size and the value of α .

The generalisation procedure is arguably the most important contribution of the thesis. As Resnik [10] comments, "It has been widely noted that the selection of an appropriate level of abstraction is a difficult problem" [p. 10]. We have tried to devise a procedure that has a clearer statistical interpretation than that of Resnik and also one that overcomes some of the shortcomings of Li and Abe's approach, such as the uniform distribution assumption [2]. An advantage of our approach is that treating α as a parameter gives the procedure a level of flexibility since α can be set to produce a level of generalisation that is appropriate for the task in hand.

An alternative to using a single class to estimate the probability of a concept, which was suggested by Jason Eisner at COLING [2], is to use all the classes dominated by the hypernyms of a concept. An estimate would be obtained for each hypernym and the estimates combined in a linear interpolation. An approach similar to this is taken by Bikel [2] in the context of statistical parsing.

Resnik [10] described an unsupervised reestimation algorithm for estimating sense frequencies. We first explained how splitting the count for a noun equally among its senses works better than might be expected, at least for the frequencies associated with sets of senses. The reason is that counts tend to accumulate in the right places in WordNet, namely for sets of senses that are positively associated with the predicate. This accumulation effect motivated the reestimation algorithm, in which the count for a noun is split equally on the first iteration, but on subsequent iterations more count is given to those noun senses that belong to positively associated sets. A feature of the algorithm is that it employs the generalisation procedure described in Chapter 4 and this led to a new interpretation of the procedure as one that finds sets of semantically similar senses, or homogeneous sets of senses, in the hierarchy. The results on a pseudo disambiguation task showed that the reestimation can be beneficial in some cases.

The performance of the reestimation algorithm is limited by the fact that highly accurate WSD is unlikely to be achieved using preferences alone. Other work that has attempted to use prefer

Corpus Conjunction

ences for sense disambiguation has achieved little success Resnik ; Carroll and McCarthy
2 Thus one way to further this work would be to see how other knowledge sources could
be used to aid the reestimation The surrounding context of a noun is an obvious source of addi
tional information There also needs to be more research int

the original method of Hindle and Rooth. It was discovered that in order to perform well the disambiguation method requires more training data than currently exist in treebanks but that with appropriate amounts of data the method is highly accurate. It was also shown that the generalisation procedure introduced in Chapter 4 outperforms a simple approach of choosing a fixed level in the hierarchy.

A further evaluation using a pseudo disambiguation task showed that our class based estimation method outperforms two alternative approaches based on the work of Resnik (1999) and Li and Abe (2001). It was discovered that the alternative methods appeared to be over generalising at least for this task. As we have argued a useful feature of our estimation procedure is that the level of significance used in the chi squared test α can be used to guard against over or under generalisation. But even when the results did vary with α our method was found to outperform the alternatives across the whole range of α values.

A further useful result was that the performance on the task was at least as good when using the Pearson chi squared statistic as when using the log likelihood chi squared statistic. This result is at odds with the currently accepted wisdom that the log likelihood chi squared statistic is a better statistic for use in corpus based NLP. We suggested an explanation for this finding which also explains the results of Dunning (1994) who initially argued for the use of the log likelihood statistic.

An important question that has yet to be addressed in the literature is whether class based estimation methods perform better when the classes are automatically acquired or when they are part of a man made hierarchy. One way to investigate this would be to perform the pseudo disambiguation task but using clustering algorithms to estimate the probabilities. Pereira et al (1999) and Rooth et al (2001) have already used a similar task to evaluate their clustering algorithms; the results depended on the number of clusters induced and ranged between 70% and 80% for both approaches compared to the 85% reported here. Unfortunately different test and training data were used in each case and so it is difficult to draw any conclusions from these results. A related issue is how the structure of WordNet affects the accuracy of the probability estimates. We have taken the structure of the hierarchy for granted without any analysis but it may be that an alternative design would be more conducive to probability estimation.

— — — — —

4 Bibliography

Charniak E. 2000. A maximum entropy inspired parser. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, pp. 132-139. Seattle, WA.

Chen Y. and Goodman J. 2009. An empirical study of smoothing techniques for language modeling. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

Collins M and Brooks J Propositional phrase attachment through a backed off model
In *Proceedings of the 2005 Conference on Empirical Methods in Natural Language Processing* (EMNLP 2005), Cambridge MA

Cover T and Thomas J *Elements of Information Theory* Wiley New York

Daelemans W Van Den Bosch A and Zavrel Z Forgetting exceptions is harmful in
language learning *Proceedings of the 2005 Conference on Empirical Methods in Natural Language Processing* 4

Dagan I Lee L and Pereira F Similarity based models of word cooccurrence probabilities
Proceedings of the 2005 Conference on Empirical Methods in Natural Language Processing 4

Dagan I Marcus M and Markovitch S Contextual word similarity and estimation from
sparse data *Computational Linguistics* 28(2)

Drange T *Crosslingual NLP: Bridging the Gap* Borror Aronson *Proceedings of the 2005 Conference on Empirical Methods in Natural Language Processing* Mouton The Hague

Dunning T Accurate methods for the statistics of surprise and coincidence *Computational Linguistics* 4

Eisner J An empirical comparison of probability models for dependency grammar *Technical Report*

B. J. P.

Good I J The population frequencies of species and the estimation of population parameters
B. J. P. 4 2 2 4

Goodman J Probabilistic feature grammars In *Proceedings of the Conference on Computational Linguistics*
or Symposium on Computational Linguistics pp Boston MA

Grishman R Macleod C and Meyers A 4 Complex syntax Building a computational
lexicon In *Proceedings of the Conference on Computational Linguistics*
pp 2 2 2 Kyoto Japan

Harrison P Abney S Black E Flickinger D Gdaniec C Grishman R Hindle D Ingria
B Marcus M Santorini B and Trzalkowski T

Joshi A and Chabes Y 2 Tree adjoining grammars and lexicalized grammars In Nivat M and Podelski A Eds *Discrete Natural Language Processing* Elsevier Princeton NJ

Katz J and Fodor J 4 The structure of a semantic theory In Fodor J and Katz J Eds *Semantics: A Reader* chap pp 4 Prentice Hall Englewood Cliffs NJ

Katz J Estimation of probabilities from sparse data for the language model component of a speech recognizer *IEEE Transactions on Acoustics, Speech, and Signal Processing* 4 4

Kilgarriff A Which words are particularly characteristic of a text A survey of statistical approaches In *Corpus Linguistics: A Reader* pp 4 Sussex UK


Kilgarriff A and Rose T Measures for corpus similarity and homogeneity In *Corpus Linguistics: A Reader* pp 2 Granada Spain

Kimball J Even principles of surface structure parsing in natural language

Bibliography

- Lin D Automatic retrieval and clustering of similar words In *Roberts S O*
Annals of Association for Computational Linguistics pp Montreal
Canada
- Magerman D *A probabilistic model for parsing* Ph D thesis
University of Stanford
- Magerman D *Statistical decision tree models for parsing* In *Roberts S O*
Annals of Association for Computational Linguistics pp 2 2 Cambridge
MA

- Ribas F. On learning more appropriate selectional restrictions. In *Proceedings of the European Conference on Computational Linguistics*, Dublin, Ireland, 2002. pp. 2-11.
- Riezler F., Prescher D., Kuhn J. and Johnson M. Lexicalized stochastic modeling of constraint based grammars using log linear measures and EM training. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, Hong Kong, 2002. pp. 4-11.
- Rooth M., Riezler F., Prescher D., Carroll G. and Beil F. Inducing a semantically annotated lexicon via EM based clustering. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, University of Maryland MD, 2002. pp. 4-11.
- Sampson G. *English as a Second Language*. Oxford University Press, Oxford, UK, 2002.
- Steedman M. *Phrasal Structure*. The MIT Press, Cambridge, MA, 2002.
- Tetina J. and Nagao M. Corpus based PP attachment.



Some of the descriptions given here are taken directly from Carroll et al (1986) and the same notation is used. Many of the examples also come directly from that paper.

C The relation between a head and a modifier; **mod** is used to indicate the word introducing the dependent where appropriate. Examples include the following:

- | | |
|------------------------------|--------------------------------|
| mod a red ag | a red ag |
| mod with walk John | walk with John |
| mod while walk talk | walk while talking |
| mod a Picasso painter | Picasso the painter |
| mod of examination patient | the examination of the patient |

The relation between a predicate and a non clausal subject; where appropriate is obj after passivisation; for example

ncsubj arrive John _ John arrived in Paris
 ncsbj employ Microsoft _ Microsoft employed C programmers
 ncsbj employ Paul obj Paul was employed by IBM

The relation between a predicate and a clausal subject controlled from within and from without respectively; for example

csubj mean leave _ that Nellie left without saying good bye meant she was angry
 csubj astonish owe _ that he owed anything would have astonished his mother
 xsubj require win _ to win the America s Cup requires heaps of cash

The relation between a predicate and a direct object; where appropriate is after dative shift; e g

dobj read book _ read books
 dobj mail Mary iobj mail Mary the contract

The relation between a predicate and the second non clausal