

This excerpt from

Foundations of Statistical Natural Language Processing.  
Christopher D. Manning and Hinrich Schütze.  
© 1999 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact [cognetadmin@cognet.mit.edu](mailto:cognetadmin@cognet.mit.edu).

# 4

## *Corpus-Based Work*

THIS CHAPTER begins with some brief advice on getting set up to do corpus-based work. The main requirements for Statistical NLP work are computers, corpora, and software. Many of the details of computers and corpora are subject to rapid change, and so it does not make sense to dwell on these. Moreover, in many cases, one will have to make do with the computers and corpora at one's local establishment, even if they are not in all respects ideal. Regarding software, this book does not attempt to teach programming skills as it goes, but assumes that a reader interested in implementing any of the algorithms described herein can already program in *some* programming language. Nevertheless, we provide in this section a few pointers to languages and tools that may be generally useful.

After that the chapter covers a number of interesting issues concerning the formats and problems one encounters when dealing with 'raw data' - plain text in some electronic form. A very important, if often neglected, issue is the low-level processing which is done to the text before the real work of the research project begins. As we will see, there are a number of difficult issues in determining what is a word and what is a sentence. In practice these decisions are generally made by imperfect heuristic methods, and it is thus important to remember that the inaccuracies of these methods affect all subsequent results.

Finally the chapter turns to marked up data, where some process - often a human being - has added explicit markup to the text to indicate something of the structure and semantics of the document. This is often helpful, but raises its own questions about the kind and content of the markup used. We introduce the rudiments of SGML markup (and thus

also XML) and then turn to substantive issues such as the choice of tag sets used in corpora marked up for part of speech.

## 4.1 Getting Set Up

### 4.1.1 Computers

Text corpora are usually big. It takes quite a lot of computational resources to deal with large amounts of text. In the early days of computing, this was the major limitation on the use of corpora. For example in the earliest years of work on constructing the Brown corpus (the 1960s), just sorting all the words in the corpus to produce a word list would take 17 hours of (dedicated) processing time. This was because the computer (an IBM 7070) had the equivalent of only about 40 kilobytes of memory, and so the sort algorithm had to store the data being sorted on tape drives. Today one can sort this amount of data within minutes on even a modest computer.

As well as needing plenty of space to store corpora, Statistical NLP methods often consist of a step of collecting a large number of counts from corpora, which one would like to access speedily. This means that one wants a computer with lots of hard disk space, and lots of memory. In a rapidly changing world, it does not make much sense to be more precise than this about the hardware one needs. Fortunately, all the change is in a good direction, and often all that one will need is a decent personal computer with its RAM cheaply expanded (whereas even a few years ago, a substantial sum of money was needed to get a suitably fast computer with sufficient memory and hard disk space).

### 4.1.2 Corpora

A selection of some of the main organizations that distribute text corpora for linguistic purposes are shown in table 4.1. Most of these organizations charge moderate sums of money for corpora.<sup>1</sup> If your budget does not extend to this, there are now numerous sources of free text, ranging from email and web pages, to the many books and (maga)zines

---

1. Prices vary enormously, but are normally in the range of US\$100-2000 per CD for academic and nonprofit organizations, and reflect the considerable cost of collecting and processing material.

Linguistic Data Consortium (LDC)	<a href="http://www ldc.upenn.edu">http://www ldc.upenn.edu</a>
European Language Resources Association (ELRA)	<a href="http://www.icp.grenet.fr/ELRA/">http://www.icp.grenet.fr/ELRA/</a>
International Computer Archive of Modern English (ICAME)	<a href="http://nora.hd.uib.no/icame.html">http://nora.hd.uib.no/icame.html</a>
Oxford Text Archive (OTA)	<a href="http://ota.ahds.ac.uk/">http://ota.ahds.ac.uk/</a>
Child Language Data Exchange System (CHILDES)	<a href="http://childes.psy.cmu.edu/">http://childes.psy.cmu.edu/</a>

**Table 4.1** Major suppliers of electronic corpora with contact URLs.

that are available free on the web. Such free sources will not bring you linguistically-marked-up corpora, but often there are tools that can do the task of adding markup automatically reasonably well, and at any rate, working out how to deal with raw text brings its own challenges. Further resources for online text can be found on the website.

When working with a corpus, we have to be careful about the validity of estimates or other results of statistical analysis that we produce. A corpus is a special collection of textual material collected according to a certain set of criteria. For example, the Brown corpus was designed as a representative sample of written American English as used in 1961 (Francis and Kučera 1982: 5–6). Some of the criteria employed in its construction were to include particular texts in amounts proportional to actual publication and to exclude verse because “it presents special linguistic problems” (p. 5).

As a result, estimates obtained from the Brown corpus do not necessarily hold for British English or spoken American English. For example, the estimates of the entropy of English in section 2.2.7 depend heavily on the corpus that is used for estimation. One would expect the entropy of poetry to be higher than that of other written text since poetry can flaunt semantic expectations and even grammar. So the entropy of the Brown corpus will not help much in assessing the entropy of poetry. A more mundane example is text categorization (see chapter 16) where the performance of a system can deteriorate significantly over time because a sample drawn for training at one point can lose its representativeness after a year or two.

REPRESENTATIVE  
SAMPLE

The general issue is whether the corpus is a *representative sample* of the population of interest. A sample is representative if what we find for the sample also holds for the general population. We will not discuss methods for determining representativeness here since this issue is dealt with at length in the corpus linguistics literature. We also refer

## BALANCED CORPUS

the reader to this literature for creating *balanced corpora*, which are put together so as to give each subtype of text a share of the corpus that is proportional to some predetermined criterion of importance. In Statistical NLP, one commonly receives as a corpus a certain amount of data from a certain domain of interest, without having any say in how it is constructed. In such cases, having more training text is normally more useful than any concerns of balance, and one should simply use all the text that is available.

In summary, there is no easy way of determining whether a corpus is representative, but it is an important issue to keep in mind when doing Statistical NLP work. The minimal questions we should attempt to answer when we select a corpus or report results are what type of text the corpus is representative of and whether the results obtained will transfer to the domain of interest.

▼ The effect of corpus variability on the accuracy of part-of-speech tagging is discussed in section 10.3.2.

### 4.1.3 Software

There are many programs available for looking at text corpora and analyzing the data that you see. In general, however, we assume that readers will be writing their own software, and so all the software that is really needed is a plain text editor, and a compiler or interpreter for a language of choice. However, certain other tools, such as ones for searching through text corpora can often be of use. We briefly describe some such tools later.

#### Text editors

You will want a plain text editor that shows fairly literally what is actually in the file. Fairly standard and cheap choices are Emacs for Unix (or Windows), TextPad for Windows, and BBEdit for Macintosh.

#### Regular expressions

In many places and in many programs, editors, etc., one wishes to find certain patterns in text, that are often more complex than a simple match against a sequence of characters. The most general widespread notation for such matches are *regular expressions* which can describe patterns

## REGULAR EXPRESSIONS

REGULAR LANGUAGE

that are a *regular language*, the kind that can be recognized by a finite state machine. If you are not already familiar with regular expressions, you will want to become familiar with them. Regular expressions can be used in many plain text editors (Emacs, TextPad, Nisus, BBEdit, ...), with many tools (such as `grep` and `sed`), and as built-ins or libraries in many programming languages (such as Perl, C, ...). Introductions to regular expressions can be found in (Hopcroft and Ullman 1979; Sipser 1996; Friedl 1997).

### Programming languages

Most Statistical NLP work is currently done in C/C++. The need to deal with large amounts of data collection and processing from large texts means that the efficiency gains of coding in a language like C/C++ are generally worth it. But for a lot of the ancillary processing of text, there are many other languages which may be more economical with human labor. Many people use Perl for general text preparation and reformatting. Its integration of regular expressions into the language syntax is particularly powerful. In general, interpreted languages are faster for these kinds of tasks than writing everything in C. Old timers might still use `awk` rather than Perl – even though what you can do with it is rather more limited. Another choice, better liked by programming purists is Python, but using regular expressions in Python just is not as easy as Perl. One of the authors still makes considerable use of Prolog. The built-in database facilities and easy handling of complicated data structures makes Prolog excel for some tasks, but again, it lacks the easy access to regular expressions available in Perl. There are other languages such as SNOBOL/SPITBOL or Icon developed for text computing, and which are liked by some in the humanities computing world, but their use does not seem to have permeated into the Statistical NLP community. In the last few years there has been increasing uptake of Java. While not as fast as C, Java has many other appealing features, such as being object-oriented, providing automatic memory management, and having many useful libraries.

### Programming techniques

This section is not meant as a substitute for a general knowledge of computer algorithms, but we briefly mention a couple of useful tips.

**Coding words.** Normally Statistical NLP systems deal with a large number of words, and programming languages like C++ provide only quite limited facilities for dealing with words. A method that is commonly used in Statistical NLP and Information Retrieval is to map words to numbers on input (and only back to words when needed for output). This gives a lot of advantages because things like equality can be checked more easily and quickly on numbers. It also maps all tokens of a word to its type, which has a single number. There are various ways to do this. One good way is to maintain a large hash table (a hash function maps a set of objects into a specified range of integers, for example,  $[0, \dots, 127]$ ). A hash table allows one to see efficiently whether a word has been seen before, and if so return its number, or else add it and assign a new number. The numbers used might be indices into an array of words (especially effective if one limits the application to 65,000 or fewer words, so they can be stored as 16 bit numbers) or they might just be the address of the canonical form of the string as stored in the hashtable. This is especially convenient on output, as then no conversion back to a word has to be done: the string can just be printed.

There are other useful data structures such as various kinds of trees. See a book on algorithms such as (Cormen et al. 1990) or (Frakes and Baeza-Yates 1992).

**Collecting count data.** For a lot of Statistical NLP work, there is a first step of collecting counts of various observations, as a basis for estimating probabilities. The seemingly obvious way to do that is to build a big data structure (arrays or whatever) in which one counts each event of interest. But this can often work badly in practice since this model requires a huge memory address space which is being roughly randomly accessed. Unless your computer has enough memory for all those tables, the program will end up swapping a lot and will run very slowly. Often a better approach is for the data collecting program to simply emit a token representing each observation, and then for a follow on program to sort and then count these tokens. Indeed, these latter steps can often be done by existing system utilities (such as `sort` and `uniq` on Unix systems). Among other places, such a strategy is very successfully used in the CMU-Cambridge Statistical Language Modeling toolkit which can be obtained from the web (see website).

## 4.2 Looking at Text

MARKUP Text will usually come in either a raw format, or marked up in some way. *Markup* is a term that is used for putting codes of some sort into a computer file, that are not actually part of the text in the file, but explain something of the structure or formatting of that text. Nearly all computer systems for dealing with text use mark-up of some sort. Commercial word processing software uses markup, but hides it from the user by employing WYSIWYG (What You See Is What You Get) display. Normally, when dealing with corpora in Statistical NLP, we will want explicit markup that we can see. This is part of why the first tool in a corpus linguist's toolbox is a plain text editor.

There are a number of features of text in human languages that can make them difficult to process automatically, even at a low level. Here we discuss some of the basic problems that one should be aware of. The discussion is dominated by, but not exclusively concerned with, the most fundamental problems in *English* text.

### 4.2.1 Low-level formatting issues

#### Junk formatting/content

Depending on the source of the corpus, there may be various formatting and content that one cannot deal with, and is just junk that needs to be filtered out. This may include: document headers and separators, type-setter codes, tables and diagrams, garbled data in the computer file, etc. OCR If the data comes from *OCR* (Optical Character Recognition), the OCR process may have introduced problems such as headers, footers and floating material (tables, figures, and footnotes) breaking up the paragraphs of the text. There will also usually be OCR errors where words have been misrecognized. If your program is meant to deal with only connected English text, then other kinds of content such as tables and pictures need to be regarded as junk. Often one needs a filter to remove junk content before any further processing begins.

#### Uppercase and lowercase

The original Brown corpus was all capitals (a \* before a letter was used to indicate a capital letter in the original source text). All uppercase text is



PROPER NAMES

rarely seen these days, but even with modern texts, there are questions of how to treat capitalization. In particular, if we have two tokens that are identical except that one has certain letters in uppercase, should we treat them as the same? For many purposes we would like to treat *the*, *The*, and *THE* as the same, for example if we just want to do a study of the usage of definite articles, or noun phrase structure. This is easily done by converting all words to upper- or lowercase, but the problem is that at the same time we would normally like to keep the two types of *Brown* in *Richard Brown* and *brown paint* distinct. In many circumstances it is easy to distinguish *proper names* and hence to keep this distinction, but sometimes it is not. A simple heuristic is to change to lowercase letters capital letters at the start of a sentence (where English regularly capitalizes all words) and in things like headings and titles when there is a series of words that are all in capitals, while other words with capital letters are assumed to be names and their uppercase letters are preserved. This heuristic works quite well, but naturally, there are problems. The first problem is that one has to be able to correctly identify the ends of sentences, which is not always easy, as we discuss later. In certain genres (such as *Winnie the Pooh*), words may be capitalized just to stress that they are making a Very Important Point, without them indicating a proper name. At any rate, the heuristic will wrongly lowercase names that appear sentence initially or in all uppercase sequences. Often this source of error can be tolerated (because regular words are usually more common than proper names), but sometimes this would badly bias estimates. One can attempt to do better by keeping lists of proper names (perhaps with further information on whether they name a person, place, or company), but in general there is not an easy solution to the problem of accurate proper name detection.

#### 4.2.2 Tokenization: What is a word?

TOKENS  
WORD  
TOKENIZATION

Normally, an early step of processing is to divide the input text into units called *tokens* where each is either a *word* or something else like a number or a punctuation mark. This process is referred to as *tokenization*. The treatment of punctuation varies. While normally people want to keep sentence boundaries (see section 4.2.4 below), often sentence-internal punctuation has just been stripped out. This is probably unwise. Recent work has emphasized the information contained in all punctuation. No matter how imperfect a representation, punctuation marks like commas and

dashes give some clues about the macro structure of the text and what is likely to modify what.

GRAPHIC WORD

The question of what counts as a word is a vexed one in linguistics, and often linguists end up suggesting that there are words at various levels, such as phonological words versus syntactic words, which need not all be the same. What is a humble computational linguist meant to do? Kučera and Francis (1967) suggested the practical notion of a *graphic word* which they define as “a string of contiguous alphanumeric characters with space on either side; may include hyphens and apostrophes, but no other punctuation marks.” But, unfortunately, life is not that simple, even if one is just looking for a practical, workable definition. Kučera and Francis seem in practice to use intuition, since they regard as words numbers and monetary amounts like \$22.50 which do not strictly seem to obey the definition above. And things get considerably worse. Especially if using online material such as newsgroups and web pages for data, but even if sticking to newswires, one finds all sorts of oddities that should presumably be counted as words, such as references to *MicroSoft* or the web company *C|net*, or the various forms of smilies made out of punctuation marks, such as :-). Even putting aside such creatures, working out word tokens is a quite difficult affair. The main clue used in English is the occurrence of *whitespace* – a space or tab or the beginning of a new line between words – but even this signal is not necessarily reliable. What are the main problems?

WHITESPACE

### Periods

Words are not always surrounded by white space. Often punctuation marks attach to words, such as commas, semicolons, and periods (full stops). It at first seems easy to remove punctuation marks from word tokens, but this is problematic for the case of periods. While most periods are end of sentence punctuation marks, others mark an abbreviation such as in *etc.* or *Calif.* These abbreviation periods presumably should remain as part of the word, and in some cases keeping them might be important so that we can distinguish *Wash.*, an abbreviation for the state of Washington, from the capitalized form of the verb *wash*. Note especially that when an abbreviation like *etc.* appears at the end of the sentence, then only one period occurs, but it serves both functions of the period, simultaneously! An example occurred with *Calif.* earlier in this paragraph. Within morphology, this phenomenon is referred to as *haplogy*.

HAPLOGY

The issue of working out which punctuation marks do indicate the end of a sentence is discussed further in section 4.2.4.

### Single apostrophes

It is a difficult question to know how to regard English contractions such as *I'll* or *isn't*. These count as one *graphic word* according to the definition above, but many people have a strong intuition that we really have two words here as these are contractions for *I will* and *is not*. Thus some processors (and some corpora, such as the Penn Treebank) split such contractions into two words, while others do not. Note the impact that not splitting them has. The traditional first syntax rule:

$$S \rightarrow NP VP$$

stops being obviously true of sentences involving contractions such as *I'm right*. On the other hand, if one does split, there are then funny words like *'s* and *n't* in your data.

Phrases such as *the dog's* and *the child's*, when not abbreviations for *the dog is* or *the dog has*, are commonly seen as containing *dog's* as the genitive or possessive case of *dog*. But as we mentioned in section 3.1.1, this is not actually correct for English where *'s* is a *clitic* which can attach to other elements in a noun phrase, such as in *The house I rented yesterday's garden is really big*. Thus it is again unclear whether to regard *dog's* as one word or two, and again the Penn Treebank opts for the latter. Orthographic-word-final single quotations are an especially tricky case. Normally they represent the end of a quotation – and so should not be part of a word, but when following an *s*, they may represent an (unpronounced) indicator of a plural possessive, as in *the boys' toys* – and then should be treated as part of the word, if other possessives are being so treated. There is no easy way for a tokenizer to determine which function is intended in many such cases.

### Hyphenation: Different forms representing the same word

Perhaps one of the most difficult areas is dealing with hyphens in the input. Do sequences of letters with a hyphen in between count as one word or two? Again, the intuitive answer seems to be sometimes one, sometimes two. This reflects the many sources of hyphens in texts.

One source is typographical. Words have traditionally been broken and hyphens inserted to improve justification of text. These line-breaking hyphens may be present in data if it comes from what was actually typeset. It would seem to be an easy problem to just look for hyphens at the end of a line, remove them and join the part words at the end of one line and the beginning of the next. But again, there is the problem of haplology. If there is a hyphen from some other source, then after that hyphen is regarded as a legitimate place to break the text, and only one hyphen appears not two. So it is not always correct to delete hyphens at the end of a line, and it is difficult in general to detect which hyphens were line-breaking hyphens and which were not.

Even if such line-breaking hyphens are not present (and they usually are not in truly electronic texts), difficult problems remain. Some things with hyphens are clearly best treated as a single word, such as *e-mail* or *co-operate* or *A-1-plus* (as in *A-1-plus commercial paper*, a financial rating). Other cases are much more arguable, although we usually want to regard them as a single word, for example, *non-lawyer*, *pro-Arab*, and *so-called*. The hyphens here might be termed lexical hyphens. They are commonly inserted before or after small word formatives, sometimes for the purpose of splitting up vowel sequences.

The third class of hyphens is ones inserted to help indicate the correct grouping of words. A common copy-editing practice is to hyphenate compound pre-modifiers, as in the example earlier in this sentence or in examples like these:

- (4.1)
- a. the once-quiet study of superconductivity
  - b. a tough regime of business-conduct rules
  - c. the aluminum-export ban
  - d. a text-based medium

And hyphens occur in other places, where a phrase is seen as in some sense quotative or as expressing a quantity or rate:

- (4.2)
- a. the idea of a child-as-required-yuppie-possession must be motivating them
  - b. a final “take-it-or-leave-it” offer
  - c. the 90-cent-an-hour raise

## d. the 26-year-old

In these cases, we would probably want to treat the things joined by hyphens as separate words. In many corpora this type of hyphenation is very common, and it would greatly increase the size of the word vocabulary (mainly with items outside a dictionary) and obscure the syntactic structure of the text if such things were not split apart into separate words.<sup>2</sup>

A particular problem in this area is that the use of hyphens in many such cases is extremely inconsistent. Some texts and authorities use *cooperate*, while others use *co-operate*. As another example, in the Dow Jones newswire, one can find all of *database*, *data-base* and *data base* (the first and third are commonest, with the former appearing to dominate in software contexts, and the third in discussions of company assets, but without there being any clear semantic distinction in usage). Closer to home, look back at the beginning of this section. When we initially drafted this chapter, we (quite accidentally) used all of *markup*, *mark-up* and *mark(ed) up*. A careful copy editor would catch this and demand consistency, but a lot of the text we use has never been past a careful copy editor, and at any rate, we will commonly use texts from different sources which often adopt different conventions in just such matters. Note that this means that we will often have multiple forms, perhaps some treated as one word and others as two, for what is best thought of as a single *lexeme* (a single dictionary entry with a single meaning).

LEXEME

Finally, while British typographic conventions put spaces between dashes and surrounding words, American typographic conventions normally have a long dash butting straight up against the words—like this. While sometimes this dash will be rendered as a special character or as multiple dashes in a computer file, the limitations of traditional computer character sets means that it can sometimes be rendered just as a hyphen, which just further compounds the difficulties noted above.

### The same form representing multiple ‘words’

In the main we have been collapsing distinctions and suggesting that one may wish to regard variant sequences of characters as really the

2. One possibility is to split things apart, but to add markup, as discussed later in this chapter, which records that the original was hyphenated. In this way no information is lost.

## HOMOGRAPHS

same word. It is important to also observe the opposite problem, where one might wish to treat the identical sequence of characters as different words. This happens with *homographs*, where two lexemes have overlapping forms, such as *saw* as a noun for a tool, or as the past tense of the verb *see*. In such cases we might wish to assign occurrences of *saw* to two different lexemes.

▼ Methods of doing this automatically are discussed in chapter 7.

**Word segmentation in other languages**

## WORD SEGMENTATION

Many languages do not put spaces in between words at all, and so the basic word division algorithm of breaking on whitespace is of no use at all. Such languages include the major East-Asian languages/scripts such as Chinese, Japanese, Korean, and Thai. Ancient Greek was also written by Ancient Greeks without word spaces. Spaces were introduced (together with accent marks, etc.) by those who came afterwards. In such languages, *word segmentation* is a much more major and challenging task.

While maintaining most word spaces, in German compound nouns are written as a single word, for example *Lebensversicherungsgesellschafts-angestellter* ‘life insurance company employee.’ In many ways this makes linguistic sense, as compounds *are* a single word, at least phonologically. But for processing purposes one may wish to divide such a compound, or at least to be aware of the internal structure of the word, and this becomes a limited word segmentation task. While not the rule, joining of compounds sometimes also happens in English, especially when they are common and have a specialized meaning. We noted above that one finds both *data base* and *database*. As another example, while *hard disk* is more common, one sometimes finds *harddisk* in the computer press.

**Whitespace not indicating a word break**

Until now, the problems we have dealt with have mainly involved splitting apart sequences of characters where the word divisions are not shown by whitespace. But the opposite problem of wanting to lump things together also occurs. Here, things are separated by whitespace but we may wish to regard them as a single word. One possible case is the reverse of the German compound problem. If one decides to treat *database* as one word, one may wish to treat it as one word even when it is written as *data base*. More common cases are things such as phone numbers, where we

may wish to regard *9365 1873* as a single ‘word,’ or in the cases of multi-part names such as *New York* or *San Francisco*. An especially difficult case is when this problem interacts with hyphenation as in a phrase like this one:

- (4.3) the New York-New Haven railroad

Here the hyphen does not express grouping of just the immediately adjacent graphic words – treating *York-New* as a semantic unit would be a big mistake.

Other cases are of more linguistic interest. For many purposes, one would want to regard phrasal verbs (*make up*, *work out*) as a single lexeme (section 3.1.4), but this case is especially tricky since in many cases the particle is separable from the verb (*I couldn’t **work** the answer **out***), and so in general identification of possible phrasal verbs will have to be left to subsequent processing. One might also want to treat as a single lexeme certain other fixed phrases, such as *in spite of*, *in order to*, and *because of*, but typically a tokenizer will regard them as separate words. A partial implementation of this approach occurs in the LOB corpus where certain pairs of words such as *because of* are tagged with a single part of speech, here preposition, by means of using so-called *ditto tags*.

DITTO TAGS

#### Variant coding of information of a certain semantic type

Many readers may have felt that the example of a phone number in the previous section was not very recognizable or convincing because *their* phone numbers are written as 812-4374, or whatever. However, even if one is not dealing with multilingual text, any application dealing with text from different countries or written according to different stylistic conventions has to be prepared to deal with typographical differences. In particular, some items such as phone numbers are clearly of one semantic sort, but can appear in many formats. A selection of formats for phone numbers with their countries, all culled from advertisements in one issue of the magazine *The Economist*, is shown in table 4.2. Phone numbers variously use spaces, periods, hyphens, brackets, and even slashes to group digits in various ways, often not consistently even within one country. Additionally, phone numbers may include international or national long distance codes, or attempt to show both (as in the first three UK entries in the table), or just show a local number, and there may or may not be explicit indication of this via other marks such as brackets and plus

Phone number	Country	Phone number	Country
0171 378 0647	UK	+45 43 48 60 60	Denmark
(44.171) 830 1007	UK	95-51-279648	Pakistan
+44 (0) 1225 753678	UK	+411/284 3797	Switzerland
01256 468551	UK	(94-1) 866854	Sri Lanka
(202) 522-2230	USA	+49 69 136-2 98 05	Germany
1-925-225-3000	USA	33 1 34 43 32 26	France
212. 995.5402	USA	++31-20-5200161	The Netherlands

**Table 4.2** Different formats for telephone numbers appearing in an issue of *The Economist*.

INFORMATION  
EXTRACTION

signs. Trying to deal with myriad formats like this is a standard problem in *information extraction*. It has most commonly been dealt with by building carefully handcrafted regular expressions to match formats, but given the brittleness of such an approach, there is considerable interest in automatic means for learning the formatting of semantic types.

▼ We do not cover information extraction extensively in this book, but there is a little further discussion in section 10.6.2.

### Speech corpora

Our discussion has concentrated on written text, but the transcripts of speech corpora provide their own additional challenges. Speech corpora normally have more contractions, various sorts of more phonetic representations, show pronunciation variants, contain many sentence fragments, and include fillers like *er* and *um*. Example (4.4) - from the Switchboard corpus available from the LDC - shows a typical extract from a speech transcript:

- (4.4) Also I [cough] not convinced that the, at least the kind of people that I work with, I'm not convinced that that's really, uh, doing much for the progr-, for the, uh, drug problem.

### 4.2.3 Morphology

Another question is whether one wants to keep word forms like *sit*, *sits* and *sat* separate or to collapse them. The issues here are similar to those in the discussion of capitalization, but have traditionally been regarded



STEMMING  
 LEMMATIZATION  
 LEMMA

as more linguistically interesting. At first, grouping such forms together and working in terms of lexemes feels as if it is the right thing to do. Doing this is usually referred to in the literature as *stemming* in reference to a process that strips off affixes and leaves you with a stem. Alternatively, the process may be referred to as *lemmatization* where one is attempting to find the *lemma* or *lexeme* of which one is looking at an inflected form. These latter terms imply disambiguation at the level of lexemes, such as whether a use of *lying* represents the verb *lie-lay* 'to prostrate oneself' or *lie-lied* 'to fib.'

Extensive empirical research within the Information Retrieval (IR) community has shown that doing stemming does not help the performance of classic IR systems when performance is measured as an average over queries (Salton 1989; Hull 1996). There are always some queries for which stemming helps a lot. But there are others where performance goes down. This is a somewhat surprising result, especially from the viewpoint of linguistic intuition, and so it is important to understand why that is. There are three main reasons for this.

One is that while grouping the various forms of a stem seems a good thing to do, it often costs you a lot of information. For instance, while *operating* can be used in a periphrastic tense form as in *Bill is operating a tractor* (section 3.1.3), it is usually used in noun- and adjective-like uses such as *operating systems* or *operating costs*. It is not hard to see why a search for *operating systems* will perform better if it is done on inflected words than if one instead searches for all paragraphs that contain *operat-* and *system*. Or to consider another example, if someone enters *business* and the stemmer then causes retrieval of documents with *busy* in them, the results are unlikely to be beneficial.

Secondly, morphological analysis splits one token into several. However, often it is worthwhile to group closely related information into chunks, notwithstanding the blowout in the vocabulary that this causes. Indeed, in various Statistical NLP domains, people have been able to improve system performance by regarding frequent multiword units as a single distinctive token. Often inflected words are a useful and effective chunk size.

Thirdly, most information retrieval studies have been done on English - although recently there has been increasing multilingual work. English has very little morphology, and so the need for dealing intelligently with morphology is not acute. Many other languages have far richer systems of inflection and derivation, and then there is a pressing need for mor-

FULL-FORM LEXICON

phological analysis. A *full-form lexicon* for such languages, one that separately lists all inflected forms of all words, would simply be too large. For instance, Bantu languages (spoken in central and southern Africa) display rich verbal morphology. Here is a form from KiHaya (Tanzania). Note the prefixes for subject and object agreement, and tense:

- (4.5) akabimúha  
 a-ka-bi-mú-ha  
 1SG-PAST-3PL-3SG-give  
 'I gave them to him.'

For historical reasons, some Bantu language orthographies write many of these morphemes with whitespace in between them, but in the languages with 'conjunctive' orthographies, morphological analysis is badly needed. There is an extensive system of pronoun and tense markers appearing before the verb root, and quite a few other morphemes that can appear after the root, yielding a large system of combinatoric possibilities. Finnish is another language famous for millions of inflected forms for each verb.

One might be tempted to conclude from the paragraphs above that, in languages with rich morphology, one would gain by stripping inflectional morphology but not derivational morphology. But this hypothesis remains to be carefully tested in languages where there is sufficient inflectional morphology for the question to be interesting.

It is important to realize that this result from IR need not apply to any or all Statistical NLP applications. It need not even apply to all of IR. Morphological analysis might be much more useful in other applications. Stemming does not help in the non-interactive evaluation of IR systems, where a query is presented and processed without further input, and the results are evaluated in terms of the appropriateness of the set of documents returned. However, principled morphological analysis is valuable in IR in an interactive context, the context in which IR should really be evaluated. A computer does not care about weird stems like *busy* from *business*, but people do. They do not understand what is going on when *business* is stemmed to *busy* and a document with *busy* in it is returned.

It is also the case that nobody has systematically studied the possibility of letting people interactively influence the stemming. We believe that this could be very profitable, for cases like *saw* (where you want to stem for the sense 'see,' but not for the sense 'cutting implement'), or derivational cases where in some cases you want the stems (*arbitrary*

from *arbitrariness*), but in some you do not (*busy* from *business*). But the suggestion that human input may be needed does show the difficulties of doing automatic stemming in a knowledge-poor environment of the sort that has often been assumed in Statistical NLP work (for both ideological and practical reasons).

▼ Stemming and IR in general are further discussed in chapter 15.

#### 4.2.4 Sentences

##### What is a sentence?

The first answer to what is a sentence is “something ending with a ‘.’, ‘?’ or ‘!’.” We have already mentioned the problem that only some periods mark the end of a sentence: others are used to show an abbreviation, or for both these functions at once. Nevertheless, this basic heuristic gets one a long way: in general about 90% of periods are sentence boundary indicators (Riley 1989). There are a few other pitfalls to be aware of. Sometimes other punctuation marks split up what one might want to regard as a sentence. Often what is on one or the other or even both sides of the punctuation marks colon, semicolon, and dash (‘:’, ‘;’, and ‘—’) might best be thought of as a sentence by itself, as ‘:’ in this example:

- (4.6) The scene is written with a combination of unbridled passion and sure-handed control: In the exchanges of the three characters and the rise and fall of emotions, Mr. Weller has captured the heartbreaking inexorability of separation.

Related to this is the fact that sometimes sentences do not nicely follow in sequence, but seem to nest in awkward ways. While normally nested things are not seen as sentences by themselves, but clauses, this classification can be strained for cases such as the quoting of direct speech, where we get subsentences:

- (4.7) “You remind me,” she remarked, “of your mother.”

A second problem with such indirect speech is that it is standard typesetting practice (particularly in North America) to place quotation marks after sentence final punctuation. Therefore, the end of the sentence is not after the period in the example above, but after the close quotation mark that follows the period.

The above remarks suggest that the essence of a heuristic sentence division algorithm is roughly as in figure 4.1. In practice most systems

- Place putative sentence boundaries after all occurrences of . ? ! (and maybe ; : —)
- Move the boundary after following quotation marks, if any.
- Disqualify a period boundary in the following circumstances:
  - If it is preceded by a known abbreviation of a sort that does not normally occur word finally, but is commonly followed by a capitalized proper name, such as *Prof.* or *vs.*
  - If it is preceded by a known abbreviation and not followed by an uppercase word. This will deal correctly with most usages of abbreviations like *etc.* or *Jr.* which can occur sentence medially or finally.
- Disqualify a boundary with a ? or ! if:
  - It is followed by a lowercase letter (or a known name).
- Regard other putative sentence boundaries as sentence boundaries.

**Figure 4.1** Heuristic sentence boundary detection algorithm.

have used heuristic algorithms of this sort. With enough effort in their development, they can work very well, at least within the textual domain for which they were built. But any such solution suffers from the same problems of heuristic processes in other parts of the tokenization process. They require a lot of hand-coding and domain knowledge on the part of the person constructing the tokenizer, and tend to be brittle and domain-specific.

There has been increasing research recently on more principled methods of sentence boundary detection. Riley (1989) used statistical classification trees to determine sentence boundaries. The features for the classification trees include the case and length of the words preceding and following a period, and the a priori probability of different words to occur before and after a sentence boundary (the computation of which requires a large quantity of labeled training data). Palmer and Hearst (1994; 1997) avoid the need for acquiring such data by simply using the part of speech distribution of the preceding and following words, and using a neural network to predict sentence boundaries. This yields a

robust, largely language independent boundary detection algorithm with high performance (about 98–99% correct). Reynar and Ratnaparkhi (1997) and Mikheev (1998) develop Maximum Entropy approaches to the problem, the latter achieving an accuracy rate of 99.25% on sentence boundary prediction.<sup>3</sup>

▼ Sentence boundary detection can be viewed as a classification problem. We discuss classification, and methods such as classification trees and maximum entropy models in chapter 16.

### What are sentences like?

In linguistics classes, and when doing traditional computational linguistics exercises, sentences are generally short. This is at least in part because many of the parsing tools that have traditionally been used have a runtime exponential in the sentence length, and therefore become impractical for sentences over twelve or so words. It is therefore important to realize that typical sentences in many text genres are rather long. In newswire, the modal (most common) length is normally around 23 words. A chart of sentence lengths in a sample of newswire text is shown in table 4.3.

## 4.3 Marked-up Data

While much can be done from plain text corpora, by inducing the structure present in the text, people have often made use of corpora where some of the structure is shown, since it is then easier to learn more. This markup may be done by hand, automatically, or by a mixture of these two methods. Automatic means of learning structure are covered in the remainder of this book. Here we discuss the basics of markup. Some texts mark up just a little basic structure such as sentence and paragraph boundaries, while others mark up a lot, such as the full syntactic structure in corpora like the Penn Treebank and the Susanne corpus. However, the most common grammatical markup that one finds is a coding of words for part of speech, and so we devote particular attention to that.

---

3. *Accuracy* as a technical term is defined and discussed in section 8.1. However, the definition corresponds to one's intuitive understanding: it is the percent of the time that one is correctly classifying items.

Length	Number	Percent	Cum. %
1-5	1317	3.13	3.13
6-10	3215	7.64	10.77
11-15	5906	14.03	24.80
16-20	7206	17.12	41.92
21-25	7350	17.46	59.38
26-30	6281	14.92	74.30
31-35	4740	11.26	85.56
36-40	2826	6.71	92.26
41-45	1606	3.82	96.10
46-50	858	2.04	98.14
51-100	780	1.85	99.99
101+	6	0.01	100.00

**Table 4.3** Sentence lengths in newswire text. Column “Percent” shows the percentage in each range, column “Cum. %” shows the cumulative percentage below a certain length.

#### 4.3.1 Markup schemes

Various schemes have been used to mark up the structure of text. In the early days, these were developed on an ad hoc basis, as needs arose. One of the more important early examples was the COCOA format, which was used for including header information in texts (giving author, date, title, etc.). This information was enclosed within angle brackets with the first letter indicating the broad semantics of the field. Some other ad hoc systems of this sort are still in quite common use. The most common form of grammatical markup, which we discuss in great detail below, is indicating the part of speech of words by adding a part of speech tag to each word. These tags are commonly indicated by devices such as following each word by a slash or underline and then a short code naming the part of speech. The Penn Treebank uses a form of Lisp-like bracketing to mark up a tree structure over texts.

However, currently by far the most common and supported form of markup is to use SGML (the *Standard Generalized Markup Language*). SGML is a general language that lets one define a grammar for texts, in particular for the type of markup they contain. The now-ubiquitous HTML is an instance of an SGML encoding. The Text Encoding Initiative (TEI) was

STANDARD  
GENERALIZED MARKUP  
LANGUAGE

a major attempt to define SGML encoding schemes suitable for marking up various kinds of humanities text resources ranging from poems and novels to linguistic resources like dictionaries. Another acronym to be aware of is XML. XML defines a simplified subset of SGML that was particularly designed for web applications. However, the weight of commercial support behind XML and the fact that it avoids some of the rather arcane, and perhaps also archaic, complexities in the original SGML specification means that the XML subset is likely to be widely adopted for all other purposes as well.

DOCUMENT TYPE  
DEFINITION  
DTD

This book does not delve deeply into SGML. We will give just the rudimentary knowledge needed to get going. SGML specifies that each document type should have a *Document Type Definition* (DTD), which is a grammar for legal structures for the document. For example, it can state rules that a paragraph must consist of one or more sentences and nothing else. An SGML parser verifies that a document is in accordance with this DTD, but within Statistical NLP the DTD is normally ignored and people just process whatever text is found. An SGML document consists of one or more elements, which may be recursively nested. Elements normally begin with a begin tag and end with an end tag, and have document content in between. Tags are contained within angle brackets, and end tags begin with a forward slash character. As well as the tag name, the begin tag may contain additional attribute and value information. A couple of examples of SGML elements are shown below:

- (4.8) a. `<p><s>And then he left.</s>`  
`<s>He did not say another word.</s></p>`
- b. `<utt speak="Fred" date="10-Feb-1998">That is an ugly`  
`couch.</utt>`

The structure tagging shown in (4.8a), where the tag `s` is used for sentences and `p` for paragraphs, is particularly widespread. Example (4.8b) shows a tag with attributes and values. An element may also consist of just a single tag (without any matching end tag). In XML, such empty elements must be specially marked by ending the tag name with a forward slash character.

In general, when making use of SGML-encoded text in a casual way, one will wish to interpret some tags within angle brackets, and to simply ignore others. The other SGML syntax that one must be aware of is character and entity references. These begin with an ampersand and end with

a semicolon. Character references are a way of specifying characters not available in the standard ASCII character set (minus the reserved SGML markup characters) via their numeric code. Entity references have symbolic names which were defined in the DTD (or are one of a few predefined entities). Entity references may expand to any text, but are commonly used just to encode a special character via a symbolic name. A few examples of character and entity references are shown in (4.9). They might be rendered in a browser or when printed as shown in (4.10).

- (4.9) a. `&#x43;` is the less than symbol  
 b. `r&eacute;`; `sum&eacute;`;  
 c. This chapter was written on `&docdate;`.
- (4.10) a. `<` is the less than symbol  
 b. résumé  
 c. This chapter was written on January 21, 1998.

There is much more to know about SGML, and some references appear in the Further Reading below, but this is generally enough for what the XML community normally terms the ‘Desperate Perl Hacker’ to get by.

### 4.3.2 Grammatical tagging

A common first step of analysis is to perform automatic grammatical tagging for categories roughly akin to conventional parts of speech, but often considerably more detailed (for instance, distinguishing comparative and superlative forms of adjectives, or singular from plural nouns). This section examines the nature of tag sets. What tag sets have been used? Why do people use different ones? Which one should you choose? ▼ How tagging is done automatically is the subject of chapter 10.

#### Tag sets

##### BROWN TAG SET

Historically, the most influential tag sets have been the one used for tagging the American Brown corpus (the *Brown tag set*) and the series of tag sets developed at the University of Lancaster, and used for tagging the Lancaster-Oslo-Bergen corpus and more recently the British National



Sentence	CLAWS c5	Brown	Penn Treebank	ICE
she	PNP	PPS	PRP	PRON(pers,sing)
was	VBD	BEDZ	VBD	AUX(pass,past)
told	VVN	VBN	VBN	V(ditr,edp)
that	CJT	CS	IN	CONJUNC(subord)
the	ATO	AT	DT	ART(def)
journey	NN1	NN	NN	N(com,sing)
might	VM0	MD	MD	AUX(modal,past)
kill	VVI	VB	VB	V(montr,infin)
her	PNP	PPO	PRP	PRON(poss,sing)
.	PUN	.	.	PUNC(per)

**Figure 4.2** A sentence as tagged according to several different tag sets.

Tag set	Basic size	Total tags
Brown	87	179
Penn	45	
CLAWS1	132	
CLAWS2	166	
CLAWS c5	62	
London-Lund	197	

**Table 4.4** Sizes of various tag sets.

C5 TAG SET  
PENN TREEBANK TAG  
SET

Corpus (CLAWS1 through CLAWS5; CLAWS5 is also referred to as the *c5 tag set*). Recently, the *Penn Treebank tag set* has been the one most widely used in computational work. It is a simplified version of the Brown tag set. A brief summary of tag set sizes is shown in table 4.4. An example sentence shown tagged via several different tag sets is shown in figure 4.2. These tag sets are all for English. In general, tag sets incorporate morphological distinctions of a particular language, and so are not directly applicable to other languages (though often some of the design ideas can be transferred). Many tag sets for other languages have also been developed.

An attempt to align some tag sets, roughly organized by traditional parts of speech appears in tables 4.5 and 4.6, although we cannot guarantee that they are accurate in every detail. They are mostly alphabetical, but we have deviated from alphabetical order a little so as to group cat-

Category	Examples	Claws c5	Brown	Penn
Adjective	happy, bad	AJO	JJ	JJ
Adjective, ordinal number	sixth, 72nd, last	ORD	OD	JJ
Adjective, comparative	happier, worse	AJC	JJR	JJR
Adjective, superlative	happiest, worst	AJS	JJT	JJS
Adjective, superlative, semantically	chief, top	AJO	JJS	JJ
Adjective, cardinal number	3, fifteen	CRD	CD	CD
Adjective, cardinal number, one	one	PNI	CD	CD
Adverb	often, particularly	AVO	RB	RB
Adverb, negative	not, n't	XXO	*	RB
Adverb, comparative	faster	AVO	RBR	RBR
Adverb, superlative	fastest	AVO	RBT	RBS
Adverb, particle	up, off, out	AVP	RP	RP
Adverb, question	when, how, why	AVQ	WRB	WRB
Adverb, degree & question	how, however	AVQ	WQL	WRB
Adverb, degree	very, so, too	AVO	QL	RB
Adverb, degree, postposed	enough, indeed	AVO	QLP	RB
Adverb, nominal	here, there, now	AVO	RN	RB
Conjunction, coordination	and, or	CJC	CC	CC
Conjunction, subordinating	although, when	CJS	CS	IN
Conjunction, complementizer <i>that</i>	that	CJT	CS	IN
Determiner	this, each, another	DTO	DT	DT
Determiner, pronoun	any, some	DTO	DTI	DT
Determiner, pronoun, plural	these, those	DTO	DTS	DT
Determiner, prequalifier	quite	DTO	ABL	PDT
Determiner, prequantifier	all, half	DTO	ABN	PDT
Determiner, pronoun or double conj.	both	DTO	ABX	DT (CC)
Determiner, pronoun or double conj.	either, neither	DTO	DTX	DT (CC)
Determiner, article	the, a, an	ATO	AT	DT
Determiner, postdeterminer	many, same	DTO	AP	JJ
Determiner, possessive	their, your	DPS	PP\$	PRP\$
Determiner, possessive, second	mine, yours	DPS	PP\$	PRP
Determiner, question	which, whatever	DTQ	WDT	WDT
Determiner, possessive & question	whose	DTQ	WPS	WPS
Noun	aircraft, data	NN0	NN	NN
Noun, singular	woman, book	NN1	NN	NN
Noun, plural	women, books	NN2	NNS	NNS
Noun, proper, singular	London, Michael	NPO	NP	NNP
Noun, proper, plural	Australians, Methodists	NPO	NPS	NNPS
Noun, adverbial	tomorrow, home	NN0	NR	NN
Noun, adverbial, plural	Sundays, weekdays	NN2	NRS	NNS
Pronoun, nominal (indefinite)	none, everything, one	PNI	PN	NN
Pronoun, personal, subject	you, we	PNP	PPSS	PRP
Pronoun, personal, subject, 3SG	she, he, it	PNP	PPS	PRP
Pronoun, personal, object	you, them, me	PNP	PPO	PRP
Pronoun, reflexive	herself, myself	PNX	PPL	PRP
Pronoun, reflexive, plural	themselves, ourselves	PNX	PPLS	PRP
Pronoun, question, subject	who, whoever	PNQ	WPS	WP
Pronoun, question, object	who, whoever	PNQ	WPO	WP
Pronoun, existential there	there	EXO	EX	EX

**Table 4.5** Comparison of different tag sets: adjective, adverb, conjunction, determiner, noun, and pronoun tags.

Category	Examples	Claws c5	Brown	Penn
Verb, base present form (not infinitive)	take, live	VVB	VB	VBP
Verb, infinitive	take, live	VVI	VB	VB
Verb, past tense	took, lived	VVD	VBD	VBD
Verb, present participle	taking, living	VVG	VBG	VBG
Verb, past/passive participle	taken, lived	VVN	VBN	VBN
Verb, present 3SG -s form	takes, lives	VVZ	VBZ	VBZ
Verb, auxiliary <i>do</i> , base	do	VDB	DO	VBP
Verb, auxiliary <i>do</i> , infinitive	do	VDB	DO	VB
Verb, auxiliary <i>do</i> , past	did	VDD	DOD	VBD
Verb, auxiliary <i>do</i> , present part.	doing	VDG	VBG	VBG
Verb, auxiliary <i>do</i> , past part.	done	VDN	VBN	VBN
Verb, auxiliary <i>do</i> , present 3SG	does	VDZ	DOZ	VBZ
Verb, auxiliary <i>have</i> , base	have	VHB	HV	VBP
Verb, auxiliary <i>have</i> , infinitive	have	VHI	HV	VB
Verb, auxiliary <i>have</i> , past	had	VHD	HVD	VBD
Verb, auxiliary <i>have</i> , present part.	having	VHG	HVG	VBG
Verb, auxiliary <i>have</i> , past part.	had	VHN	HVN	VBN
Verb, auxiliary <i>have</i> , present 3SG	has	VHZ	HVZ	VBZ
Verb, auxiliary <i>be</i> , infinitive	be	VBI	BE	VB
Verb, auxiliary <i>be</i> , past	were	VBD	BED	VBD
Verb, auxiliary <i>be</i> , past, 3SG	was	VBD	BEDZ	VBD
Verb, auxiliary <i>be</i> , present part.	being	VBG	BEG	VBG
Verb, auxiliary <i>be</i> , past part.	been	VBN	BEN	VBN
Verb, auxiliary <i>be</i> , present, 3SG	is, 's	VBZ	BEZ	VBZ
Verb, auxiliary <i>be</i> , present, 1SG	am, 'm	VBB	BEM	VBP
Verb, auxiliary <i>be</i> , present	are, 're	VBB	BER	VBP
Verb, modal	can, could, 'll	VM0	MD	MD
Infinitive marker	to	TOO	TO	TO
Preposition, to	to	PRP	IN	TO
Preposition	for, above	PRP	IN	IN
Preposition, of	of	PRF	IN	IN
Possessive	's, '	POS	\$	POS
Interjection (or other isolate)	oh, yes, mmm	ITJ	UH	UH
Punctuation, sentence ender	. ! ?	PUN	.	.
Punctuation, semicolon	;	PUN	.	:
Punctuation, colon or ellipsis	: ...	PUN	:	:
Punctuation, comma	,	PUN	,	,
Punctuation, dash	-	PUN	-	-
Punctuation, dollar sign	\$	PUN	not	\$
Punctuation, left bracket	( [ {	PUL	(	(
Punctuation, right bracket	) ] }	PUR	)	)
Punctuation, quotation mark, left	' "	PUQ	not	"
Punctuation, quotation mark, right	' "	PUQ	not	"
Foreign words (not in English lexicon)		UNC	(FW-)	FW
Symbol	[fj] *		not	SYM
Symbol, alphabetical	A, B, c, d	ZZ0		
Symbol, list item	A A. First			LS

**Table 4.6** Comparison of different tag sets: Verb, preposition, punctuation and symbol tags. An entry of 'not' means an item was ignored in tagging, or was not separated off as a separate token.

egories that are sometimes collapsed. In this categorization, we use an elsewhere convention where the least marked category is used in all cases where a word cannot be placed within one of the more precise subclassifications. For instance, the plain Adjective category is used for adjectives that aren't comparatives, superlatives, numbers, etc. The complete Brown tag set was made larger by two decisions to augment the tag set. Normal tags could be followed by a hyphen and an attribute like TL (for a title word), or in the case of foreign words, the FW foreign word tag was followed by a hyphen and a part of speech assignment. Secondly, the Brown tag scheme makes use of 'combined tags' for graphic words that one might want to think of as multiple lexemes, such as *you'll*.<sup>4</sup> Normally such items were tagged with two tags joined with a plus sign, but for negation one just adds \* to a tag. So *isn't* is tagged BEZ\* and *she'll* is tagged PPS+MD. Additionally, possessive forms like *children's* are tagged with a tag ending in 'S'. Normally, these tags are transparently derived from a base non-possessive tag, for instance, NNS\$ in this case. These techniques of expanding the tag set are ignored in the comparison.

Even a cursory glance will show that the tag sets are very different. Part of this can be attributed to the overall size of the tag set. A larger tag set will obviously make more fine-grained distinctions. But this is not the only difference. The tag sets may choose to make distinctions in different areas. For example, the c5 tag set is larger overall than the Penn Treebank tag set, and it makes many more distinctions in some areas, but in other areas it has chosen to make many fewer. For instance, the Penn tag set distinguishes 9 punctuation tags, while c5 makes do with only 4. Presumably this indicates some difference of opinion on what is considered important. Tag sets also disagree more fundamentally in how to classify certain word classes. For example, while the Penn tag set simply regards subordinating conjunctions as prepositions (consonant with work in generative linguistics), the c5 tag set keeps them separate, and moreover implicitly groups them with other types of conjunctions. The notion of implicit grouping referred to here is that all the tag sets informally show relationships between certain sets of tags by having them begin with the same letter or pair of letters. This grouping is implicit in that although it is obvious to the human eye, they are formally just distinct symbolic

4. Compare the discussion above. This is also done in some other corpora, such as the London-Lund corpus, but the recent trend seems to have been towards dividing such graphic words into two for the purposes of tagging.

tags, and programs normally make no use of these families. However, in some other tag sets, such as the one for the International Corpus of English (Greenbaum 1993), an explicit system of high level tags with attributes for the expression of features has been adopted. There has also been some apparent development in people's ideas of what to encode. The early tag sets made very fine distinctions in a number of areas such as the treatment of certain sorts of qualifiers and determiners that were relevant to only a few words, albeit common ones. More recent tag sets have generally made fewer distinctions in such areas.

### The design of a tag set

What features should guide the design of a tag set? Standardly, a tag set encodes both the target feature of classification, telling the user the useful information about the grammatical class of a word, and the predictive features, encoding features that will be useful in predicting the behavior of other words in the context. These two tasks should overlap, but they are not necessarily identical.

#### PART OF SPEECH

The notion of part of speech is actually complex, since parts of speech can be motivated on various grounds, such as semantic (commonly called notional) grounds, syntactic distributional grounds, or morphological grounds. Often these notions of part of speech are in conflict. For the purposes of prediction, one would want to use the definition of part of speech that best predicts the behavior of nearby words, and this is presumably strictly distributional tags. But in practice people have often used tags that reflect notional or morphological criteria. For example one of the uses of English present participles ending in *-ing* is as a gerund where they behave as a noun. But in the Brown corpus they are quite regularly tagged with the VBG tag, which is perhaps better reserved for verbal uses of participles. This happens even within clear noun compounds such as this one:

- (4.11) Fulton/NP-TL County/NN-TL Purchasing/VBG Department/NN

Ideally, we would want to give distinctive tags to words that have distinctive distributions, so that we can use that information to help processing elsewhere. This would suggest that some of the tags in, for example, the Penn Treebank tag set are too coarse to be good predictors. For instance, the complementizer *that* has a very distinct distribution from regular prepositions, and degree adverbs and the negative *not* have

very different distributions from regular adverbs, but neither of these distinctions show up in the tag set. People have frequently made changes to add or remove distinctions according to their intuitions – for example, Charniak (1996) questions the decision of the Penn Treebank to tag auxiliaries with the same tags as other verbs, given that auxiliary verbs have a very distinctive distribution, and proceeds to retag them with an AUX tag. In general, the predictive value of making such changes in the set of distinctions in part of speech systems has not been very systematically evaluated. So long as the same tag set is used for prediction and classification, making such changes tends to be a two-edged sword: splitting tags to capture useful distinctions gives improved information for prediction, but makes the classification task harder.<sup>5</sup> For this reason, there is not necessarily a simple relationship between tag set size and the performance of automatic taggers.

#### 4.4 Further Reading

The Brown corpus (the Brown University Standard Corpus of Present-Day American English) consists of just over a million words of written American English from 1961. It was compiled and documented by W. Nelson Francis and Henry Kučera (Francis and Kučera 1964; Kučera and Francis 1967; Francis and Kučera 1982). The details on early processing of the Brown corpus are from an email from Henry Kučera (posted to the corpora mailing list by Katsuhide Sonoda on 26 Sep 1996). The LOB (Lancaster-Oslo-Bergen) corpus was built as a British-English replication of the Brown Corpus during the 1970s (Johansson et al. 1978; Garside et al. 1987).

Identifying proper names is a major issue in Information Extraction. See (Cardie 1997) for an introduction.

A carefully designed and experimentally tested set of tokenization rules is the set used for the Susanne corpus (Sampson 1995: 52–59).

Nunberg (1990) provides a linguistic perspective on the importance of *punctuation*. An introductory discussion of what counts as a word in linguistics can be found in (Crowley et al. 1995: 7–9). Lyons (1968: 194–206) provides a more thorough discussion. The examples in the section on hyphenation are mainly real examples from the Dow Jones newswire.

PUNCTUATION

5. This is unless one category groups two very separate distributional clusters, in which case splitting the category can actually sometimes make classification easier.

Others are from e-mail messages to the corpora list by Robert Amsler and Mitch Marcus, 1996, and are used with thanks.

There are many existing systems for morphological analysis available, and some are listed on the website. An effective method of doing stemming in a knowledge-poor way can be found in Kay and Röscheisen (1993). Sproat (1992) contains a good discussion of the problems morphology presents for NLP and is the source of our German compound example.

The COCOA (COunt and CONcordance on Atlas) format was used in corpora from ICAME and in related software such as LEXA (Hickey 1993).

SGML and XML are described in various books (Herwijnen 1994; McGrath 1997; St. Laurent 1998), and a lot of information, including some short readable introductions, is available on the web (see website).

The guidelines of the Text Encoding initiative (1994 P3 version) are published as McQueen and Burnard (1994), and include a very readable introduction to SGML in chapter 2. In general, though, rather than read the actual guidelines, one wants to look at tutorials such as Ide and Véronis (1995), or on the web, perhaps starting at the sites listed on the website. The full complexity of the TEI overwhelmed all but the most dedicated standard bearers. Recent developments include TEILite, which tries to pare the original standard down to a human-usable version, and the Corpus Encoding Standard, a TEI-conformant SGML instance especially designed for language engineering corpora.

Early work on CLAWS (Constituent-Likelihood Automatic Word-tagging System) and its tag set is described in (Garside et al. 1987). The more recent c5 tag set presented above is taken from (Garside 1995). The Brown tag set is described in (Francis and Kučera 1982) while the Penn tag set is described in (Marcus et al. 1993), and in more detail in (Santorini 1990).

This book is not an introduction to how corpora are used in linguistic studies (even though it contains a lot of methods and algorithms useful for such studies). However, recently there has been a flurry of new texts on *corpus linguistics* (McEnery and Wilson 1996; Stubbs 1996; Biber et al. 1998; Kennedy 1998; Barnbrook 1996). These books also contain much more discussion of corpus design issues such as sampling and balance than we have provided here. For an article specifically addressing the problem of designing a representative corpus, see (Biber 1993).

More details about different tag sets are collected in Appendix B of (Garside et al. 1987) and in the web pages of the AMALGAM project (see website). The AMALGAM website also has a description of the tokenizing

rules that they use, which can act as an example of a heuristic sentence divider and tokenizer. Grefenstette and Tapanainen (1994) provide another discussion of tokenization, showing the results of experiments employing simple knowledge-poor heuristics.

## 4.5 Exercises

### Exercise 4.1

[★★]

As discussed in the text, it seems that for most purposes, we'd want to treat some hyphenated things as words (for instance, *co-worker*, *Asian-American*), but not others (for instance, *ain't-it-great-to-be-a-Texan*, *child-as-required-yuppie-possession*). Find hyphenated forms in a corpus and suggest some basis for which forms we would want to treat as words and which we would not. What are the reasons for your decision? (Different choices may be appropriate for different needs.) Suggest some methods to identify hyphenated sequences that should be broken up – e.g., ones that only appear as non-final elements of compound nouns:

*[<sub>N</sub>[child-as-required-yuppie-possession] syndrome]*

### Exercise 4.2

[★★ For linguists]

Take some linguistic problem that you are interested in (non-constituent coordination, ellipsis, idioms, heavy NP shift, pied-piping, verb class alternations, etc.). Could one hope to find useful data pertaining to this problem in a general corpus? Why or why not? If you think it might be possible, is there a reasonable way to search for examples of the phenomenon in either a raw corpus or one that shows syntactic structures? If the answer to both these questions is yes, then look for examples in a corpus and report on anything interesting that you find.

### Exercise 4.3

[★★]

Develop a sentence boundary detection algorithm. Evaluate how successful it is. (In the construction of the *Wall Street Journal* section of the ACL-DCI CD-ROM (Church and Liberman 1991), a rather simplistic sentence boundary detection algorithm was used, and the results were not hand corrected, so many errors remain. If this corpus is available to you, you may want to compare your results with the sentence boundaries marked in the corpus. With luck, you should be able to write a system that performs considerably better!)



This excerpt from

Foundations of Statistical Natural Language Processing.  
Christopher D. Manning and Hinrich Schütze.  
© 1999 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact [cognetadmin@cognet.mit.edu](mailto:cognetadmin@cognet.mit.edu).