

## Chapter 4

# Word Frequency in Context: Alternative Architectures for Examining Related Words, Register Variation and Historical Change

Mark Davies

### Introduction

Since the advent of 'mega-corpora' that are 100 million words in size or larger, there have been challenges in terms of extracting large amounts of data economically. For example, several years ago it was sufficient to create a query engine that would perform a linear scan of the text, and such an architecture might return the results from a one million word corpus in one to two seconds. Using that same architecture, however, a similar query on a 100 million word corpus might take 100 to 200 seconds. As a result of these performance issues, in the last ten to 15 years a number of alternative architectures have been developed. These include the use of large numbers of indices that contain offset values for each word in the corpus and the use of large hash operations to find nearby words (e.g. SARAYXAIRA<sup>1</sup>) and the relational database architecture of the IMS Corpus Workbench.<sup>2</sup>

Over the past five years, we have employed a modified relational database architecture for a number of corpora that we have created. In contrast to the IMS Corpus Workbench approach, however, these corpora rely heavily on an 'n-gram' architecture, which will be one of the major topics of this paper. These corpora include the 100 million word *Corpus del Español*<sup>3</sup> and a new architecture and interface for the 100 million word *British National Corpus* (BNC)<sup>4</sup> both of which will be discussed below.

---

1 G. Burnage and D. Dunlop, 'Encoding the British National Corpus', in J. Aarts, P. de Haan and N. Oostdijk (eds), *English Language Corpora: Design, Analysis and Exploitation*, papers from the 13th International Conference on English Language Research on Computerized Corpora, Nijmegen 1992 (Amsterdam: Rodopi, 1993), pp. 79–95; L. Burnard, *Reference Guide for the British National Corpus: World Edition* (Oxford: Oxford University Computing Services, 2000).

2 O. Christ, *The IMS Corpus Workbench Technical Manual* (Institut für Maschinelle Sprachverarbeitung, Stuttgart, Germany: Universität Stuttgart, 1994).

3 <<http://www.corpusdelespanol.org>>

4 <<http://view.bnu.edu/>>

As with some other competing architectures, this relational database/n-gram approach allows queries like the following:

- the overall frequency of a given word or phrase in the corpus ('mysterious', 'blue skies')
- the frequency of words with a given substring (\*'able, \*heart\*', etc.)
- queries involving part of speech or lemma (e.g. 'utter' NNI, 'as' ADJ, 'barely realized' ADV VVD)

Unlike some other architectures, however, our approach is quite fast: any of the preceding queries on a 100 million word corpus would take less than one second.

In addition, as we will demonstrate, our approach allows a number of types of query that would be difficult or impossible to carry out directly in one step with competing architectures. These include – but certainly are not limited to – the following:

- comparison of frequency with related words, e.g. nouns occurring immediately after 'utter' but not after 'complete' or 'sheer', or adjectives within ten words of 'woman' but not 'man';
- one simple query to find the frequency of words in separate databases, such as user-defined, customized lists (clothing, emotions, technology terms, etc.) or synsets from WordNet;
- register variation, e.g. all words ending in '\*icity', or all verbs, or all three-word lexical bundles, which are more common in academic texts than in works of fiction, or in legal or medical texts;
- historical variation, e.g. all words, phrases or collocates of a given word or part of speech, which are more common in the 1900s than in the 1800s.

In the discussion that follows, we will first present the basic architecture (relational databases and n-grams) and provide concrete examples of some of the types of queries that this architecture allows. We will then discuss shortcomings of this architecture, and consider how these issues have been handled in some of the newer interfaces that we have created, such as the VIEW interface for the BNC.<sup>5</sup>

### A simple n-gram architecture

Let us first consider the 'first-generation' approach to relational databases and n-grams, which was used for the *Corpus del Español*<sup>6</sup> that we created in 2001–2002.<sup>7</sup>

- 5 <<http://view.byu.edu/>>
- 6 <<http://www.corpusdelespanol.org/>>
- 7 M. Davies, Advanced Research on Syntactic and Semantic Change with the Corpus del Español, in J. Kabatek, C.D. Fusch and W. Raible (eds), *Romance Corpus Linguistics II*:

Table 4.1 Example of 3-grams where *lem1* = 'break' and *word2* = 'the'

FREQ.	WORD1	WORD2	WORD#
106	breaking	the	law
98	break	the	law
56	broke	the	silence
53	break	the	news
46	broke	the	news
40	break	the	deadlock
24	broken	the	law
23	break	the	habit

Table 4.2 Example of 3-grams where *lem1* = 'break' and *word2* = 'the'

FREQ	WORD1	LEM1	POS1	W2	LEM2	POS2	W3	LEM3	POS3
106	breaking	break	VVG	the	AT0	law	law	law	NN1
98	Break	break	VVI	the	AT0	law	law	law	NN1
56	broke	break	VVD	the	AT0	silence	silence	silence	NN1
53	Break	break	VVI	the	AT0	news	news	news	NN1
46	broke	break	VVD	the	AT0	news	news	news	NN1
40	Break	break	VVI	the	AT0	deadlock	deadlock	deadlock	NN1
24	broken	break	VVN	the	AT0	law	law	law	NN1
23	Break	break	VVI	the	AT0	habit	habit	habit	NN1

and the subsequent BNC-based 'Phrases in English' database and interface that was based on the same architecture and which was created by William Fletcher in 2003.<sup>8</sup>

In this approach, one uses a program to create the n-grams of a given corpus, such as the Wordlist module of Wordsmith<sup>9</sup> or the KFN-gram program.<sup>10</sup> For example, with Wordsmith, one would simply create separate lists of all of the 1-

*Corpora and Diachronic Linguistics* (Tübingen: Gunter Naeff, 2005), pp. 203–14; M. Davies, 'The Advantage of Using Relational Databases for Large Corpora: Speed, Advanced Queries, and Unlimited Annotation', *International Journal of Corpus Linguistics*, 10 (2005): 301–28.

8 W. Fletcher, 'Exploring Words and Phrases from the British National Corpus', *Phrases in English* (2005) <<http://pic.usna.edu/>>, accessed June 2005.

9 M. Scott, *Wordsmith Tools, Version 4*.

10 <<http://minneapolis.com/KWICFinder/kNgramHelp.html>>

grams, 2-grams, 3-grams, etc., in the corpus, and then import these into a relational database. In the case of 3-grams for the BNC, for example, a small section of the 3-grams table would look like Table 4.1. Each unique three-word string in the corpus appears in the database, with its associated frequency. For example, in the BNC 'breaking the law' occurs 106 times, 'broke the law' occurs 56 times, and so on. It is also possible to create frequency tables that include POS (part of speech) and lemmatization information. In this case, the table might look like Table 4.2.

With these n-gram/frequency tables, it is a relatively simple process to use SQL queries to extract the desired data. For example, to extract the 100 most common singular nouns (NN1) in the BNC, the SQL query would be as follows:

```
(1)
SELECT TOP 100 * FROM [TABLE_NAME] WHERE
    POS1 = 'NN1'
ORDER BY FREQ DESC
```

ORDER BY FREQ DESC

To select the 100 most common three-word strings where the first word is a form of 'break', the second word is 'a' or 'the', and the third word is a noun, the SQL query would be as follows, and the results would be those seen in Table 4.2, above:

```
(2)
SELECT TOP 100 * FROM [TABLE_NAME] WHERE
    LEM1 = 'BREAK' AND
    WORD2 = 'THE' AND
    POS3 LIKE 'NN%'
ORDER BY FREQ DESC
```

Either of these two queries would take less than half a second to retrieve the 100 most frequent matching words or strings from the 100 million word corpus. This is the approach used in our *Corpus del Español* and in the 'Phrases in English' interface, and it represents an early approach to the use of n-grams.

#### Accounting for register or historical variation

There is a serious problem, however, associated with a strict n-gram architecture. Once the frequencies are calculated for each unique n-gram in the corpus, one then loses all contextual information for that n-gram – in other words, in which part of the corpus each n-gram occurs. For example, Table 4.2 above shows that

Table 4.3 Example of 3-grams where LEM1 = 'break' and WORD2 = 'the'

WORD1	LEM1	POS1	...	W3	LEM3	POS3	REG1	REG2	REG3
breaking	break	VVG	...	law	law	NN1	x1	y1	z1
break	break	VVI	...	law	law	NN1	x2	y2	z2
break	break	VVD	...	silence	silence	NN1	x3	y3	z3
break	break	VVI	...	news	news	NN1	x4	y4	z4
break	break	VVD	...	news	news	NN1	x5	y5	z5
break	break	VVI	...	deadlock	deadlock	NN1	x6	y6	z6
broken	break	VVN	...	law	law	NN1	x7	y7	z7
break	break	VVI	...	habit	habit	NN1	x8	y8	z8

Table 4.4 Hard N' (N duro) in the Corpus del Español, by century

WORD1	WORD2	x12	...	x17	x18	x19	x19SP	x19FIC	x19NF
MANO	DURA	0	...	0	10	22	6	9	7
LINEA	DURA	0	...	0	0	12	0	3	9
SER	DURO	0	...	1	9	10	7	1	2
CUELLO	DURO	1	...	0	0	10	9	1	0
MADERA	DURA	0	...	0	5	10	7	0	3
CARA	DURA	0	...	0	3	10	6	4	0

'breaking the law' occurs 106 times in the corpus, but at this point we have no idea how many of these are in the *SPOKEN* texts, or *FICTION* or *NEWSPAPERS*. Therefore, it would probably be impossible to find the most frequent words or phrases in a given register using this approach, or to compare the frequency of words or phrases in two competing (sets of) registers.

There is a way around the lack of context for each n-gram, however. One could calculate the overall n-gram frequency for a set of different registers (as in Table 4.2), and then create n-gram frequency tables for each register individually. One would then 'merge' the information from the 'register' tables into the overall frequency table, which would contain separate columns (for each n-gram) showing the frequency in each register. In other words, the resulting table might look like Table 4.3. This is in fact the approach taken in the construction of the 100 million word *Corpus del Español*. For each n-gram, there are columns that show the frequency of the string in each century from the 1200s to the 1900s (x12–x19, below). There are also separate columns containing the frequency of the n-gram in each of the three registers *SPOKEN*, *FICTION* and *NON-FICTION* in the 1900s. Table 4.4

shows a small section of bi-grams, containing a few of the n-grams that match the query `NOUN + lemma DURO 'hard N'`.

The advantage of using such an approach should be readily apparent. Since each n-gram has the associated frequency in each of the different historical periods and the different registers, this frequency information can be accessed directly as part of the query. For example, in the case of the *Corpus del Español*, we can find which nouns occur with *duro* 'hard' in the 1900s, but not in the 1800s. The SQL query would look something like the following (here simplified from how it would appear in the actual database):

```
(3)
SELECT TOP 100 * FROM [TABLE_NAME] WHERE
POS1 = 'NOUN' AND
LEM2 = 'DURO' AND
X19 <> 0 AND
X18 = 0
ORDER BY X19 DESC
```

This gives us results like 'línea dura' ('hard line'), 'disco duro' ('hard drive') and 'años duros' ('hard years'), etc.

In spite of the advantages of this approach, one problem is that it is quite costly to run the SQL UPDATE commands that copy the frequency information (for tens of millions of n-grams) from each of the separate tables (1200s, 1500s, 1900s–FIC, etc.) into the main n-gram tables. In addition, this approach may only be practical when there are a limited number of frequency columns, such as the 11 columns in the *Corpus del Español* (1200s–1900s, and three additional registers for the 1900s). In the case of the BNC, on the other hand, there are nearly 70 different registers, according to the categorization made by David Lee.<sup>11</sup>

In summary, the frequency information from each sub-corpus can be quite valuable, in terms of being able to compare different historical periods, different registers, and so on. However, because of the difficulty of creating such tables, they are not used in some other competing architectures and interfaces. As a result, it is only possible to look at word and phrase frequency across the entire corpus with these approaches.

11 <[http://opinion.nuchba.ac.jp/~davidlee/devotedlocorpora/home/BNC\\_WORLD\\_INDEX.ZIP](http://opinion.nuchba.ac.jp/~davidlee/devotedlocorpora/home/BNC_WORLD_INDEX.ZIP)>

### The issue of size

In addition to the problem of 'granularity' in terms of frequency in sub-registers, another problem with a strict n-gram approach has to do with the size of the tables. While there are only about 800,000 rows in the [1-grams] table of the BNC (i.e. 800,000+ unique types in the corpus), this increases to about 11 million unique [2-grams] and 40 million unique [3-grams], and it would move towards 90 to 95 million unique 7-grams for the 100 million words.

The problem with this approach, then, is that the n-gram tables become quite unmanageable in terms of size. Even with efficient clustered indexes on the tables, it may take ten to 15 seconds to return the results from a particularly difficult query, such as the most frequent n-grams for ['the NNI that']. The second problem is that once we add up all of the different n-gram tables (1-grams + 2-grams + 3-grams, etc.), we soon find that the total number of rows in these tables is larger than the total number of words in the corpus, especially if we include 4-grams to 7-grams.

As a result of the size issue, the approach taken in the construction of the 'Phrases in English' database and interface<sup>12</sup> is to include just those n-grams that occur three times or more. By eliminating from the tables all n-grams that occur just once or twice, the size of the tables is reduced dramatically – by 75 per cent in the case of the 3-grams, and even more for 4-grams to 7-grams. Therefore, there is a large performance gain by eliminating all n-grams that occur just once or twice. Unfortunately, with this approach, one also completely loses the ability to analyze these less common strings. If one is interested in only the highly-frequent n-grams, this is probably not a significant problem. But for detailed comparison of sub-corpora, or to see which phrases are entering into or leaving the language, it is problematic to exclude 75 per cent or more of all n-gram types.

As we have seen, there are two basic problems with a strict n-gram approach. First, in terms of sub-corpora, we either have to ignore the context in which each n-gram appears (register, historical period, etc.), or else we have to create tables that are quite difficult to construct, by merging the frequency for each sub-corpus into a column in the main n-gram table. Secondly, in terms of size and speed, we either create tables that eliminate n-grams that occur just once or twice (and thus lose 75 per cent or more of all unique n-grams), or we have a number of separate n-gram tables (2-grams, 5-grams, etc.) whose combined size in terms of rows is probably much larger than the total number of words in the corpus.

### An alternate architecture – all sequential n-grams

There is a 'second-generation' approach, however, that avoids both of the problems associated with a strict n-gram approach. In this architecture, we create one single

12 <<http://pie.usna.edu>>



Table 4.5 Sequential n-gram 'tokens'

ID	text	...	word3	pos3	word4	pos4	word5	pos5	...
50891887	EAA	...	a	ATO	small	AJO	group	NNI	...
50891888	EAA	...	small	AJO	group	NNI	of	PRF	...
50891889	EAA	...	group	NNI	of	PRF	people	NN0	...
50891890	EAA	...	of	PRF	people	NN0	at	PRP	...
50891891	EAA	...	People	NN0	at	PRP	work	NNI	...
50891892	EAA	...	at	PRP	work	NNI	.	PUN	...
50891893	EAA	...	Work	NNI	.	PUN	there	EX0	...
50891894	EAA	...	.	PUN	there	EX0	will	VM0	...
50891895	EAA	...	there	EX0	will	VM0	be	VBI	...
50891896	EAA	...	will	VM0	be	VBI	significant	AJO	...

Table 4.6 Text meta information table

Text	Register	Topics	Title
EA8	W commerce	leadership; commerce	Making it happen...
EA9	W commerce	hotel management; tourism	The hotel receptionist...
EAA	W soc science	-	Managing people at...
EAI	W soc science	Public law; politics	Public law and political...
EAK	W nat science	scientific research	Nature; London; Macmillan...

database table that has as many rows as the total number of words in the corpus. Each row contains a sequential word in the corpus (along with part of speech and other information, if desired). Most likely, these sequential words will be one of the central columns of the table. This column is then surrounded on the left and right by a number of 'contextual' columns, to create a 'context window' for each word in the corpus.

For example, Table 4.5 represents the main table in our BNC/VIEW database. This table contains about 100 million rows (one for each sequential word in the BNC), with each word in corpus (column word4) surrounded by three words to the left (word1 to word3) and three words to the right (word5 to word7), as well as the 'word offset' ID and the text from which the n-gram is taken (for the sake of brevity, only word3 to word5 is shown here).

In addition to this main sequential n-gram table, there is also a small table that contains 'meta-data' for each of the 4,000+ texts in the BNC. For example, Table 4.6 shows just a few of the columns of information for a handful of entries from this table, based on the [EAA] text.

13 <<http://view.bnc.ucl.ac.uk/>>

Let us now consider how these tables can be used to create quick, powerful queries even on large corpora like the BNC.

#### Basic frequency data for words, phrases, substrings and collocations

Basic queries, even on large corpora such as the 100 million BNC are quite fast with the new architecture. They also have the added benefit of including *all* matching strings – not just the most frequent n-grams, as is the case with some other architectures. For example, to find the most frequent nouns following 'break' + 'a/'the' + NOUN ('break the news', 'break a promise'), the user would enter the following string through the web-based interface. This would be converted to the following SQL command, which would produce the results seen in Table 4.2 above.

(4)

```

QUERY: break a/the [n*]
SQL:
SELECT TOP 100
COUNT(*) WORD4.POS4.WORD5.POS5.WORD6.POS6
FROM [TABLE_NAME] WHERE
WORD4 = 'BREAK' AND
WORDS IN ('THE','A') AND
POS6 LIKE 'NN%';
GROUP BY WORD4,POS4,WORD5,POS5,WORD6,POS6
ORDER BY COUNT(*) DESC

```

Any combination of substrings, words, or part of speech in a seven-word window can be used for the query.

In addition to seeing just raw frequencies, one can also see a 'relevancy-based' display, which is based on a modified *F*-score. With a simple raw-frequency sorting (SORT BY: RELEVANCE] in the search form), a query like [ \* chair ] would produce results like 'a chair', 'the chair', 'his chair', etc. However, if the user selects [ SORT BY: RELEVANCE] in the search form, then the results will be 'high-backed chair', 'sedan chair', 'wicker chair', 'swivel chair', etc. To produce this set of results, the script, (1) finds the raw percentage for all strings matching [ \* chair], (2) finds the overall frequency of the words in the [\*] slot (e.g. 'the', 'wooden', 'swivel', etc.), and then, (3) divides the first figure by the second. In this case, though, the three-

step query still only takes about half of one second to return the results from the 100 million word corpus.

### Comparing related words

A major advantage of storing the n-grams in a relational database is that this architecture lends itself well to frequency comparisons between sets of words. For example, it is possible to find all of the collocates of a given WORD 1 and the collocates of a different WORD 2, and then compare these two lists within the database itself. For example, a user can see the difference between two or more synonyms by finding the collocates that occur with one but not with the other, and all of this can be carried out in just one or two seconds.

As a concrete example, assume that a non-native speaker of English is interested in the difference between 'utter', 'sheer' and 'absolute'. The user would enter the following into the web-based search form:

(5)

QUERY: {utter/sheer/absolute} [n\*]

The curly brackets around the three synonyms indicate that the user wants to find the collocates for each of these three words, and then 'group' these collocates together for each individual word. The results would look like those shown in Table 4.7.

This would indicate to the language learner, for example, that 'sheer' occurs with 'weight' (60 tokens), 'force' (31 tokens) and 'luck' (26 tokens), but that none of these words occur with either 'utter' or 'absolute'. 'Utter', on the other hand, is the only one of the three synonyms that occurs with 'confusion' (ten tokens), 'condemnation' (five tokens) and 'devastation' (five tokens), and 'absolute' is the only synonym that occurs with 'majority' (98 tokens), 'terms' (84 tokens) and 'zero' (54 tokens).

To process this query, the script carries out a number of tasks: (1) it finds all nouns following 'sheer'; (2) it finds all nouns following 'utter'; (3) it finds all nouns following 'absolute'; (4) it stores all of the results for these three queries in a temporary table; and, (5) it runs three separate, sequential SQL queries to find the most frequent collocates for each of these words, which do not occur with either of the other two. (It is also possible to sort by raw frequency with each word, rather than WORD 1 v. WORD 2 / WORD 3, as shown above.)

Notice that comparisons such as these might be possible, but they would certainly be much more cumbersome with other architectures. For example, users of those interfaces could carry out a query with WORD 1, and then WORD 2, and so on. They would then copy and paste the results into separate tables of a database, and then (assuming some skill in SQL), carry out cross-table JOINS to determine the relative frequency with the different words. In this case, however,

Table 4.7 Grouping by synonyms

%	+	-	SHEER	%	+	-	UTTER	%	+	-	ABSOLUTE
1.00	60	-	weight	1.00	19	-	confusion	1.00	98	-	majority
1.00	31	-	force	1.00	5	-	condemnation	1.00	84	-	terms
1.00	26	-	luck	1.00	5	-	devastation	1.00	54	-	zero
1.00	23	-	quantity	1.00	5	-	disregard	1.00	53	-	minimum
1.00	13	-	cliff	1.00	5	-	helplessness	1.00	51	-	value
1.00	12	-	cliffs	1.00	4	-	loneliness	1.00	39	-	egallitarianism
1.00	11	-	coincidence	1.00	3	-	dejection	1.00	33	-	right
1.00	10	-	employment	1.00	3	-	ruthlessness	1.00	27	-	price

the process would probably take at least three to four minutes, whereas with our interface it takes only about 1.2 seconds.

### Integration with other databases

Another important advantage of the relational database approach is that the central n-gram tables can be integrated with other relational databases. For example, the frequency information from the BNC/VIEW databases can be joined with semantic information from other databases such as WordNet,<sup>14</sup> or with personalized lists (relating to semantic fields) created by the user.

Let us take a concrete example, which we have already discussed elsewhere.<sup>15</sup> In order to add a strong semantic component to the BNC/VIEW database, we have imported into a separate database the entire contents of WordNet – a semantically-based hierarchy of hundreds of thousands of words in English. The 'synsets' that make up WordNet indicate roughly equivalent meanings (synonyms), more specific words (hyponyms) and more general words (hyponyms) for a given word, as well as 'parts' (meronyms) of a larger group (e.g. parts of a body) or the larger groups to which a certain word belongs (holonyms, e.g. leg > table/body).

14 C. Fellbaum, *WordNet: An Electronic Lexical Database* (Cambridge, MA: MIT Press, 1998); S. Landes, C. Leacock and R. Tengli, 'Building Semantic Concordances', in C. Fellbaum (ed.), *WordNet: An Electronic Lexical Database* (Cambridge, MA: The MIT Press, 1998), pp. 199–216.

15 M. Davies, 'Semantically-based Queries with a Joint BNC/WordNet Database', in R. Fachinetti (ed.), *Corpus Linguistics Twenty-five Years On* (Amsterdam: Rodopi, 2007), pp. 149–67.

Table 4.8 WordNet/BNC integration: frequency of synonyms of 'sad'

1	Sorry	10,767
2	Sad	3,322
3	Distressing	359
4	Pitiful	199
5	Deplorable	144
6	Lamentable	69

At the most basic level, a user can simply submit as a query like the following:

(6)

QUERY: [=sad]

The script sees the equals sign and interprets this as a query to find all synonyms of 'sad'. The script then retrieves all of the matching words from WordNet and finds the frequency of each of these words in the BNC, *see* Table 4.8. Users can also limit the hits by part of speech, and can also intervene before seeing the frequency listing to select just certain synsets (or meanings) from WordNet.

The WordNet data can also be integrated into more advanced queries. A user can – with one simple query – compare which nouns occur with the different synonyms of a given word. For example, the following query finds all cases of a noun following a synonym of 'bad':

(7)

QUERY: [=bad ] [n\*]

In less than four seconds, the user would then see something similar to Table 4.9. (Note that the format on the web interface is somewhat different to the abbreviated listing shown here.) Such collocational data can be very useful for a language learner, who is probably unsure of the precise semantic range of each adjective. The type of listing given above, which shows the most common nouns with each of the adjectives, can easily permit the language learner to make inferences about the semantic differences between each of the competing adjectives. For example, he or she would see that 'severe illness' occurs but 'wicked illness' does not, and that 'terrible mistake' is common, whereas 'foul mistake' is not.

Finally, it is again worth noting that such a query would be extremely difficult with an architecture that does not allow the user to access other databases that can be tied into the main frequency/n-gram databases. Without our approach, users would probably have to create the list of synonyms in another program, perhaps run queries with each of these words individually, and then collate and sort the

Table 4.9 Synonyms of [bad] + NOUN

PHRASE	FREQ.	PHRASE	FREQ.
disgusting thing	16	severe shortage	26
disgusting way	5	severe weather	43
distasteful species	5	severe winter	49
evil empire	10	terrible accident	26
evil eye	24	terrible blow	13
evil influence	9	terrible danger	14
foul language	29	terrible feeling	19
foul mood	21	terrible mistake	48
foul play	72	terrible shock	41
foul temper	14	wicked grin	6
severe blow	44	wicked people	13
severe burn	28	wicked thing	27
severe damage	59	wicked way	14
severe drought	30	wicked witch	12
severe illness	23		

results in a spreadsheet. Using our approach, all of this is done by the web-based script in less than two seconds.

### Register-based queries

As was discussed previously, one serious shortcoming of many other approaches is that it is either difficult or impossible to compare the results of different sub-corpora – for example, different registers or different historical periods. With the newer BNC/VIEW architecture, however, this is both quite simple and quite fast.

Recall that with the 'sequential n-gram' architecture, each word in the corpus appears on its own line in the database, as shown in Table 4.10. As shown previously, there is also a small table, *see* Table 4.11, that contains 'meta-data' for each of the 4,000+ texts in the BNC.

To limit the query to one particular register (or set of registers), or to compare registers, the script relies on an SQL JOIN between these two tables. It looks for all rows that match the n-gram table (Table 4.10), all texts that match the specific register(s) (Table 4.11), and then limits these hits to just those that have a particular text in common (e.g. EAA). Examples of the types of register-based queries might be queries to find which nouns, or verbs, or three-word lexical

Table 4.10 Sequential n-gram table

ID	Text	word3	pos3	word4	pos4	word5	pos5	...
50891887	EAA ...	a	ATO	small	AJO	group	NN1	...
50891888	EAA ...	small	AJO	group	NN1	of	PRF	...
50891889	EAA ...	group	NN1	of	PRF	people	NN0	...
50891890	EAA ...	of	PRF	people	NN0	at	PRP	...
50891891	EAA ...	people	NN0	at	PRP	work	NN1	...
50891892	EAA ...	at	PRP	work	NN1	,	PUN	...
50891893	EAA ...	work	NN1	,	PUN	there	EX0	...
50891894	EAA ...	,	PUN	there	EX0	will	VM0	...
50891895	EAA ...	there	EX0	will	VM0	be	VBI	...
50891896	EAA ...	will	VM0	be	VBI	significant	AJO	...

Table 4.11 Text meta information table

Text	Register	Topics	Title
EAS	W_commerce	leadership; commerce	Making it happen...
EAG	W_commerce	hotel management; tourism	The hotel receptionist...
EAA	W_soc_science	—	Managing people at...
EAI	W_soc_science	public law; politics	Public law and political...
EAK	W_nat_science	scientific research	Nature. London: Macmillan...

bundles, or collocates with 'chair' occur more in one register (e.g. FICTION) than in another (e.g. ACADEMIC).

For example, suppose that a user wants to find which verbs are more common in legal texts than in academic texts generally. He or she would simply select the following in the web-based form:

- (8)
- QUERY: [w|]
- REGISTER 1 [w\_ac\_politLaw\_edu] (from pull-down menu)
- REGISTER 2 [ACADEMIC]

This searches for all infinitival lexical verbs (VVI) in law texts (w\_ac\_politLaw\_edu), and then all infinitival verbs in ACADEMIC texts as a whole, and then compares the two sets of words. The words – representing verbs that are highly frequent in a legal context – are found in Table 4.12.

Again, note that such comparisons would possibly be quite cumbersome for competing architectures. To the degree that these architectures allow queries by

Table 4.12 Searching by register: lexical verbs in legal texts

Word/phrase	Tokens REG1	Tokens REG2	Per mil in REG1 [4,640,346 words]	Per mil in REG2 [110,789,236 words]	Ratio
Sue	331	10	71.33	0.93	76.96
Cerify	27	2	5.82	0.19	31.39
Adjourn	26	2	5.60	0.19	30.23
Notify	38	3	8.19	0.28	29.45
Waive	38	3	8.19	0.28	29.45
Disclose	190	16	40.95	1.48	27.61
Overrule	23	2	4.96	0.19	26.74
Prohibit	53	5	11.42	0.46	24.65
Plead	62	6	13.36	0.56	24.03

sub-corpora (and not all do), users could carry out a query for [VVI] in Register 1 and then a subsequent query in Register 2. They would then copy and paste the results into separate tables in a database, and then carry out cross-table JOINS to determine the relative frequency in the two registers. Again, however, the process would probably take at least three to four minutes, but using our interface it takes only about two seconds.

### Expanding collocate searches

The architecture described to this point works well for determining the frequency of words, substrings, phrases and 'slot-based' searches (e.g. ADJ + 'world', synonym of ['tired'] + NOUN, etc.). One fundamental disadvantage of the architecture, however, is that it is limited by its very nature to strings that occur within a seven-word window. This is because the 'sequential n-gram' table (see Table 4.5 above) contains that number of columns. Therefore, it would not be possible to find, for example, all the nouns that occur in the wider context of a given adjective using this architecture.

Recently, however, we changed the searching algorithm to allow for the recovery of collocates from a much wider context – up to ten words to the left and to the right of a given word. There are several steps in the script that carry out such queries. Let us take a concrete example to see how this would work. Assume that we are looking for adjectives that are collocates – ten words to the left or right – of the word 'woman'. First, the script finds all of the ID values (=word offset values in the corpus) for 'woman' – more than 22,000 occurrences in the 100 million word corpus – and these are stored in a temp table. We then find all of the adjectives in the corpus whose ID value is between ten more and ten less than the ID values in the temp table. These are placed in a second temp table, and an



SQL command then finds the most common words in this table. Overall, the script takes about five to six seconds to run, and yields adjectives like 'young', 'old', 'beautiful', 'married', etc.

As with the 'slot-based' queries (e.g. ADJ + 'woman') these collocates can then be compared to the collocates of a competing word, such as 'man', to determine with collocates are used with 'woman' much more than with 'man' (e.g. 'childless', 'pretty', 'pregnant', 'distracted', 'fragile', 'desirable'), or more with 'man' than with 'woman' (e.g. 'honourable', 'reasonable', 'military', 'modest', 'rational').

Likewise, one can compare collocates across registers, to look for possible polysemy with a given word. For example, one could look for adjectives with 'chair' that are more common in fiction than in academic text (e.g. 'small', 'hard', 'rocking', 'asleep') or which are more common in academic texts than in fiction (e.g. 'senior', 'philosophical', 'established', 'powerful'). Again, this script only takes about 1.5 seconds to produce the list of contrasting collocates.

### Conclusion

We hope to have demonstrated two fundamental facts in this chapter. First, it is often insightful and advantageous to look at word frequency in context – by register, in historical terms, in syntagmatic terms (collocations) and paradigmatic terms (a given word contrasted with competing words that fill the same slot, such as synonyms). Secondly, the highly-structured relational databases lend themselves well to the comparison of contexts. Word frequency, then, can be analyzed not just as the overall frequency of a given word or lemma in a certain corpus, but, rather, as the frequency of words in a wide range of related contexts.

## Chapter 5

# Issues for Historical and Regional Corpora: First Catch Your Word

Christian Kay

### Introduction

My interest in what words can tell us about a text stems mainly from two electronic projects: the *Historical Thesaurus of English* (HTE)<sup>1</sup> and the *Scottish Corpus of Texts and Speech* (SCOTS).<sup>2</sup> I have a further involvement in the *Linguistic and Cultural Heritage Electronic Network* project (LICHEN), headed by Lisa-Lena Opas-Hänninen at the University of Oulu, Finland, which aims to collect and display languages of the circumarctic region.<sup>3</sup>

### TOE and SCOTS

HTE was replaced in 2008, and already has a daughter project, *A Thesaurus of Old English* (TOE),<sup>4</sup> a conceptually organized thesaurus of the surviving vocabulary of Old English (OE, c. AD 700–1100). Both TOE and SCOTS, which are freely available on the internet, make some attempt to deal with word frequency. Much of the vocabulary of OE is assumed not to have survived, and that which does is unlikely to be representative of the whole. As a result, the editors made use of four flags, somewhat similar to the register labels in modern dictionaries. These are 'o' indicating infrequent use, 'p' for poetic register, 'q' for doubtful forms, and 'g' for words occurring only in glossed texts or glossaries. Unlabelled words

1 C. Kay, J. Roberts, M. Samuels and I. Wotherspoon (eds), *The Historical Thesaurus of the OED* (Oxford: Oxford University Press, forthcoming 2009), see also <http://ibira.englang.arts.gla.ac.uk/historicalthesaurus/>.

2 <http://www.scottishcorpus.ac.uk/>; J. Corbett, J. Anderson, C. Kay and J. Stuart-Smith, funded by AHRB grant B/RE/AN9984/APN17387.

3 See I. Jauso, J. Anderson, W. Anderson, D. Beavan, J. Corbett et al., 'The LICHEN Project: Creating an Electronic Framework for the Collection, Management, Online Display, and Exploitation of Corpora', in A. Hardie (ed.), *Digital Resources in the Humanities Conference 2005 Abstracts* (DRH, 2005), pp. 27–9.

4 An electronic version, supported by British Academy grant LRG-37362, can be seen at <http://ibira.englang.arts.gla.ac.uk/oehtesaurus/>.