Tutorial 6: Corpus Querying

Informatics 1 Data & Analysis

Week 8, Semester 2, 2013–2014

This worksheet has three parts: tutorial Questions, followed by some Examples and their Solutions.

- Before your tutorial, work through and attempt all of the Questions in the first section.

- The Examples are there for additional preparation, practice, and revision.

- Use the Solutions to check your answers, and read about possible alternatives.

You must bring your answers to the main questions along to your tutorial. You will need to be able to show these to your tutor, and may be exchanging them with other students, so it is best to have them printed out on paper.

If you cannot do some questions, write down what it is that you find challenging and use this to ask your tutor in the meeting.

Tutorials will not usually cover the Examples, but if you have any questions about those then write them down and ask your tutor, or go along to InfBASE during the week.

It's important both for your learning and other students in the group that you come to tutorials properly prepared. If you have not attempted the main tutorial questions, then you may be sent away from the tutorial to do them elsewhere.

Some exercise sheets contain material marked with a star ⋆. These are optional extensions.

Data & Analysis tutorials are not formally assessed, but they are a compulsory and important part of the course. If you do not do the exercises then you are unlikely to pass the exam.

Attendance at tutorials is obligatory: if you are ill or otherwise unable to attend one week then email your tutor, and if possible attend another tutorial group in the same week.

Please send any corrections and suggestions to Ian.Stark@ed.ac.uk

---

Introduction

In this tutorial you will be working with a real corpus and analysing it with the *Corpus Query Processor* (CQP), a query engine that searches corpora based on user queries about words, parts of speech, or other markup.

Reading

The following exercises are based largely upon a tutorial written by Stefan Evert and the other maintainers of the *IMS Open Corpus Workbench* at `http://cwb.sourceforge.net`. Before starting these exercises, you must read Sections 1 and 2 of the CQP Query Language tutorial. Download this as a PDF document from the course web page, which hosts the correct version for this worksheet.

Setting up the working environment

These exercises must be carried out on one of the DICE machines in the Informatics computer labs, as those have access to both the CQP engine and the necessary corpora. In particular, you will be using a corpus comprising the works of Charles Dickens, annotated with part-of-speech tags and syntactic structure.

[If connecting remotely, please log in first to student.ssh.inf.ed.ac.uk and then student.compute which has CQP installed. A command-line connection through ssh or PuTTy should be sufficient. See the Informatics Computing Support help pages for more information on remote working.]

You must first declare the location of a corpus index. Type the following at the command line in a terminal window:

```
export CORPUS_REGISTRY=/group/corpora/public/corpus_workbench/registry
```

Take care to copy this exactly: it is case sensitive and there are two underscores.

To start CQP itself, enter the command 'cqp -e'. (The annotation '-e' enables line-editing; try 'cqp -eC' for colour.) You should then see the following in your command prompt:

```
[no corpus]>
```

This means that cqp is active but no specific *corpus* has been selected yet. In order to select the DICKENS corpus type

```
[no corpus]> DICKENS;
```

You will then see the following:

```
DICKENS>
```

Every command line in cqp must finish with a semicolon ';'. You can now get more specific information about this particular corpus by typing:

```
DICKENS> info;
```

Press the space bar to display the next page and 'q' to get back to the command prompt. To leave cqp, enter the command 'exit;'.

1 Concordances

1.1 Concordance of a given word

Having set up a working environment, we can start performing some queries on the DICKENS corpus. Searching for *concordances* (all occurrences of a given word, displayed in context) is a basic form of querying a corpus. To search for a specific word, such as "dear", you should type:

```
DICKENS> [word = "dear"];
```

which can be abbreviated to

```
DICKENS> "dear";
```

As previously, press space for the next page and 'q' to get back to the command prompt.

We can make any search case insensitive with the '%c' modifier:

```
DICKENS> "dear" %c;
```

We can search for more than one word after another by adding more queries:

```
DICKENS> "dear" "friend";
```

1.2 More complex queries

The use of *regular expressions* considerably extends the power of the CQP query language. CQP uses the following format for regular expressions:

`exp1 exp2`: first `exp1` then `exp2` in sequence
`exp*`: zero or more occurrences of `exp`
`exp?`: zero or one occurrences of `exp`
`exp+`: one or more occurrences of `exp`
`exp1|exp2`: either `exp1` or `exp2`

For example, the following query retrieves four variations of the word "regular": regular, regulars, regularly, regularity.

```
DICKENS> "regular(|s|ly|ity)";
```

A single dot ('.') can be used as a wildcard in CQP, matching any single character. For example, this query finds words such as "than", "then" and "thin":

```
DICKENS> "th.n";
```

while the following query finds words starting "un..." and ending "...ly", , such as "unusually" and "unsuccessfully":

```
DICKENS> "un.*ly";
```

The use of boolean logic extends further the power of CQP queries. The operators '&' (and) and '|' (or) can be used to combine two tests, while '!=' tests for inequality. For instance, the following query would find all occurrences of words beginning "un" and ending "ly", but different from "unlikely":

```
DICKENS> [(word = "un.*ly") & (word != "unlikely")];
```

Similarly, the following query would find all occurrences of the words "man" and "woman":

```
DICKENS> [(word = "man") | (word = "woman")];
```

Question 1: Concordance of a given bigram. A *bigram* is any two-word sequence in some text. Retrieve all occurrences of the bigram "great deal" in the Dickens corpus. Is this query affected by making it case insensitive?

Question 2: Concordance of variations of a given word. Retrieve variations of the word "kiss" in the corpus: "kisses", "kissing", ... Is this query affected by making it case insensitive?

2 Annotations

2.1 Displaying linguistic annotations

When working with `cqp` you may want to customise how the results of queries are displayed. For example, we can select whether or not to display part-of-speech (POS) annotations on the corpus. By default these are turned off, but they can be enabled with

```
DICKENS> show +pos;
```
(Query results will *show* annotations)

and disabled with

```
DICKENS> show -pos;
```
(Query results will *hide* annotations)

Try some of your queries again, and look at the annotations.

2.2 Accessing linguistic annotations

Once we have annotations on text, we can use these make more refined searches. For example, with the POS annotations we can search for uses of adjectives, verbs or nouns in a given corpus. Switch on POS tagging with "`show +pos;`" and do some searches to identify the different annotations that used for verbs, adjectives and nouns.

These tags are taken from the *Penn Treebank* project. You can find out more about the Treebank at `http://www.cis.upenn.edu/~treebank` and a list of tags at `http://is.gd/treebanktagset`.

In the same way that a query "`[word = "Dear"]`" will find a particular word in the corpus, a query "`[pos = "NN"]`" will find all (singular) nouns in the corpus. Both "word" and "pos" queries can use regular expressions and logical operations to build more complex queries.

Question 3. What is the common pattern for the POS annotations used for different types of verbs? Use a regular expression to formulate a query that retrieves all occurrences of verbs in the corpus.

Question 4. Write a query to retrieve all tokens which are *not nouns*.

Some of the following questions are marked with a star ⋆. This indicates that they are optional — you are encouraged to attempt all you can, but they are not a requirement for tutorials.

⋆ **Question 5.** Retrieve all occurrences of the word "lock" where it is used as a noun.

⋆ **Question 6.** The "lemma" annotation in `cqp` contains the *headword*: the principal dictionary word associated with a word. Switch on lemma tagging with "`show +lemma;`" and try some queries.

Use this to formulate an alternative (case-insensitive) query for all variations of the word "kiss" in the corpus. Does this return the same results as your answer to Question 2?

3 Frequencies of a given word

Frequency information obtained from corpora can be useful for language research. Here we look at finding the absolute frequency — the number of occurrences — of certain features in the corpus.

To find the total number of tokens in the corpus, type this:

```
DICKENS> Q1 = []; size Q1;
```

Here we have named the result of the query Q1, and then retrieved its size. In this case, the query is one that matches every word in the corpus (and, in fact, every punctuation mark too).

A similar query counts the absolute frequency of a specific word:

```
DICKENS> Q2 = [word = "dear"]; size Q2;
```

This query finds the absolute frequency of all bigrams consisting of the word "dear" followed by a noun:

```
DICKENS> Q3 = [word = "dear"] [pos = "N.*"]; size Q3;
```

We can refine this to list the absolute frequency of the bigram grouped by which noun it is that follows "dear":

```
DICKENS> Q3 = [word = "dear"] [pos = "N.*"]; count Q3 by word;
```

A slightly different presentation is:

```
DICKENS> Q3 = [word = "dear"] [pos = "N.*"];
DICKENS> group Q3 matchend word by match word;
```

Here "`matchend word`" refers to the word at the end of the bigram match (the noun) and "`match word`" refers to the word at the start of the bigram (in this case, "dear"). You can also group by POS, or lemma.

A tricky point: Suppose we want to find the most common word in the corpus. A natural query would be:

```
DICKENS> Q4 = []; count Q4 by word;
```

This works, but also includes punctuation marks and white space amongst words. Instead, the following works better:

```
DICKENS> Q4 = [word = "[a-zA-Z].*"]; count Q4 by word;
```

Here the range expression "`[a-zA-Z]`" matches any letter, either lower or upper case.

Question 7 - Absolute Frequency. Find the absolute frequency of the word "girl" in the corpus, ignoring any case differences.

Question 8 - Relative Frequency. Calculate the relative frequency of the word "girl" in the corpus (the number of occurrences as a fraction of the total number of words).

⋆ **Question 9.** Find the most common bigram in the Dickens corpus.

4 Collocations

A *collocation* is a sequence of words that occur in text more often than would be expected by chance; more generally, that form expressions through repeated use in certain contexts.

Question 10. List the 10 most common adjective-noun pairs in the Dickens corpus. Which of these would you classify as collocations?

Question 11. For a given adjective-noun pair AB (i.e., adjective A followed by noun B) we can create a *contingency table* of frequency observations. Such contingency tables will be used later in the course. The table for AB looks as follows:

	A	$\neg A$
B	F_{AB}	$F_{\neg A B}$
$\neg B$	$F_{A \neg B}$	$F_{\neg A \neg B}$

where:

F_{AB} is the absolute frequency of adjective A followed by noun B
$F_{\neg A B}$ is the absolute frequency of any adjective but A followed by noun B
$F_{A \neg B}$ is the absolute frequency of adjective A followed by any noun but B
$F_{\neg A \neg B}$ is the absolute frequency of any adjective but A followed by any noun but B

Construct the contingency table of case-insensitive frequency observations for the bigram "great deal". What (if anything) does this table tell us about the bigram?

5 Tutorial Discussion

- Suggest some current and possible future uses of corpora in Informatics.

- Suggest some current and possible future uses of corpora in Linguistics.

- How do you think these two kinds of use might be related?

Examples

This section contains further exercises on Corpus Querying. These examples are similar to the main tutorial questions: they involve querying the Dickens corpus with CQP, and they cover topics like concordances, annotations, word frequencies and collocations. Instructions on how to use CQP are provided in the tutorial introduction.

Example 1: Concordances

(a) Retrieve all occurrences of the word "proud" in the Dickens corpus.

(b) Retrieve all occurrences of the word "proud" where the query is case-insensitive.

(c) Retrieve variations of the word "proud" in the corpus: "proudly", "proudest", ...

(d) Retrieve all occurrences of the bigram "sweet girl" in the corpus.

Example 2: Annotations

(a) Write a query to retrieve all adjective occurrences in the corpus.

(b) Write a query to retrieve all tokens which are adjectives or adverbs.

(c) Retrieve all occurrences of the word "face" where it is not used as a noun.

Example 3: Word frequencies

(a) Find the absolute frequency of the word "face" in the Dickens corpus.

(b) Find the absolute frequency of the word "face" in the Dickens corpus, where it is not used as a noun.

(c) Find the absolute frequency of all bigrams consisting of an adjective followed by the word "girl".

(d) Find the absolute frequency of all bigrams consisting of an adjective followed by the word "girl", grouped by the adjective that precedes "girl".

Example 4: Collocations

Below you can see a list of bigrams in the Dickens corpus, ordered by their frequency.

```
111     next morning  [#64236-#64346]
100     little while  [#53448-#53547]
97      other people  [#73043-#73139]
96      other side  [#73396-#73491]
94      own room  [#75644-#75737]
93      many times  [#57550-#57642]
88      old friend  [#66892-#66979]
85      such things  [#98498-#98582]
85      young friend  [#111398-#111482]
78      young men  [#113500-#113577]
76      long way  [#55010-#55085]
```

For each of the following bigrams, use CQP to construct a concordance for the word pair, and assess whether you think it forms a collocation.

(a) "old friend"

(b) "other people"

(c) "next morning"

Solutions to Examples

These are not entirely "model" answers; instead, they indicate a possible solution. Remember that not all of these questions will have a single "right" answer. If you have difficulties with a particular example, or have trouble following through the solution, please raise this as a question in your tutorial.

Solution 1

(a) Retrieve all occurrences of the word "proud" in the Dickens corpus.

```
DICKENS> [word = "proud"];
```

which can be abbreviated to

```
DICKENS> "proud";
```

(b) Retrieve all occurrences of the word "proud" where the query is case-insensitive.

```
DICKENS> "proud" %c;
```

(c) Retrieve variations of the word "proud" in the corpus: "proudly", "proudest", ...

```
DICKENS> "proud(|er|est|ly)";
```

or

```
DICKENS> "proud|prouder|proudest|proudly";
```

or

```
DICKENS> "proud.*";
```

(d) Retrieve all occurrences of the bigram "sweet girl" in the corpus.

```
DICKENS> "sweet" "girl";
```

or

```
DICKENS> [word = "sweet"] [word = "girl"];
```

Solution 2

(a) Write a query to retrieve all adjective occurrences in the corpus.

```
DICKENS> [pos = "JJ.*"];
```

(b) Write a query to retrieve all tokens which are adjectives or adverbs.

```
DICKENS> [(pos = "JJ.*") | (pos = "RB.*")];
```

(c) Retrieve all occurrences of the word "face" where it is not used as a noun.

```
DICKENS> [(word = "face") & (pos != "N.*")];
```

Solution 3

(a) Find the absolute frequency of the word "face" in the Dickens corpus.

```
DICKENS> Q1 = [word = "face"]; size Q1;
```

(b) Find the absolute frequency of the word "face" in the Dickens corpus, where it is not used as a noun.

```
DICKENS> Q2 = [(word = "face") & (pos != "N.*")]; size Q2;
```

(c) Find the absolute frequency of all bigrams consisting of an adjective followed by the word "girl".

```
DICKENS> Q3 = [pos = "JJ.*"] [word = "girl"]; size Q3;
```

(d) Find the absolute frequency of all bigrams consisting of an adjective followed by the word "girl", grouped by the adjective that precedes "girl".

```
DICKENS> Q4 = [pos = "JJ.*"] [word = "girl"];
DICKENS> group Q4 matchend word by match word ;
```

Solution 4

We cannot conclude with certainty whether any of the bigrams provided are statistically true collocations — whether they actually occur more often than would be expected by chance, given the underlying frequency of individual words in the corpus — unless we do some further analysis, such as χ^2 testing (we will discuss these later on in this course). However, by looking at examples of uses of the bigrams in the corpus we can argue whether these bigrams have a semantic weight of their own, and hence whether they could be considered to be collocations.

(a) "old friend"

This bigram has in some cases the meaning of "close friend" or "beloved friend". A representative example in the corpus is "Well met , my dear old friend ! said I ." In these cases, "old friend" could be regarded as a collocation.

(b) "other people"

This is not a collocation, as the expression simply refers to other people.

(c) "next morning"

This bigram is used in some cases in the corpus to mean "tomorrow". In these cases it might be considered to be a collocation.