

AntConc 3.4.0w (Windows) 2014
Global Settings Tool Preferences Help

Concordance Plot File View Clusters/N-Grams Collocates Word List Keyword List

Concordance Hits 66
Hit KWIC

1 on her to leave her neat parlour, hung
2 ured and obliging. He had gone three miles
3 chance, he said, that he had not gone
4 times over; her compliments to all friends go
5 her compliments to all friends go round and
6 very sad--but he is always wanted all
7 the water, and she, by the sudden whirling
8 muffin last night--if it had been handed
9 piece of news indeed. I shall just go
10 Woodhouse--well, at last, I fancy, he looked
11 Hartfield, he thought I had much better go

round with fancy-work, whenever she cou
round one day in order to bring her some
round by Randalls. He did not think we ev
round and round again; and if she does bu
round again; and if she does but send her
round the country. I suppose there is not
round of something or other among the s
round once, I think it would have been en
round by Mrs. Cole's; but I shall not
round and saw me; for instead of going or
round by Mr. Cole's stables, for I should

em
em
em
em
em
em
em
em
em
em
em

450 MILLION WORDS, 1990-2012 (DOWNLOAD ALL 190,000 TEXTS)

SEE CONTEXT: CLICK ON WORD (ALL SECTIONS) OR NUMBER (SPECIFIED SECTION)

SEC 1: 11,926,481 WORDS

WORD/PHRASE	TOKENS 1	TOKENS 2	PM 1	PM 2	RATIO
1 SHALT	20	5	1.68	0.36	4.72
2 DARE	19	8	1.59	0.57	2.80
3 SHALL	768	174	64.39	12.36	2.57
4 OUGHT	355	196	25.49	13.93	1.83
5 LET	304	196	25.49	13.93	1.83
6 'S	304	196	25.49	13.93	1.83
7 'D	527	223	44.19	15.84	1.69
8 NEED	319	223	26.75	15.84	1.41
9 MIGHT	7250	6047	607.89	429.62	1.1
10 CA	820	810	68.75	67.55	1.1
11 MUST	8100	889	8688	68.50	1.1
12 'LL	817	10163	68.50	722.04	1.1
13 SHOULD	9062	13	759.82	0.92	1.1
14 WILT	11	21277	0.92	1,511.65	0.1
15 WOULD	16621	415	1,393.62	25.32	0.1
16 WO	302	20288	2,454.03	1,441.39	0.1
17 MAY	14348	41902	836.12	2,976.98	0.1
18 CAN	29268	14328	57.44	73.53	0.1
19 COULD	9972	1035	1,402.01	1,972.67	0.1
20 USED	685	27766			
21 WILL	16721				

SEC 2: 14,075,316 WORDS

WORD/PHRASE	TOKENS 2	TOKENS 1	PM 2	PM 1
1 WILL	27766	16721	1,972.67	1,402.01
2 USED	1035	685	73.53	57.44
3 COULD	14328	9972	1,017.95	2,454.03
4 CAN	41902	29268	2,976.98	1,203.62
5 MAY	20288	14348	1,441.39	25.49
6 WO	415	302	1,511.65	1,393.62
7 WOULD	21277	11	0.92	759.82
8 WILT	13	11	722.04	68.50
9 SHOULD	10163	9062	63.16	617.25
			57.55	607.89
			429.62	607.89
			15.84	26.75
			25.29	44.19
			13.93	25.49
			13.93	25.49
			12.36	29.77
			25.08	64.39

Log-likelihood
Mutual information
MI3
Z-score
T-score
Log-likelihood
Dice coefficient
Rank by frequency

PRACTICAL

Corpus Linguistics

An Introduction to Corpus-Based Language Analysis

MARTIN WEISSER



WILEY Blackwell

Practical Corpus Linguistics

Practical Corpus Linguistics

An Introduction to Corpus-Based
Language Analysis

Martin Weisser

WILEY Blackwell

This edition first published 2016
© 2016 John Wiley & Sons, Inc

Registered Office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

Editorial Offices

350 Main Street, Malden, MA 02148-5020, USA

9600 Garsington Road, Oxford, OX4 2DQ, UK

The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

For details of our global editorial offices, for customer services, and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com/wiley-blackwell.

The right of Martin Weisser to be identified as the author of this work has been asserted in accordance with the UK Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and authors have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. It is sold on the understanding that the publisher is not engaged in rendering professional services and neither the publisher nor the author shall be liable for damages arising herefrom. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Weisser, Martin, author.

Practical corpus linguistics : an introduction to corpus-based language analysis / Martin Weisser. – First edition.

pages cm

Includes bibliographical references and index.

ISBN 978-1-118-83187-8 (hardback) – ISBN 978-1-118-83188-5 (paper) 1. Linguistic analysis (Linguistics)–Databases. 2. Linguistic analysis (Linguistics)–Software. 3. Corpora (Linguistics)–Methodology. 4. Corpora (Linguistics)–Technological innovations. 5. Computational linguistics–Methodology. 6. Computer network resources–Evaluation. 7. Citation of electronic information resources. I. Title.

P128.D37W45 2016

410.1'88–dc23

2015023709

A catalogue record for this book is available from the British Library.

Cover image: © Martin Weisser

Set in 10.5/13pt Galliard by Aptara Inc., New Delhi, India

To Ye & Emma,
who've had to suffer
from an undue lack of attention
throughout the final months
of writing this book

Contents

List of Figures	xiii
List of Tables	xv
Acknowledgements	xvii
1 Introduction	1
1.1 Linguistic Data Analysis	3
1.1.1 What's data?	3
1.1.2 Forms of data	3
1.1.3 Collecting and analysing data	7
1.2 Outline of the Book	8
1.3 Conventions Used in this Book	10
1.4 A Note for Teachers	11
1.5 Online Resources	11
2 What's Out There?	13
2.1 What's a Corpus?	13
2.2 Corpus Formats	13
2.3 Synchronic vs. Diachronic Corpora	15
2.3.1 'Early' synchronic corpora	15
2.3.2 Mixed corpora	18
2.3.3 Examples of diachronic corpora	20
2.4 General vs. Specific Corpora	21
2.4.1 Examples of specific corpora	22
2.5 Static Versus Dynamic Corpora	25
2.6 Other Sources for Corpora	26

	Solutions to/Comments on the Exercises	26
	Note	28
	Sources and Further Reading	28
3	Understanding Corpus Design	29
3.1	Food for Thought – General Issues in Corpus Design	29
3.1.1	Sampling	30
3.1.2	Size	31
3.1.3	Balance and representativeness	32
3.1.4	Legal issues	32
3.2	What’s in a Text? – Understanding Document Structure	33
3.2.1	Headers, ‘footers’ and meta-data	34
3.2.2	The structure of the (text) body	36
3.2.3	What’s (in) an electronic text? – understanding file formats and their properties	37
3.3	Understanding Encoding: Character Sets, File Size, etc.	38
3.3.1	ASCII and legacy encodings	38
3.3.2	Unicode	39
3.3.3	File sizes	40
	Solutions to/Comments on the Exercises	41
	Sources and Further Reading	42
4	Finding and Preparing Your Data	43
4.1	Finding Suitable Materials for Analysis	44
4.1.1	Retrieving data from text archives	44
4.1.2	Obtaining materials from Project Gutenberg	44
4.1.3	Obtaining materials from the Oxford Text Archive	45
4.2	Collecting Written Materials Yourself (‘Web as Corpus’)	46
4.2.1	A brief note on plain-text editors	46
4.2.2	Browser text export	48
4.2.3	Browser HTML export	49
4.2.4	Getting web data using ICEweb	50
4.2.5	Downloading other types of files	52
4.3	Collecting Spoken Data	53
4.4	Preparing Written Data for Analysis	56
4.4.1	‘Cleaning up’ your data	56
4.4.2	Extracting text from proprietary document formats	58
4.4.3	Removing unnecessary header and ‘footer’ information	58
4.4.4	Documenting what you’ve collected	59
4.4.5	Preparing your data for distribution or archiving	60
	Solutions to/Comments on the Exercises	62
	Sources and Further Reading	66
5	Concordancing	67
5.1	What’s Concordancing?	67

5.2	Concordancing with AntConc	69
5.2.1	Sorting results	74
5.2.2	Saving, pruning and reusing your results	75
	Solutions to/Comments on the Exercises	78
	Sources and Further Reading	81
6	Regular Expressions	82
6.1	Character Classes	84
6.2	Negative Character Classes	86
6.3	Quantification	86
6.4	Anchoring, Grouping and Alternation	87
6.4.1	Anchoring	87
6.4.2	Grouping and alternation	88
6.4.3	Quoting and using special characters	90
6.4.4	Constraining the context further	91
6.5	Further Exercises	92
	Solutions to/Comments on the Exercises	93
	Sources and Further Reading	100
7	Understanding Part-of-Speech Tagging and Its Uses	101
7.1	A Brief Introduction to (Morpho-Syntactic) Tagsets	103
7.2	Tagging Your Own Data	109
	Solutions to/Comments on the Exercises	113
	Sources and Further Reading	120
8	Using Online Interfaces to Query Mega Corpora	121
8.1	Searching the BNC with BNCweb	122
8.1.1	What is BNCweb?	122
8.1.2	Basic standard queries	123
8.1.3	Navigating through and exploring search results	124
8.1.4	More advanced standard query options	126
8.1.5	Wildcards	126
8.1.6	Word and phrase alternation	128
8.1.7	Restricting searches through PoS tags	129
8.1.8	Headword and lemma queries	131
8.2	Exploring COCA through the BYU Web-Interface	132
8.2.1	The basic syntax	133
8.2.2	Comparing corpora in the BYU interface	135
	Solutions to/Comments on the Exercises	137
	Sources and Further Reading	145
9	Basic Frequency Analysis – or What Can (Single) Words Tell Us About Texts?	146
9.1	Understanding Basic Units in Texts	146
9.1.1	What’s a word?	147
9.1.2	Types and tokens	149

9.2	Word (Frequency) Lists in AntConc	151
9.2.1	Stop words – good or bad?	156
9.2.2	Defining and using stop words in AntConc	158
9.3	Word Lists in BNCweb	160
9.3.1	Standard options	160
9.3.2	Investigating subcorpora	162
9.3.3	Keyword lists	169
9.4	Keyword Lists in AntConc and BNCweb	169
9.4.1	Keyword lists in AntConc	169
9.4.2	Keyword lists in BNCweb	172
9.5	Comparing and Reporting Frequency Counts	175
9.6	Investigating Genre-Specific Distributions in COCA	178
	Solutions to/Comments on the Exercises	179
	Sources and Further Reading	192
10	Exploring Words in Context	193
10.1	Understanding Extended Units of Text	194
10.2	Text Segmentation	195
10.3	N-Grams, Word Clusters and Lexical Bundles	196
10.4	Exploring (Relatively) Fixed Sequences in BNCweb	198
10.5	Simple, Sequential Collocations and Colligations	198
10.5.1	‘Simple’ collocations	198
10.5.2	Colligations	200
10.5.3	Contextually constrained and proximity searches	201
10.6	Exploring Colligations in COCA	202
10.7	N-grams and Clusters in AntConc	205
10.8	Investigating Collocations Based on Statistical Measures in AntConc, BNCweb and COCA	207
10.8.1	Calculating collocations	207
10.8.2	Computing collocations in AntConc	209
10.8.3	Computing collocations in BNCweb	210
10.8.4	Computing collocations in COCA	211
	Solutions to/Comments on the Exercises	212
	Sources and Further Reading	226
11	Understanding Markup and Annotation	227
11.1	From SGML to XML – A Brief Timeline	229
11.2	XML for Linguistics	230
11.2.1	Why bother?	230
11.2.2	What does markup/annotation look like?	230
11.2.3	The ‘history’ and development of (linguistic) markup	232
11.2.4	XML and style sheets	234
11.3	‘Simple XML’ for Linguistic Annotation	236
11.4	Colour Coding and Visualisation	240
11.5	More Complex Forms of Annotation	246

Solutions to/Comments on the Exercises	248
Sources and Further Reading	253
12 Conclusion and Further Perspectives	254
Appendix A: The CLAWS C5 Tagset	259
Appendix B: The Annotated Dialogue File	261
Appendix C: The CSS Style Sheet	269
Glossary	271
References	277
Index	283

List of Figures

3.1	Illustration of basic document structure	34
4.1	The ICEweb interface	50
5.1	Example of a KWIC concordance output	69
5.2	The AntConc startup screen	70
5.3	AntConc file opening options	71
5.4	AntConc file settings	71
5.5	AntConc ‘Corpus Files’ window (two files loaded)	72
5.6	AntConc ‘Search Term’ and search options	72
5.7	AntConc results for <i>round</i> in two novels by Jane Austen	73
5.8	AntConc ‘Search Window Size’ options	73
5.9	AntConc ‘Kwic Sort’ options	74
6.1	Sample paragraph for practising and understanding regex patterns	84
7.1	Sample output of the Simple PoS Tagger	110
8.1	The BNCweb startup screen	122
8.2	Results for simple search for <i>assume</i>	123
8.3	BNCweb query follow-on options	125
8.4	The basic COCA interface	133
8.5	Display of antonyms <i>thoughtful</i> and <i>thoughtless</i> as alternatives	134
8.6	Side-by-side comparison for the lemma of <i>movie</i> in the COCA and BNC	136
9.1	Output of a basic frequency list in AntConc	152
9.2	Token (word) (re-)definition in AntConc	154
9.3	AntConc Word List preferences	159
9.4	BNCweb frequency list selection options	161
9.5	Options for defining subcorpora in BNCweb	162
9.6	Excel text import wizard (stage 1)	164

9.7	Excel text import wizard (stage 3)	165
9.8	Sort options in Excel	166
9.9	Options for defining subcorpora according to genre	167
9.10	BNCweb keyword and title scan	168
9.11	AntConc Keyword preferences	170
9.12	Keyword options in BNCweb	172
9.13	Keyword comparison of university essays and written component of the BNC (top 31 entries)	173
10.1	Illustration of a collocation span	208
10.2	Options for statistical collocation measures in BNCweb	209
11.1	A brief SGML sample	232
11.2	CSS sample paragraph styling	235
11.3	TEI header for BNC file KST	247

List of Tables

2.1	Extract from <i>Beowulf</i> , encoded/represented in two different ways	14
2.2	Early written corpora	16
2.3	Composition of the Brown Corpus	17
2.4	Some examples of (earlier) spoken corpora	18
2.5	Early mixed corpora	19
2.6	Modern mega corpora	20
2.7	Examples of diachronic corpora	21
2.8	Examples of academic corpora	22
2.9	Examples of learner corpora	23
2.10	Selection of pragmatically annotated corpora	24
3.1	Common file formats and their properties	37
7.1	Ambiguous PoS tags in the Brown Corpus	102
7.2	The Penn Treebank tagset (based on Taylor et al. 2003: 8)	103
7.3	The CLAWS 7 (C7) tagset	105
8.1	Wildcards and their uses for investigating linguistic features	129
8.2	Simplified tags in BNCweb	132
8.3	Word + PoS tags breakdown for <i>mind</i>	142
9.1	Top 15 most frequent word types in section A of the LOB Corpus	157
9.2	Recalculated norming sample from Biber, Conrad & Reppen (1998: 263)	175
11.1	Annotation types listed in Garside et al. 1997	228
11.2	CSS properties for XML visualisation exercise	240
11.3	Colour semantics and CSS styles	243

Acknowledgements

I'd first like to start by thanking my former students in Chemnitz, Bayreuth and Hong Kong who 'suffered through' the initial sets of teaching materials that eventually formed the basis for writing this textbook. The next big thanks needs to go to my colleagues here at Guangdong University of Foreign Studies, who attended a series of workshops where I tested out the materials from the preliminary drafts of several chapters and who provided me with highly useful feedback. Particular mention here deserves to go to Junyu (Mike) ZHANG, who not only commented on the content of several chapters, but also pointed out certain issues of style that have hopefully helped me to make the writing more accessible to an international readership.

My next round of thank yous goes to Yanping DONG and Hai XU, for allowing me to join the National Key Research Center for Linguistics and Applied Linguistics at Guangdong University of Foreign Studies, which has provided me with more of the desperately needed time to focus on writing this book, while also allowing me to conduct other types of research that have influenced the contents of the book in various ways. To Hai XU, I give additional thanks for sharing his experience in, and knowledge of, corpora of Chinese, which has, unfortunately, only partially found its way into this book, due to limits of space. To my other colleagues, especially Yiqiong ZHANG, I also give thanks for providing a more moral type of support through engaging in further discussions and making me feel at home in the Center.

I also owe a great debt to Laurence Anthony for allowing me to pester him with a series of questions about and suggestions for improving AntConc. More or less the same, though possibly to a slightly lesser extent, goes for Sebastian Hoffmann and Mark Davies for answering questions about particular features of, and again

partly responding to requests for improving, BNCweb and the COCA interface, respectively.

Next, I'd sincerely like to thank the anonymous reviewers of this book, who, through their many invaluable constructive comments, have not only encouraged me in writing the book, but also hopefully allowed me to improve the contents substantially.

My final – but most important and heartfelt – note of thanks goes to Geoff Leech. Although credit for introducing me to the study of corpus linguistics has to go to someone else, he's certainly been the single most important influence on my career and thinking as a corpus linguist. I'll forever be grateful for his ongoing support throughout my years spent at Lancaster, working with him, as external examiner of my PhD, and later up to his untimely demise in August 2014. I sincerely hope that he would have appreciated the design and critical aims of this textbook, and perhaps also recognised his implicit hand in shaping it...

I

Introduction

This textbook aims to teach you how to analyse and interpret language data in written or orthographically transcribed form (i.e. represented as if it were written, if the original data is spoken). It will do so in a way that should not only provide you with the technical skills for such an analysis for your own research purposes, but also raise your awareness of how corpus evidence can be used in order to develop a better understanding of the forms and functions of language. It will also teach you how to use corpus data in more applied contexts, such as e.g. in identifying suitable materials/examples for language teaching, investigating socio-linguistic phenomena, or even trying to verify existing linguistic theories, as well as to develop your own hypotheses about the many different aspects of language that can be investigated through corpora. The focus will primarily be on English-language data, although we may occasionally, whenever appropriate, refer to issues that could be relevant to the analysis of other languages. In doing so, we'll try to stay as theory-neutral as possible, so that no matter which 'flavour(s)' of linguistics you may have been exposed to before, you should always be able to understand the background to all the exercises or questions presented here.

The book is aimed at a variety of readers, ranging mainly from linguistics students at senior undergraduate, Masters, or even PhD levels who are still unfamiliar with corpus linguistics, to language teachers or textbook developers who want to create or employ more real-life teaching materials. As many of the techniques we'll be dealing with here also allow us to investigate issues of style in both literary and non-literary text, and much of the data we'll initially use actually consists of fictional works because these are easier to obtain and often don't cause any copyright

issues, the book should hopefully also be useful to students of literary stylistics. To some extent, I also hope it may be beneficial to computer scientists working on language processing tasks, who, at least in my experience, often lack some crucial knowledge in understanding the complexities and intricacies of language, and frequently tend to resort to mathematical methods when more linguistic (symbolic) ones would be more appropriate, even if these may make the process of writing ‘elegant’ and efficient algorithms more difficult.

You may also be asking yourself why you should still be using a textbook at all in this day and age, when there are so many video tutorials available, and most programs offer at least some sort of online help to get you started. Essentially, there are two main reasons for this: a) such sources of information are only designed to provide you with a basic overview, but don’t actually teach you, simply demonstrating how things are done. In other words they may do a relatively good job in showing you one or more ways of doing a few things, but often don’t really allow you to use a particular program independently and for more complex tasks than the author of the tutorial/help file may actually have envisaged. And b) online tutorials, such as the ones on YouTube, may not only take a rather long time to (down)load, but might not even be (easily) accessible in some parts of the world at all, due to internet censorship.

If you’re completely new to data analysis on the computer and working with – as opposed to simply opening and reading – different file types, some of the concepts and methods we’ll discuss here may occasionally make you feel like you’re doing computer science instead of working with language. This is, unfortunately, something you’ll need to try and get used to, until you begin to understand the intricacies of working with language data on the computer better, and, by doing so, will also develop your understanding of the complexity inherent in language (data) itself. This is by no means an easy task, so working with this book, and thereby trying to develop a more complete understanding of language and how we can best analyse and describe it, be it for linguistic or language teaching purposes, will often require us to do some very careful reading and thinking about the points under discussion, so as to be able to develop and verify our own hypotheses about particular language features. However, doing so is well worth it, as you’ll hopefully realise long before reaching the end of the book, as it opens up possibilities for understanding language that go far beyond a simple manual, small-scale, analysis of texts.

In order to achieve the aims of the book, we’ll begin by discussing which types of data are already readily available, exploring ways of obtaining our own data, and developing an understanding of the nature of electronic documents and what may make them different from the more traditional types of printed documents we’re all familiar with. This understanding will be developed further throughout the book, as we take a look at a number of computer programs that will help us to conduct our analyses at various levels, ranging from words to phrases, and to even larger units of text. At the same time, of course, we cannot ignore the fact that there may be issues in corpus linguistics related to lower levels, such

as that of morphology, or even phonology. Having reached the end of the book, you'll hopefully be aware of many of the different issues involved in collecting and analysing a variety of linguistic – as well as literary – data on the computer, which potential problems and pitfalls you may encounter along the way, and ideally also how to deal with them efficiently. Before we start discussing these issues, though, let's take a few minutes to define the notion of (linguistic) data analysis properly.

1.1 Linguistic Data Analysis

1.1.1 What's data?

In general, we can probably see all different types of language manifestation as language data that we may want/need to investigate, but unfortunately, it's not always possible to easily capture all such 'available' material for analysis. This is why, apart from the 'armchair' data available through introspection (cf. Fillmore 1992: 35), we usually either have to collect our materials ourselves or use data that someone else has previously collected and provided in a suitable form, or at least a form that we can adapt to our needs with relative ease. In both of these approaches, there are inherent difficulties and problems to overcome, and therefore it's highly important to be aware of these limitations in preparing one's own research, be it in order to write a simple assignment, a BA dissertation, MA/PhD thesis, research paper, etc.

Before we move on to a more detailed discussion of the different forms of data, it's perhaps also necessary to clarify the term *data* itself a little more, in order to avoid any misunderstandings. The word itself originally comes from the plural of the Latin word *datum*, which literally means '(something) given', but can usually be better translated as 'fact'. In our case, the data we'll be discussing throughout this book will therefore represent the 'facts of language' we can observe. And although the term itself, technically speaking, is originally a plural form referring to the individual facts or features of language (and can be used like this), more often than not we tend to use it as a singular mass noun that represents an unspecified amount or body of such facts.

1.1.2 Forms of data

Essentially, linguistic data comes in two general forms, written or spoken. However, there are also intermediate categories, such as texts that are written to be spoken (e.g. lectures, plays, etc.), and which may therefore exhibit features that are in between the two clear-cut variants. The two main media types often require rather radically different ways of 'recording' and analysis, although at least some of the techniques for analysing written language can also be used for analysing transliterated or (orthographically) transcribed speech, as we'll see later when looking at some dialogue data. Beyond this distinction based on medium, there are of

course other classification systems that can be applied to data, such as according to *genre*, *register*, *text type*, etc., although these distinctions are not always very clearly formalised and distinguished from one another, so that different scholars may sometimes be using distinct, but frequently also overlapping, terminology to represent similar things. For a more in-depth discussion of this, see Lee (2002).

To illustrate some of the differences between the various forms of language data we might encounter, let's take a look at some examples, taken from the Corpus of English Novels (CEN) and Corpus of Late Modern English Texts, version 3.0 (CLMET3.0; De Smet, 2005), respectively. To get more detailed information on these corpora, you can go to <https://perswww.kuleuven.be/~u0044428/>, but for our purposes here, it's sufficient for you to know that these are corpora that are mainly of interest to researchers engaged in literary stylistic analyses or historical developments within the English language. However, as previously stated, throughout the book, we'll often resort to literary data to illustrate specific points related to both the mechanics of processing language and as examples of genuinely linguistic features. In addition to being fictional, this data will often not be contemporary, simply because much contemporary data is often subject to copyright. Once you understand more about corpora and how to collect and compile them yourself, though, you'll be able to gather your own contemporary data, should you wish so, and explore actual, modern language in use.

Apart from being useful examples of register differences, the extracts provided below also exhibit some characteristics that make them more difficult to process using the computer. We'll discuss these further below, but I've here highlighted them with boxes.

Sample A – from *The Glimpses Of The Moon* by Edith Wharton, published 1922

IT rose for them--their honey-moon--over the waters of a lake so famed as the scene of romantic raptures that they were rather proud of not having been afraid to choose it as the setting of their own.

"It required a total lack of humour, or as great a gift for it as ours, to risk the experiment," Susy Lansing opined, as they hung over the inevitable marble balustrade and watched their tutelary orb roll its magic carpet across the waters to their feet.

"Yes--or the loan of Strefford's villa," her husband emended, glancing upward through the branches at a long low patch of paleness to which the moonlight was beginning to give the form of a white house-front.

Sample B – from: *Eminent Victorians* by Lytton Strachey, published 1918

Preface

THE history of the Victorian Age will never be written; we know too much about it. For ignorance is the first requisite of the historian—ignorance, which simplifies and clarifies, which selects and omits, with a placid perfection unattainable by the highest art. Concerning the Age which has just passed, our fathers and our grandfathers have poured forth and accumulated so vast a quantity of information that the industry of a Ranke would be submerged by it, and the perspicacity of a Gibbon would quail

before it. It is not by the direct method of a scrupulous narration that the explorer of the past can hope to depict that singular epoch. If he is wise, he will adopt a subtler strategy. He will attack his subject in unexpected places; he will fall upon the flank, or the rear; he will shoot a sudden, revealing searchlight into obscure recesses, hitherto undivined. He will row out over that great ocean of material, and lower down into it, here and there, a little bucket, which will bring up to the light of day some characteristic specimen, from those far depths, to be examined with a careful curiosity.

Sample C – from *The Big Drum* by Arthur Wing Pinero, published 1915

Noyes.

[Announcing Philip.] Mr. Mackworth.

Roope.

[A simple-looking gentleman of fifty, scrupulously attired—jumping up and shaking hands warmly with Philip as the servant withdraws.] My dear Phil!

Philip.

[A negligently—almost shabbily—dressed man in his late thirties, with a handsome but worn face.] My dear Robbie!

Roope.

A triumph, to have dragged you out! [Looking at his watch.] Luncheon isn't till a quarter-to-two. I asked you for half-past-one because I want to have a quiet little jaw with you beforehand.

Philip.

Delightful.

Roope.

[Er—I]’d better tell you at once, old chap, whom you’ll meet here [to-day].

Sample A is clearly a piece of narrative fiction, mixing narrative description and simulated reported speech, references to characters and situations that are depicted as life-like, as well as featuring a number of at least partly evaluative reporting verbs, such as *opined* and *emended*. **Sample B**, on the other hand, contains no reported speech and reporting verbs, although it’s clearly also narrative – albeit non-fictional –, with a relatively complex sentence structure, including numerous relative and adverbial clauses, and an overall high degree of formality. **Sample C**, in contrast, exhibits clear characteristics of (simulated) spoken language, much shorter and less complex syntax, even single-word ‘sentences’, with names, titles and informal terms of address (*old chap*) used when the characters are addressing/introducing each other, exclamations, contractions, and at least one hesitation marker (*Er*). And even though the language in the latter sample seems fairly natural, we can still easily see that it comes from a scripted text, partly because of the indication of speakers (which I’ve highlighted in bold-face), and partly due to the stage instructions included in square brackets.

As we haven’t discussed any of the issues in processing such text samples yet, it may not be immediately obvious to you that these different types of register may potentially require different analysis approaches, depending on what our exact

aims in analysing them are. For instance, for **Sample A**, do we want to conceptually treat the reported speech as being of the same status as the descriptive parts, and do we thus want to analyse them together or separately? Or are we possibly just interested in how the author represents the direct speech of the characters in the novel, and would therefore want to extract only that? And if so, how would we best go about this?

Sample B is probably relatively straightforward to analyse in terms of perhaps a frequency analysis of the words, but what if we're also interested in particular aspects of syntax or lexis that may be responsible for its textual complexity or the perceived level of formality, respectively? And, last but not least, concerning **Sample C**, similarly to **Sample A**, which parts of the text would we be interested in here and how would we extract them? Are the stage instructions equally important to us as the direct speech exchanges between the characters? Or, if, for example, we're interested in the average number of words uttered by each character, how do we deal with hesitation markers? Do we treat them as words or 'non-words' simply to be deleted? As I've already tried to hint at in the beginning of this paragraph, the answers to these questions really depend on our research purpose(s), and can thus not be conclusively stated here.

Something else you may have noticed when looking at the samples I've provided above is that they're all from the early 20th century. As such, the language we encounter in them may sometimes appear overly formal (or even archaic) to some extent, compared to the perhaps more 'colloquial' language we're used to from the different registers these days. I've chosen extracts from these three particular texts and period for a number of reasons: a) their authors all died more than 70 years ago so the texts are in the public domain; in other words, there are no copyright issues, even when quoting longer passages; b) they are included in corpus compilations; and c) they not only illustrate register/genre differences but also how the conventions for these may change over time, as can be seen, for example, in the spelling of *to-day* in the final extract.

As pointed out above, another interesting aspect of these samples is that they exhibit particular formatting issues, which again may not be immediately apparent to you yet, but are due to somewhat bizarre typographical conventions. If you look closely at the samples, you can see that in **Sample A** there are double dashes marking the parenthetical counterpart (i.e. reference resolution) "their honey-moon" to the sentence-initial cataphoric pronoun "IT". What is in fact problematic to some extent for processing the text is that these dashes actually look like double hyphens, i.e. they're not surrounded by spaces on either side, as would be the normal convention. Now, many computer programs designed to count words will split the input text on spaces and punctuation. Unfortunately, though, this would leave us with some very strange 'words' (that superficially look like hyphenated compounds), *them—their* and *honey-moon—over*, in any resulting word-frequency list. This is obviously something we do not want and which introduces errors into any automatic analysis of the data. Something similar, albeit not to signal a parenthetical but instead some kind of pseudo-punctuation, happens

again for “Yes—or” a little further down in the text. We can already see, therefore, from this relatively short sample of text that a failure to deal with this feature could cause issues in a number of places throughout the text. The same problem occurs in the other two samples, only that there the dash doesn’t actually consist of two separate characters, but one single *m-dash*.

A different problem occurs in the use of initial capitals in **Samples A** and **B**. As you can see, the words *it* and *the* appear in capital letters throughout, signalling the beginning of the chapter typographically. Again, as ‘human consumers’ of the text, this will not cause any processing problems, but for the computer, *the*, *The*, and *THE* are in fact three different ‘words’, or at least word forms. Thus, even single initial capitals at the beginning of sentences may become issues in identifying and counting words on the computer. We’ll talk more about this type of issue in Section 4.4.1, where we’ll explore ways of dealing with such features of the text in order to retain relatively ‘clean’ data.

1.1.3 Collecting and analysing data

When collecting our own data, we obviously need to consider *methodologies* that allow us to collect the right types and amount(s) of data to answer our particular *research questions*. This, however, isn’t the only type of consideration necessary, but we also need to bear in mind *ethical issues* involved in the collection – such as asking people for permission to record them or to publish their recordings, etc. – and which type of *format* that data should be stored in so as to be most useful to us, and potentially also other researchers.

When using other people’s existing data, there are usually issues in accessing data stored in their specific format(s) or converting the data to a format that is more suitable to one’s own needs, as we’ve just seen above, such as removing unwanted types of information or transforming overly specific information into simpler forms of representation. In this textbook, we’ll also look at some of the important aspects of collecting or adapting data to one’s needs, as well as how to go about analysing and presenting them in various ways, once a suitable format has been established.

In order to be able to work with electronic data, we also need to become familiar with a variety of different programs, some written specifically for linguistic analysis, some for more general purposes of working with texts. One of the key features of this book is that the programs I’ll recommend to you are almost exclusively obtainable free of charge, i.e. so-called *freeware*. This doesn’t mean that there aren’t other excellent programs out there that may do some of the tasks we want to perform even better, or in simpler or more powerful ways, but simply reflects the fact that there are already many useful free programs available, and also my own philosophy that we shouldn’t need to spend substantial amounts of money just to enable us to do research. This is at least part of the reason why I make most of my own programs available to the research community in this way, apart from the fact that this makes my own research (results) more easily reproducible by

others, and therefore caters for the aims of satisfying *accountability* and academic honesty. For the sake of completeness, though, I'll generally try to at least refer to alternative commercial programs, but without discussing them in any detail.

Corpus linguistics, as a form of data analysis methodology, can of course be carried out on a number of different operating systems, so I'll also try to make recommendations as to which programs may be useful for the most commonly used ones, Windows, Mac OS X, and Linux. Because there are many different 'flavours' of Linux, though, with a variety of different windowing interfaces, I'll restrict my discussions to two of the most popular ones, KDE and Gnome. Unfortunately, I won't be able to provide detailed support on how to actually install the programs themselves, as this may sometimes involve relatively detailed information about your system that I cannot predict. Instead, however, I'll actually try to avoid/pre-empt such issues by recommending default programs that are probably already installed, provided that they do in fact fulfil all or at least most of our needs.

1.2 Outline of the Book

This book is organised into four sections. The first section (comprising Chapters 1 and 2) begins with a very brief introduction to the history and general design of corpora, simply to 'set the scene', rather than to provide an extensive coverage of the multitude of corpora that have been created for different purposes and possibly also made available for free or in the form of various types of interfaces. More extensive coverage on the subject, including more theoretical implications, is already provided in books like Kennedy (1998), Meyer (2002), or Lindquist (2009), so these texts can always be consulted for reference if necessary, and we can instead focus on more practical issues. For a more detailed interesting discussion of some of the different 'philosophical' approaches to corpus linguistics, you can consult McEnery and Hardie (2012).

The introductory section is followed by an overview of different methods to compile and prepare corpora from available online resources, such as *text archives* or the *WWW*. This section (spanning Chapters 3 and 4) should essentially provide the basis for you to start building your own corpora, but also introduces you to various issues related to handling language on the computer, including explanations of different file types you may encounter or want to use, as well as certain types of meta-information about texts.

Section 3 (Chapters 5 to 10) then deals with different approaches to corpus-based linguistic data analysis, ranging from basic searching (concordancing) via learning about more complex linguistic patterns, expressed in the form of regular expressions, to simple and extended word (frequency) list analyses. This part already contains information on how to tag your data morpho-syntactically, using freely available tagging resources, and how to make use of tagging in your analyses. The final section then takes the notion of adding linguistic information to

your data further, and illustrates how to enrich corpus data using basic forms of XML in order to cyclically improve your analyses or publish/visualise analysis results effectively.

As corpus linguistics is a methodology that allows us to develop insights into how language works by ‘consulting’ real-life data, it should be fairly obvious that we cannot learn how to do corpus research on a purely theoretical basis. Therefore, as far as possible, all sections of this book will be accompanied by practical exercises. Some of these will appear to be relatively straightforward, almost mechanical, ones where you simply get to follow a sequence of steps in order to learn how to use a specific function inside a program or web interface, while others are more explicitly designed to enable you to develop your own strategies for solving problems and testing hypotheses in linguistics. Please bear in mind, though, that for the former type of exercise, simply following the steps blindly without trying to understand why you’re doing them will not allow you to learn properly. So, as far as possible, at each point you should try to understand what we’re trying to achieve and how the particular program we’re using only gives us a handle on producing the relevant data, but does not actually answer our research questions for us. In the same vein, it’s also important to understand that once we actually have extracted some relevant data from a corpus, this is rarely ever the ‘final product’. Such data generally either still needs to be interpreted, filtered, or evaluated as to its usefulness, if necessary by (re-)adjusting the search strategy or initial hypotheses and/or conclusions, or, if it’s to be used for more practical purposes, such as in the creation of teaching materials or exercises, to be brought into an appropriate form.

As we move on and you learn more and more techniques, the exercises will also get more complex, sometimes assuming the size of small research projects, if carried out in full detail. As a matter of fact, as these exercises require and consolidate a lot of the knowledge gained in prior sections, they might well be suitable for small research projects to be set by teachers, and possibly even form the basis of BA theses or MA dissertations.

Of course, you won’t be left alone in figuring out the solutions to these exercises; both types will be solved at the end of each respective section, either in the form of detailed and precise explanations, or, whenever the results might be open to interpretation, by appropriate comments illustrating what you could/should be able to observe. For the more extensive exercises referred to in the previous paragraph, I’ll often start you off with suitable explanations regarding the procedures to follow, and also hint at some potential issues that may arise, but will leave the completion up to you, to help you develop your awareness independently. Furthermore, as real corpus linguistics is not just about getting some ‘impressive’ numbers but should in fact allow you to gain real insights into different aspects of language, you should always try to relate your results to what you know from established theories and other methods used in linguistics, or even other related disciplines, such as for example sociology, psychology, etc., as far as they may be relevant to answering your research questions. This is also why the solutions to,

and discussions of, the exercises may often represent those parts of the book that cover some of the more theoretical aspects of corpus linguistics, aspects that you'll hopefully be able to master once you've acquired the more practical tools of the trade. Thus, even if you may think you've already found a perfect answer to an exercise, you should probably still spend some time reading carefully through each solution.

As this textbook is more practical in nature than other textbooks on corpus linguistics, at the end of almost all chapters, I've also added a section entitled 'Sources and Further Reading'. These sections essentially provide lists of references I've consulted and/or have found most useful and representative in illustrating the particular topic(s) discussed in the chapter. You can consult these references if you want to know more about theoretical or practical issues that I am unable to cover here, due to reasons of space. These sections may not necessarily contain highly up-to-date references, for the simple reason that, unfortunately, later writings may not always represent improvements over the more fundamental works produced in some of the areas covered. Once you understand more about corpus linguistics, though, you may want to consult the individual chapters in two of the recent handbooks, O'Keefe & McCarthy (2010) and Lüdeling & Kytö (2008), so that you can evaluate the progress made over recent years yourself.

1.3 Conventions Used in this Book

In linguistics, there are many conventions that help us to distinguish between different levels of analysis and/or description, so as to better illustrate which different types of language phenomena we're dealing with at any given point in time. Throughout this book, I'm going to make use of many, if not most, of these conventions, so it's important to introduce them at this point. In addition to using these conventions as is done in linguistics, I may also use some of them to indicate special types of textual content relevant to the presentation of resources in this book, etc.

Double quotes ("...") indicate direct speech or short passages quoted from books.

Single quotes ('...') signal that an expression is being used in an unusual or unconventional way, that we're referring to the meaning of a word or construction on the semantic level, or to refer to menu items or sometimes button text in programs used. The latter may also be represented by a stylised button text, e.g.

Start

Curly brackets ({...}) are used to represent information pertaining to the level of morphology.

Angle brackets (<...>) indicate that we're dealing with issues related to orthography or spelling. Alternatively, they're also used in certain types of linguistic annotation.