FOURTH EDITION

HALLIDAY'S INTRODUCTION TO FUNCTIONAL GRAMMAR

M.A.K. Halliday
Revised by Christian M.I.M. Matthiessen

COMPANION @ WEBSITE

"Essential reading ... teeming with insights."
Michael Toolan, *University of Birmingham, UK*

ROUTLEDGE

Third edition published 2004 by Hodder Education, an Hachette UK company

This fourth edition published in 2014
by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

Simultaneously published in the USA and Canada
by Routledge
711 Third Avenue, New York, NY 10017

Routledge is an imprint of the Taylor & Francis Group, an informa business

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data
Halliday, M. A. K. (Michael Alexander Kirkwood), 1925–
[Introduction to functional grammar]
Halliday's introduction to functional grammar / M.A.K. Halliday and Christian Matthiessen. – Fourth Edition
pages cm
Previous ed. published as: Introduction to functional grammar, 2004.
Includes bibliographical references and index.
1. Functionalism (Linguistics) 2. Grammar, Comparative and general. I. Matthiessen, Christian M. I. M., author. II. Title.
P147.H35 2013
410.1'8–dc23
2013006799

ISBN: 9780415826280 (hbk)
ISBN: 9781444146608 (pbk)
ISBN: 9780203431269 (ebk)

Typeset in 10 on 12.5pt Berling
by Phoenix Photosetting, Chatham, Kent

TENTS

This page intentionally left blank

VENTIONS

Systemic description

Capitalization labels used in systems and realization statements

Operators in system specifications

Operators in realization statements

Graphic conventions in system networks

there is a system *x/y* with entry condition *a* [if *a*, then either *x* or *y*]

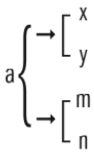

there are two simultaneous systems *x/y* and *m/n*, both having entry condition *a* [if *a*, then both either *x* or *y* and, independently, either *m* or *n*]

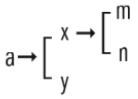

there are two systems *x/y* and *m/n*, ordered in dependence such that *m/n* has entry condition *x* and *x/y* has entry condition *a* [if *a* then either *x* or *y*, and if *x*, then either *m* or *n*]

there is a system *x/y* with compound entry condition, conjunction of *a* and *b* [if both *a* and *b*, then either *x* or *y*]

there is a system *m/n* with two possible entry conditions, disjunction of *a* and *c* [if either *a* or *c*, or both, then either *m* or *n*]

Annotation of text

Boundary markers

Other forms of annotation

Example sources

Sources of examples are given in square brackets after examples. The main types are listed in the table below.

Other conventions

Bold font is used to indicate (first mention of) technical terms, as in:

> Each foot, in turn, is made up of a number of **syllables**

Italic font is used to indicate grammatical and lexical items and examples cited in the body of the text, as in:

> Here, the Theme *this responsibility* is strongly foregrounded

This page intentionally left blank

ℜODUCTION

The first edition of *Halliday's introduction to functional grammar* (IFG) appeared in 1985. It was, among other things, an introduction to the systemic functional **theory** of grammar that M.A.K. Halliday initiated through the publication of his 1961 article 'Categories of the theory of grammar' (although his publications on the grammar of Chinese go back to 1956). It was at the same time an introduction to the **description** of the grammar of English that he had started in the early 1960s (see e.g. Halliday, 1964). Thus, the first edition of IFG was an introduction both to a functional theory of the grammar of human language in general and to a description of the grammar of a particular language, English, based on this theory. The relationship between theory and description was a dialogic one: the theory was illustrated through the description of English, and the description of English was empowered by the theory. Halliday could have used any other language for this purpose rather than English – for example, Chinese, since he had worked on Chinese since the late 1940s. The theory had been developed as a theory of grammar in general, and by the mid-1980s it had already been deployed and tested in the description of a number of languages.

Around half a century has passed since Halliday's first work on the general theory of grammar and his first work on the description of English, and around a quarter of a century has passed since IFG1 appeared: that edition represents the mid-point between the early work and today's continued theoretical and descriptive research activities, activities that were enabled by IFG1 and are reflected in IFG4. When IFG1 appeared, it was the only introduction of its kind, a summary of the work by Halliday and others undertaken since the early 1960s. It was a 'thumbnail sketch'. He had already published accounts of various areas, accounts that were in many respects more detailed than the sketches in IFG – e.g. his account of transitivity and theme (Halliday, 1967/8), his interpretation of modality (Halliday, 1970) and his description of grammar and intonation (Halliday, 1967a). He had also worked on a manuscript

presenting a comprehensive account of the grammar of English, *The meaning of modern English*; many aspects of this account such as his interpretation of tense in English were only sketched in IFG1. In addition, researchers had contributed significant text-based studies of grammar and of intonation based on his framework. These informed the description of English, but have not been published since text-based accounts were not welcomed by publishers in the period dominated by formal generative linguistics.

Since IFG1 appeared a quarter of a century ago, and IFG2 followed nine years later in 1994, systemic functional linguists have published other complementary volumes drawing on IFG in different ways, designed to serve different communities of users; these include Geoff Thompson's *Introducing functional grammar* (first edition in 1996; second in 2004, with the third about to appear), Meriel and Thomas Bloor's *Functional analysis of English: a Hallidayan approach* (first edition in 1995; second in 2004), my own *Lexicogrammatical cartography: English systems* (1995), Graham Lock's *Functional English grammar: An introduction for second language teachers* (1996), and the IFG workbook by Clare Painter, J.R. Martin and myself (first edition: *Working with functional grammar*, 1998; second edition: *Deploying functional grammar*, 2010). In addition, researchers have contributed many journal articles and book chapters to thematic volumes dealing with particular aspects of IFG or reporting on research based on the IFG framework. For a summary of the rich work in the IFG framework, see Matthiessen (2007b). However, researchers have also complemented IFG stratally, moving from the account of lexicogrammar presented in IFG to the stratum of semantics; book-length accounts include Martin's *English text* (1992) and Halliday's and my *Construing experience* (1999, republished in 2006).

By the time Halliday generously invited me to take part in the project of producing IFG3, the ecological niche in which IFG operates had thus changed considerably – certainly for the better. It had, in a sense, become more crowded; but this meant that IFG3 could develop in new ways. Thanks to Geoff Thompson's more introductory *Introducing functional grammar* and to other contributions of this kind, we were able to extend IFG in significant ways, perhaps making the third edition more of a reference work and less of a beginner's book than the previous two editions had been. We certainly included features of the grammar of English that had not been covered before, and we provided a more comprehensive sketch of the overall theoretical framework in Chapters 1 and 2. In preparing the third edition, we worked extensively with corpora of different kinds – resources that had become more accessible since IFG1, supported by computational tools that had been developed since that edition; and we included many examples drawn from corpora, and from our own archives of text. In addition, we included system networks for all the major areas of the grammar.

In my own *Lexicogrammatical cartography: English system* (1995), LexCart, I had used system networks as a cartographic tool, organizing the presentation of the description of the grammar in terms of the system networks – ranging across metafunctions and down ranks and taking a number of steps in delicacy. These system networks were derived from a system network of the clause that Halliday had put together for a computational project initiated by Nick Colby at UC Irvine and then taken over as the seed of the Nigel grammar as part of the Penman project directed by Bill Mann at the Information Sciences Institute, USC, in 1980 (this system network has now been published as part of Halliday's collected works). As a research linguist working on Mann's project since the beginning, I expanded this clause network, and added networks for other parts of the grammar – with

the help of Halliday and other systemic functional linguists (see Matthiessen, 1995a, and cf. Matthiessen, 2007b). When we added system networks to IFG3, we did not try to organize the overall presentation in terms of them as I had done in LexCart, since IFG already had its own logic of presentation, which included more reasoning about the development of the account than I had included in LexCart.

In preparing IFG4, I have followed the trajectory from IFG1 to IFG3, while at the same time keeping in mind changes in the environment in which this fourth edition will appear. I have continued working with corpora, benefiting from new resources generously made available to the research community such as COCA (see Chapter 2). A great deal of this work is, quite naturally, 'under the hood': as with IFG3, many fishing expeditions are reflected by only one or two examples, or by just a brief note in passing, and many other expeditions are only reflected indirectly. Along the way, there have been various interesting findings that there is no space to report on in IFG4, like changes in the use of 'gush' as a verb in *Time Magazine* since the 1920s, or more generally in the use of verbs of saying over that period. In working with corpora, I was at various points tempted to replace all examples from older corpora dating back to the 1960s with examples from more recent ones; but I decided against it for various reasons – an important one being that, like any other language, English is an assemblage of varieties of different kinds (cf. Chapter 2, Section 2.4), including temporal dialects: the collective system of a language typically spans a few generations – never in a state of being, always in a process of becoming. And even more than a few generations: while Chaucer is almost out of range, Shakespeare is not.

One new feature in IFG4 is the introduction of a scheme for classifying texts according to contextual variables, presented in Chapter 1. In Chapter 2 through to Chapter 10, I have classified all the short texts and text extracts according to this scheme. This is a step in the direction of illuminating the grammar at work in different text types – of supporting the understanding of a language as an assemblage of registers. We hope that the website companion to IFG4 (see below) will make it possible to provide many more text examples.

Another feature of IFG4 is the continued expansion of references to theoretical frameworks and to descriptive work on English in systemic functional linguistics but also in other frameworks. Here it is, of course, impossible to be comprehensive, or even to achieve a balanced representation of references to relevant contributions. In his preface to Volume 1 of his *Basic linguistic theory*, Dixon refers to 'quotationitis', introducing it as 'a fashion in linguistics', and characterizing it as 'attempting to cite every single thing published on or around a topic, irrespective of its quality or direct relevance', and then pointing to problems with this 'fashion'. At the same time, it is very important that readers of IFG should be able to follow up on particular points mentioned in the book and go beyond the material presented here; and these days scholars are increasingly subjected by governments to ill-conceived and destructive frameworks designed to measure their output and impact in terms of publications, so citations make a difference. At one point, I thought that the solution in the area of description might be to cite central passages in the major reference grammars of English. However, on the one hand, this would actually be a significant project in its own right, and on the other hand, these reference grammars are not, on the whole, designed as gateways to the literature. I hope that the website companion to IFG4 will be able to provide more bibliographic information. And various online search facilities are helping students and researchers find relevant references.

IFG4 can be used as a reference work supporting more introductory accounts, or as a textbook in its own right. In either case, there are a number of books that are an important part of the environment in which IFG operates – theoretical and descriptive accounts of grammar (e.g. Halliday, 2002b, 2005; Butt *et al.*, 2000; Thompson, 2004; Bloor & Bloor, 2004; Eggins, 2004; Matthiessen, 1995a; Martin, Matthiessen & Painter, 2010; Matthiessen & Halliday, 2009; Caffarel, Martin & Matthiessen, 2004), of (prosodic) phonology (e.g. Halliday & Greaves, 2008) and of semantics (e.g. Martin, 1992; Eggins & Slade, 2005; Martin & Rose, 2007; Halliday & Matthiessen, 2006). Accounts of language development both in the home and the neighbourhood before school (e.g. Halliday, 1975, 2004; Painter, 1984, 1999) and in school (see Christie & Derewianka, 2008, for a recent summary of research and report on their own research from early primary school to late secondary school in Australia) give a unique insight into the ontogenetic beginnings and continual expansion of lexicogrammar, and also a very rich understanding of the grammar at work in everyday and educational contexts. Recent overviews of systemic functional linguistics include Hasan, Matthiessen & Webster (2005, 2007), Halliday & Webster (2009); and, through the window of terminology, Matthiessen, Teruya & Lam (2010). Here it is very important to note that Systemic Functional Grammar (SFG) is only one part of Systemic Functional Linguistics (SFL). If one is working on English, it is, of course, always helpful to have the standard reference grammars of English within easy reach – Quirk *et al.* (1985), Biber *et al.* (1999) and Huddleston & Pullum (2002), as well as overviews of descriptions of English such as Aarts & McMahon (2006).

In addition, IFG4 will be supported by a dedicated website. At the time of writing, I am still working on material for the website, but it is clear that the site will offer additional examples, extended text illustrations, sources of examples cited, additional pointers to the literature, colour versions of a number of figures in IFG4 and probably additional displays, the appendices from the first two editions of IFG and the foreword, and, I hope, in-depth discussions of certain topics. I also hope that it will, at least to some extent, be possible to take account of alternative descriptions, both systemic functional ones based on the framework of the 'Cardiff grammar', developed by Robin Fawcett, Gordon Tucker and their team of colleagues, researchers and students, and functional ones from other traditions, as well as formal ones where there are interesting convergences or illuminating differences. I hope the website will make it possible to treat IFG4 as a 'live document'.

Let me round off this introduction on a personal note. When I saw the first drafts of parts of IFG1 around 1980 or 1981, I was working as a research linguist on a computational linguistic text-generation project directed by Bill Mann (cf. Matthiessen & Bateman, 1991; Matthiessen, 2005). Halliday was a consultant on the project and had (as mentioned above) already contributed an 'algebraic' representation of the core systems of the clause as a foundation of the computational grammar part of the text generation system, the 'Nigel grammar', and with the help of the first drafts and earlier published system networks, I expanded the description for the computational grammar. Halliday and I had both started on the project in mid-1980. In the course of this project and its successors, I was very fortunate to learn from him how to develop grammatical descriptions – holistically, as global outlines rather than as local grammar fragments; and I learned how to model grammar and how to produce descriptions that are explicit enough for computational modelling.

However, my interest in Halliday's work and in systemic functional linguistics more generally had started during my undergraduate days in general linguistics and English linguistics at Lund

University in the 1970s. As an undergraduate student in linguistics, I was taught to develop descriptions of fragments of grammar using the version of Chomsky's generative grammar that was current at the time (a version of the 'Extended Standard Theory'); I remember working on mood tags – without any of the insights that Halliday's account brings to this area of the grammar of English. But we were also encouraged to explore different theoretical frameworks, by the two professors of Linguistics during my time there as a student, Bertil Malmberg and then Bengt Sigurd. And in the Department of English, where I was also a student, there was a great deal of interest in Halliday & Hasan's (1976) account of cohesion – a contribution that stimulated a number of PhD theses in that department, as part of the reorientation to corpus-based research brought about by the new Professor of English linguistics, Jan Svartvik. (In those days, it was still possible for students to construct their own study paths; I had added Arabic and Philosophy to my particular mix.)

When I first came across Systemic Functional Linguistics back at Lund University, something clicked – or rather a number of things clicked. I realized that Halliday had solved a problem that had puzzled and bothered me for quite a long time – since secondary school, where I had come across Alvar Ellegård's highly original introduction to generative semantics and also Bertil Malmberg's introduction to European structuralism. Both approaches seemed full of insight and promising – one providing a deeper understanding of structure and the other showing the power of the paradigmatic axis. However, they appeared to be completely incompatible. It was only when I read Halliday's work that I understood how systemic (paradigmatic) organization could be related to structural (syntagmatic) organization through realization statements. His theory of paradigmatic organization and the relationship between the paradigmatic axis and the syntagmatic one is one of the major breakthroughs in twentieth-century theoretical linguistics. Later I became aware of other breakthroughs he had quietly made, including his theory of metafunctions, his theory of instantiation and his theory of grammatical metaphor.

In working on the description of English in a computational linguistics context, and on the description of Akan in a typological linguistic context, I also came to appreciate the descriptive power of systemic functional theory, including the heuristic value of developing a description with the help of a function-rank matrix (see Chapter 2). I still remember very clearly the quite extraordinary sensation I had when I began auditing the first seminars I had ever attended by Halliday – a course he gave at UC Irvine starting around March 1980: this was the first time anyone had ever given me a clear sense of the overall organization of language as a complex semiotic system. I thought to myself that he was the first linguist to teach me about language; previously other linguists had taught me about linguistics. There is a very significant fundamental difference between the two; and language is much harder to understand (and so to teach about) than linguistics!

I was very fortunate to start working on the systemic functional description of English in 1980 under Halliday's guidance. His descriptions were often quite 'unorthodox' in the sense that they differed significantly from 'mainstream' accounts — for example, his account of the clause as a metafunctional grammatical construct, his account of grammar and lexis as zones within a lexicogrammatical continuum (rather than as separate 'modules'), his account of transitivity in English based on the complementarity of the transitive and ergative models, his account of theme and information as complementary textual systems, his account of modality as a cline for propositions and proposal between positive and negative polarity

extended through interpersonal grammatical metaphor, his account of tense as a logical system for construing serial time (as opposed to a combination of tense and aspect), his account of hypotactic verbal group complexes and of clause complexes (contrasting with accounts based on the notion of complementation).

Naturally, in working on the computational grammar in the 1980s, I tried out more fashionable accounts that were part of the received tradition; but every time I experimented I came to realize how much more insightful Halliday's accounts were – being part of (and thus revealing patterns within) the overall system of the grammar. He never tried to convince me – never tried to pull rank (although in his position, I would've been very tempted to tell me: 'just take my word for it'), but, instead, he taught me how to work things out for myself.

One of the early areas I worked on was tense; when I finally understood his account, and was able to appreciate the advance it represented over both tense-aspect accounts that were popular at the time and Hans Reichenbach's sketch of a temporal logic from the 1940s that had been adopted in a number of more recent linguistic and computational linguistic accounts, I experienced the sense of an *Aha-Erlebnis* for the first time in my life – the term had been introduced to us in high school (I probably learned the term 'epiphany' much later), but I think I had only understood it theoretically before: I suddenly understood the deep insight embodied in Halliday's description of the English grammar of serial time.

On another occasion I was trying to come to grips with 'serial verb constructions' in Akan in the mid-1980s and I suddenly realized that Halliday's account of hypotactic verbal group complexes was a much better model than the assumption (still common at the time) that some form of complementation was involved. But I've already gone on too long … I just wanted to convey both my sense of the extraordinary intellectual excitement of being involved in the long-term research programme of which IFG has turned out to be an important part and my enormous sense of gratitude to Halliday for his mentorship, and also for his fortitude – for daring to be so dramatically different from the mainstream even at the cost of being ignored and effaced by its practitioners and for daring to develop appliable linguistics at a time when application was a sign of theoretical impurity.

As I tinker with Michael Halliday's *Introduction to functional grammar*, I am yet again reminded of my enormous debt to him — a debt that I am very happy to see increase over the decades; it will continue to accumulate interest for as long as I live. At the same time, I'm also happily aware of all the colleagues and students who have engaged with IFG, asking questions and giving comments that have informed my work on the fourth edition. I am deeply grateful to all of them. It's impossible to mention everyone; but I have benefited in particular from the researchers who have done PhDs with me developing comprehensive descriptions of the clause grammars of a rich range of languages: Alice Caffarel on French, Kazuhiro Teruya on Japanese, Minh Duc Thai on Vietnamese, Eden Li on Chinese, Pattama Patpong on Thai, Ernest Akerejola on Òkó, Abhishek Kumar on Bajjika and Mohamed Ali Bardi on Arabic.

Christian M.I.M. Matthiessen
The Hong Kong Polytechnic University
Hong Kong

PART I
THE CLAUSE

This page intentionally left blank

PTER

ONE

THE ARCHITECTURE OF LANGUAGE

1.1 Text and grammar

When people speak or write, they produce **text**; and text is what listeners and readers engage with and interpret. The term 'text' refers to any instance of language, in any medium, that makes sense to someone who knows the language; we can characterize text as language functioning in context (cf. Halliday & Hasan, 1976: Ch. 1; Halliday, 2010). Language is, in the first instance, a resource for making meaning; so text is a process of making meaning in context.

To a grammarian, text is a rich, many-faceted phenomenon that 'means' in many different ways. It can be explored from many different points of view. But we can distinguish two main angles of vision: one, focus on the text as an object in its own right; two, focus on the text as an instrument for finding out about something else. Focusing on text as an object, a grammarian will be asking questions such as: Why does the text mean what it does (to me, or to anyone else)? Why is it valued as it is? Focusing on text as instrument, the grammarian will be asking what the text reveals about the system of the language in which it is spoken or written. These two perspectives are clearly complementary: we cannot explain why a text means what it does, with all the various readings and values that may be given to it, except by relating it to the linguistic system as a whole; and, equally, we cannot use it as a window on the system unless we understand what it means and why. But the text has a different status in each case: either viewed as **artefact**, or else viewed as **specimen**.

The text itself may be lasting or ephemeral, momentous or trivial, memorable or soon forgotten. Here are three examples of text in English:

Text 1-1: Exploring text (spoken, monologic)

Today all of us do, by our presence here, and by our celebrations in other parts of our country and the world, confer glory and hope to newborn liberty.

Out of the experience of an extraordinary human disaster that lasted too long, must be born a society of which all humanity will be proud. Our daily deeds as ordinary South Africans must produce an actual South African reality that will reinforce humanity's belief in justice, strengthen its confidence in the nobility of the human soul and sustain all our hopes for a glorious life for all.

All this we owe both to ourselves and to the peoples of the world who are so well represented here today.

Text 1-2: Recommending text (written, monologic)

Cold power is the **ideal brand for any family**.

We understand that there is more than one thing you want to achieve out of every wash load.

As such, we have developed a formula capable of achieving **a wide range of benefits** for all types of wash loads.

Text 1-3: Sharing text (spoken, dialogic)

'And we've been trying different places around the island that – em, a couple of years ago we got on to this place called the Surai in East Bali and we just go back there now every time. It is –'

'Oh I've heard about this.'

'Have you heard about it? Oh.'

'Friends have been there.'

'It is the most wonderful wonderful place. Fabulous.'

Text (1-3) was a spontaneous spoken text that we are able to transpose into writing because it was recorded on audiotape. Text (1-2) is a written text, which we could (if we wanted to) read aloud. Text (1-1) is more complex: it was probably composed in writing, perhaps with some spoken rehearsal; but it was written in order to be spoken, and to be spoken on an all-important public occasion (Nelson Mandela's inaugural speech as President, 10 May 1994).

When grammarians say that from their point of view all texts are equal, they are thinking of them as specimens. If we are interested in explaining the grammar of English, all these three texts illustrate numerous grammatical features of the language, in meaningful functional contexts, all equally needing to be taken into account. Seen as artefacts, on the other hand, these texts are far from equal. Text (1-1) constituted an important moment in modern human history, and may have left its imprint on the language in a way that only a very few highly valued texts are destined to do. But here too there is a complementarity. Text (1-1) has value because we also understand texts like (1-2) and (1-3); not that we compare them, of course, but that each text gets its meaning by selecting from the same meaning-making resources. What distinguishes any one text is the way these resources are deployed.

Our aim in this book has been to describe and explain the meaning-making resources of modern English, going as far in detail as is possible within one medium-size volume. In deciding what parts of the grammar to cover, and how far to go in discussion of theory, we have had in mind those who want to use their understanding of grammar in analysing and interpreting texts. This in turn means recognizing that the contexts for analysis of discourse are numerous and varied – educational, social, literary, political, legal, clinical and so on; and in all these the text may be being analysed as specimen or as artefact, or both (specimen here might mean specimen of a particular functional variety, or **register**, such as 'legal English'). What is common to all these pursuits is that they should be grounded in an

account of the grammar that is coherent, comprehensive, and richly dimensioned. To say this is no more than to suggest that the **grammatics** – the model of grammar – should be as rich as the grammar itself (Halliday, 1984b, 1996; for educational considerations, cf. also Williams, 2005). If the account seems complex, this is because the grammar is complex – it has to be, to do all the things we make it do for us. It does no service to anyone in the long run if we pretend that **semiosis** – the making and understanding of meaning – is a simpler matter than it really is.[1]

1.1.1 Constituency: (1) phonological

Perhaps the most noticeable dimension of language is its **compositional** structure, known as 'constituency': larger units of language consist of smaller ones. The patterns of any sub-system of language such as the sub-system of sounding, or **phonology**, are distributed across units of varying size, ranging from the largest units of that sub-system to the smallest. Units of different sizes carry different kinds of pattern; for example, in phonology, the largest units carry melodic patterns, and the smallest units carry articulatory patterns.

If we listen to any of these texts – to any text, in fact – in its spoken form we will hear continuous melody with rising and falling pitch, and with certain moments of prominence marked by either relatively rapid pitch changes or extended pitch intervals (cf. Halliday & Greaves, 2008). These moments of prominence define a snatch of melody – a melodic unit, or **line**; and within this melodic progression we will be able to pick up a more or less regular beat, defining some rhythmic unit, or **foot**. We can perhaps recognize that the 'line' and the 'foot' of our traditional verse metres are simply regularized versions of these properties of ordinary speech.

Each foot, in turn, is made up of a number of units of articulatory movement, or **syllables**; and each syllable is composed of two parts, one of which enables it to rhyme. We refer to this rhyming segment, simply, as the **rhyme**; the preceding segment to which it is attached is called the **onset**. Both onset and rhyme can be further analysed as articulatory sequences of consonants and vowels: consonant and vowel **phonemes**, in technical parlance.

The stretch of speech is continuous; we stop and pause for breath from time to time, or hesitate before an uncertain choice of word, but such pauses play no part in the overall construction. None of these units – melodic line (or 'tone group'), foot (or 'rhythm group'), syllable or phoneme – has clearly identifiable boundaries, some definite point in time where it begins and ends. Nevertheless, we can hear the patterns that are being created by the spoken voice. There is a form of order here that we can call **constituency**, whereby larger units are made up out of smaller ones: a line out of feet, a foot out of syllables, a syllable out of sequences of phonemes (perhaps with 'sub-syllable' intermediate between the two). We refer to such a hierarchy of units, related by constituency, as a **rank scale**, and to each step in the hierarchy as one **rank** (cf. Halliday, 1961, 1966c; Huddleston, 1965).

---

[1] Throughout this book we will show the first mention of technical terms such as 'register', 'grammatics' and "semiosis" in bold. Most scientific disciplines use technical terms quite extensively as part of the linguistic resources for construing their field of study. Technical terms are *not* unnecessary 'jargon'; they are an essential part of construction of scientific knowledge. Many of the terms used here can be found in Matthiessen, Teruya & Lam (2010). If this introduction to functional grammar seems to have many technical terms, we recommend a comparison with a university textbook introducing, e.g., anatomy or geology!

What we have been setting up here is the rank scale for the sound system of English: the **phonological rank scale** (see Halliday, 1967a: 12ff.; Halliday & Greaves, 2008). Every language has some rank scale of phonological constituents, but with considerable variation in how the constituency is organized (cf. Halliday, 1992c, on Mandarin): in patterns of articulation (syllables, phonemes), of rhythm (feet), and of melody (tone groups), and in the way the different variables are integrated into a functioning whole. We get a good sense of the way the sounds of English are organized when we analyse children's verses, or 'nursery rhymes'; these have evolved in such a way as to display the patterns in their most regularized form. *Little Miss Muffet* can serve as an example (Figure 1-1).[2]

Fig. 1-1 Example of phonological constituency

We will say more about phonology in Section 1.2 below. Meanwhile we turn to the notion of constituency in writing.

1.1.2 Constituency: (2) graphological

As writing systems evolved, they gradually came to model the constituent hierarchy of spoken language, by developing a rank scale of their own. Thus, in modern English writing, we have a graphological rank scale of four ranks: the **sentence** (beginning with a capital letter and ending with a major punctuation mark: a full stop, question mark or exclamation mark), **sub-sentence** (bounded by some intermediate punctuation mark: colon, semicolon or comma; or a dash), **word** (bounded by spaces) and **letter**. Here is the same text written in orthographic conventional form (see Figure 1-2).

Fig. 1-2 Examples of graphological constituency: sentence, sub-sentence and word

[2] Versions of nursery rhymes are those given in Iona & Peter Opie, *The Oxford dictionary of nursery rhymes*.

The constituent structure is represented by a combination of **spelling** (combining letters to form words) and **punctuation** (using special signs, and also the case of the letter, to signal boundaries; cf. Halliday, 1985a). The system is more complex than we have illustrated here, in three respects: (1) word boundaries are somewhat fuzzy, and there is a special punctuation mark, the hyphen, brought in to allow for the uncertainty, e.g. *frying pan*, *fryingpan*, *frying-pan*; (2) there is a further rank in the hierarchy of sub-sentences, with colon and semicolon representing a unit higher than that marked off by a comma; (3) there is at least one rank above the sentence, namely the paragraph. These do not affect the principle of graphological constituency; but they raise the question of why these further orders of complexity evolved.

The simple answer is: because writing is not the representation of speech sound. While every writing system is related to the sound system of its language in systematic and non-random ways (exactly how the two are related varies from one language to another), the relationship is not a direct one. There is another level of organization in language to which both the sound system and the writing system are related, namely the level of **wording**, or 'lexicogrammar'. (We shall usually refer to this simply as 'grammar', as in the title of the book; but it is important to clarify from the start that grammar and vocabulary are not two separate components of a language – they are just the two ends of a single continuum (see Halliday, 1961; Hasan, 1987; Matthiessen, 1991b; Tucker, 1998, 2007).) The sound system and the writing system are the two modes of **expression** by which the lexicogrammar of a language is represented, or **realized** (to use the technical term).

Since language evolved as speech, in the life of the human species, all writing systems are in origin parasitic on spoken language (cf. Halliday, 1985a; Matthiessen, 2006b); and since language develops as speech, in the life of every hearing individual, this dependency is constantly being re-enacted. Even with the deaf, whose first language uses the visual channel, this is not writing; Sign is more closely analogous to spoken than to written language, signs being in a sense visible forms of articulation and facial expressions visible prosodies. But as writing systems evolve, and as they are mastered and put into practice by the growing child, they take on a life of their own, reaching directly into the wording of the language rather than accessing the wording via the sound; and this effect is reinforced by the functional complementarity between speech and writing. Writing evolved in its own distinct functional contexts of book keeping and administration as 'civilizations' first evolved – it never was just 'speech written down'; and (at least until very recent advances in technology) the two have continued to occupy complementary domains.

So, still keeping for the moment to the notion of constituency, as a way in to exploring how language is organized, let us look at the phenomenon of constituency in lexicogrammar. This will help to explain the principles that lie behind this kind of hierarchic construction, and to understand what is common to different manifestations (such as melodic unit of speech, the line of metric verse, and the sub-sentence of the written text).

1.1.3 Constituency: (3) lexicogrammatical

We will visit Little Miss Muffet just one more time. The punctuation of the text, in the previous section, clearly indicated its graphological composition, in terms of sentences, sub-sentences and words. When we now break down the same text into its grammatical constituents, we find a high degree of correspondence across the higher units: each written

sentence in the graphology is one **clause complex** in the grammar, and each sub-sentence is one **clause**. This is obviously not a coincidence: the two sets of units are related (see Figure 1-3).

Fig. 1-3 Example of grammatical constituency

But they are not identical; the correspondence will not always hold. Little Miss Muffet evolved as a spoken text, so when someone decided to write it down they chose to punctuate it according to the grammar. In Nelson Mandela's text, on the other hand, the first (written) sentence is grammatically a single clause – but it is written as five sub-sentences. Here the punctuation is telling us more about the phonological structure (the division into tone groups) than about the grammar. There is nothing unusual about this: many writers punctuate phonologically rather than grammatically, or in some mixture of the two. And there are many kinds of written text that are carefully punctuated into sentences and sub-sentences (i.e. with full stops, colons and commas) but containing no clauses or clause complexes at all, like the following:

Text 1-4: Recommending – 'Classified rates'

CLASSIFIED RATES

£5.10 per line (average six words per line); display £12 per single column centimetre; box numbers £5.

Discounts: 20 per cent for four insertions, 30 per cent for eight insertions, 50 per cent for twelve insertions.

Prices do not include VAT.

London Review of Books, 28 Little Russell Street, London WC1A 2HN.

It is often uncertain whether someone writing about grammar is talking about graphological units or grammatical units. To avoid this confusion we shall call them by different names (as has become the usual practice in systemic functional grammar). We will use **sentence** and **sub-sentence** to refer only to units of orthography. In referring to grammar we will use the term **clause**. When a number of clauses are linked together grammatically we talk of a **clause complex** (each single linkage within a clause complex can be referred to as one **clause nexus**).

Below the clause, the situation is rather different. Graphologically, sub-sentences consist of words – there is no written unit in between. The word is also a grammatical unit; and

here we shall continue to use the same term for both, because the correspondence is close enough (both categories, orthographic word and grammatical word, are equally fuzzy!). Grammatically, however, the constituent of a clause is not, in fact, a word; it is either a phrase or a word group (which we shall call simply **group** from now on). (We have not shown phrases in Little Miss Muffet; there are two examples, *on a tuffet* and *beside her*. For the important difference between a group and a phrase, see Section 6.1 'Groups and phrases' in, Chapter 6.) Grammatically, a word functions as constituent of a group.

Words have constituents of their own, **morphemes**. These are not marked off in the writing system; sometimes they can be identified as the parts of a written word, e.g. *eat* + *ing*, *curd* + *s*, *frighten* + *ed*, or else recognized as traces of its history (*beside*, *away* were both originally dimorphemic). We shall not be dealing systematically with word morphology in this book (see Matthiessen & Halliday, in prep.); but it illustrates the limits of compositional structure in language (and hence the problems of trying to explain all of grammar in constituency terms). Grammarians used to worry a lot about whether to analyse *sat*, *came* as consisting of two morphemes (*sit/come* plus an abstract morpheme 'past' realized as a vowel change); but this is a problem created by the theory. Composition is an important semogenic (meaning-creating) resource; but it should not be allowed to dominate our thinking about grammar.

Let us summarize here the five principles of constituency in lexicogrammar.

(1) There is a **scale of rank** in the grammar of every language. That of English (which is typical of many[3]) can be represented as:
    clause
    phrase/group
    word
    morpheme.
(2) Each consists of ***one or more*** units of the rank next below. For example, *Come!* is a clause consisting of one group consisting of one word consisting of one morpheme.[4]
(3) Units of every rank may form **complexes**: not only clause complexes but also phrase complexes, group complexes, word complexes and even morpheme complexes may be generated by the same grammatical resources.
(4) There is the potential for **rank shift**, whereby a unit of one rank may be down-ranked (downgraded) to function in the structure of a unit of its own rank or of a

[3] Languages vary, however, with respect to the 'division of grammatical labour' among the ranks. In particular, certain languages do relatively more grammatical work at group (and clause) rank, while other languages do relatively more work at word rank. Thus, for example, Japanese, Turkish, and Inuit do relatively more work at word rank, whereas, for example, Thai, Chinese, and Vietnamese do relatively more work at group rank. For instance, verbal affixes operating at word rank in one language may correspond to verbal auxiliaries operating at group rank in another, or even to modal particles operating at clause rank in yet another. This distribution of grammatical work across the rank scale is likely to change over time as a language evolves: there is a strong tendency for higher-ranking items to drift down the rank scale, as when pronouns and auxiliaries lose their status as free words and gradually become bound verbal affixes.

[4] This is not an arbitrary 'rule'. It is what explains the fact that such an instance is selecting simultaneously in systems of every rank: *Come!* is an 'imperative' (as opposed to 'indicative') clause, a 'positive' (as opposed to 'negative') verbal group, a base (as opposed to derived) form of the verb (word).

rank below. Most commonly, though not uniquely, a clause may be down-ranked to function in the structure of a group.

(5) Under certain circumstances it is possible for one unit to be **enclosed** within another; not as a constituent of it, but simply in such a way as to split the other one into two discrete parts.

To represent the lexicogrammatical constituents in a passage of written text we adopt the notational conventions set out in Table 1-1 Notational conventions for representing lexicogrammatical constituency (see Table 1-1).

Table 1-1 Notational conventions for representing lexicogrammatical constituency

Examples:

|| out of [the experience [of [an extraordinary human disaster [[that | lasted | too long]]]] | must be born | a society [[of which | all humanity | will be | proud]] ||

|| did <you> read | that article [the other day] [about [this woman [[[who | was driving | along | somewhere | on [this country road] || when | hail | just suddenly | started pouring down]]]] ||

||| we | understand || that | there | is | more [than [one thing]] [[you | want to achieve | out of [every wash load]]] |||

|| today | all of us | do < by [our presence here] and | by [our celebrations [in [other parts [of [our country and the world [ ]]]] > confer | glory and hope | to [newborn liberty] ||

The clause is the central processing unit in the lexicogrammar – in the specific sense that it is in the clause that meanings of different kinds are mapped into an integrated grammatical structure. For this reason the first half of this book is organized around the principal systems of the clause: theme, mood and transitivity. In Part II we move outward from the clause, to take account of what happens above and below it – systems of the clause complex, of groups and phrases, and of group and phrase complexes; and also beyond the clause, along other dimensions so to speak.

The perspective moves away from structure to consideration of grammar as system, enabling us to show the grammar as a meaning-making resource and to describe grammatical categories by reference to what they mean. This perspective is essential if the analysis of grammar is to be an insightful mode of entry to the study of discourse. But first, in the remainder of the present chapter, we will say a little more about compositional structure, including a more detailed sketch of phonology, so that we can take the relevant aspects of it for granted throughout the rest of the book.

1.2 Phonology and grammar

If we want to take a comprehensive view of English grammar, we must first make an excursion into phonology. This is because there are some grammatical systems that are realized by prosodic means: for example, by the contrast between falling and rising tone.

As we have seen in Section 1.1.1, the units of phonology are organized from largest to smallest according to the phonological rank scale – tone group (melodic line), foot (rhythm group), syllable and phoneme. Each unit is the domain of certain phonological systems, and it can be characterized in terms of a characteristic structure (the exception being the smallest unit, the phoneme): see the summary in Table 1-2. These units can be divided into two regions of articulation and prosody. Articulatory features are associated with smaller units, typically phonemes (vowels and consonants). Prosodic features are associated with larger units; they are features of intonation and rhythm (for an overview of the phonetics of prosody, see Nooteboom, 1997). The gateway between the two regions is the syllable; it realizes prosodic features of intonation and rhythm (and may carry its own, e.g. syllabic tone in 'tone languages' such as Chinese) and it 'choreographs' articulatory gestures (sequences of phonemes).

Table 1-2 The phonological rank scale

As a general principle, articulation is **arbitrary** (conventional), in the sense that there is no systematic relation between sound and meaning (as emphasized by Saussure and further developed by other European structuralist linguists, in particular Hjelmslev and Martinet; see e.g. Halliday, 1985b/2003: 196). Prosody, on the other hand, is **natural** (just as grammar is in relation to semantics; see Halliday & Matthiessen, 1999: 18–22): it is related systematically to meaning, as one of the resources for carrying contrasts in grammar. In this section we give a sketch of the prosodic region of the phonology from the standpoint of its importance to the grammatical description. For a more comprehensive account, see Matthiessen & Halliday (in prep.), Halliday (1967a) and Halliday & Greaves (2008).

1.2.1 Rhythm: the foot

Consider another well-known piece of traditional children's literature:

Text 1-5: Expounding – nursery rhyme

If all the world was apple pie,

And all the sea was ink,

And all the trees were bread and cheese,

What should we have to drink?

Say it aloud; or, better still, get a pre-literate child to say it aloud for you, so as to avoid imposing any artificial conventions in the reading of written verse. You will hear the lines of melody, and you will hear the rhythm.

The rhythm is carried by a succession of beats, occurring at more or less regular intervals (Abercrombie, 1967: 96–98; Catford, 1977: 85–91, 1985).[5] In this verse, the beats occur in alternate syllables, which happen to be the even-numbered syllables in the line: *all*, *world*, *ap-*, *pie*, *sea*, *ink*; *all*, *trees*, *bread*, *cheese*; *should*, *have*, *drink*. For contrast, here is a verse with the beat on the odd-numbered syllables in the line:

Text 1-6: Recreating – nursery rhyme

Better Botter bought some butter.

But, she said, the butter's bitter;

If I put it in my batter,

It will make the batter bitter ...

Accompanying the beat syllable are other, off-beat syllables that are rhythmically dependent on it. In these examples there was only one off-beat syllable attached to each one that carried a beat. There may be two, as in Little Miss Muffet (cf. Figure 1-1 above). We will refer to the syllable that carries the beat as a **strong** syllable (but noting at the same time that there is also a technical term, **salient**, for it) and to the off-beat syllables as **weak**. The structural unit formed by one strong, or salient, syllable together with any weak syllable(s) following on from it is called a **foot**. The foot is one of the constituents of the English sound system; it is the unit of rhythm.

The foot is easily recognized in children's nursery rhymes, (1) because the strong syllables are very regular in tempo and (2) because a pattern is set up with a fixed number of syllables in the foot. This second point does not mean that every foot contains exactly that number; there may be a moment of silence, like a rest in music, or one syllable may be lengthened to stretch over the time allowed for two, like *curds* in *eating her curds and whey*. But each verse establishes is own basic pattern, either of two or of three syllables in the foot, or sometimes of four; and every foot adapts to that pattern.

---

[5] Note that the account of rhythm presented here is based on the study of natural, connected speech. This contrasts sharply with accounts of rhythm that are based on isolated words and expressions (as in metrical phonology) or constructed examples that are read aloud. There is, however, a trace of the connection between the word (as a unit of grammar) and the foot (as unit of phonology) in the form of **accent**: a grammatical word is realized phonologically by a sequence of syllables, and one of them will be the default location of the beat. But this may be overridden in connected speech; and there are words where the beat may fall on either syllable (as with *Chinese*) to accommodate the rhythmic pattern of connected speech.

The question then is: how does this relate to the natural rhythm of English speech? All poetry ultimately derives from natural spoken language; over time it evolves a rich array of patterns of its own, at all levels of language, but all of them have their origin in speech. Every language has its own natural rhythm, some patterned way of modulating the pulse of the air-stream that comes from the diaphragm (cf. Catford, 1977: 85–91). In English, the rhythm of speech derives from the marked contrast between strong and weak syllables (cf. Grabe & Low, 2002). When you speak, naturally and spontaneously, without paying attention to the process of speaking, the strong syllables tend to occur at roughly even intervals: nothing like so exact, of course, as in children's verses, or in recitations like counting or listing the days of the week, but enough to provide a clear measure, a rhythmic progression with which the listener keeps in phase.

This rhythmic progression represents a form of constituency: the foot is a constituent of the sound system of English. The foot in poetry has its origin in the foot of the spoken language. But there are three factors that need to be taken into account when we compare the two.

(1) In natural speech, the number of syllables in the foot continually varies; there may be just one (the salient syllable), or there may be two, three, four, or even five or six in speech with rapid tempo. This was, in fact, the pattern followed in Old English and Early Middle English verse; the line had a definite number of feet (typically four), but the number of syllables in the foot could vary freely. The metric foot – that is, a foot with a fixed number of syllables – became established in Chaucer's time, largely through the influence of Chaucer himself, and it remained the norm of mainstream English verse for the next five centuries. In the twentieth century it ceased to dominate, and there has been a new wave of input into poetry from spoken language – including, in the past few decades, from speakers of new varieties of English whose rhythms are very different from those of the original native speakers of the language.

(2) Part of the tradition of metric verse was the analysis of verse forms in terms of **metrics**: this was an analysis based on the number of feet per line, and the number and distribution of syllables in the foot. A line might have two, three, four, five or six (occasionally seven or eight) feet; the favourite line, that of Chaucer, Shakespeare, Milton, Pope and Keats, was the pentameter (five feet). A foot might have two, three or four syllables, but, in addition, it might be either 'descending' or 'ascending' – that is, the salient syllable might occur either at the beginning or at the end. For example, a two-syllable foot might be **trochaic** (strong + weak) or **iambic** (weak + strong). Of the verses cited above, *Betty Botter* is trochaic, while *the world of apple pie and ink* is iambic.

This last distinction, between descending and ascending rhythm, is a property of metric verse form, accounting for how the line is organized into metric feet; it has no significance for the sound system of English. In spoken English the salient syllable always occurs at the beginning of the foot; a foot is thus like a bar in music, defined as beginning with a beat. The phonological foot, therefore, as distinct from the metric foot, consists of one strong syllable optionally followed by one or more weak syllables. The functional interpretation of this structure is

Ictus (^ Remiss)

where ^ means 'followed by' and the parentheses indicate that the Remiss element is optional. The Ictus and the Remiss are **elements** of the functional structure of the foot, and

they are realized by units of the rank below that of the foot, i.e. by syllables. The Ictus is realized by a strong syllable (or a silent beat; see immediately below), and the Remiss (if present) by one or more weak syllables.

(3) The tradition of metric analysis was faulty in certain respects, particularly in confusing the opposition of 'strong/weak' with the quite different phenomenon of 'long/short'. But its main defect in relation to our present discussion is that it failed to recognize silence. Silence is a systematic feature of the rhythm of spoken English; there are many instances of what is usually called a **silent beat**, where the Ictus is clearly present in the sound pattern (i.e. there ***is*** a beat), but it is realized in the form of silence – just as a bar might begin with a silent beat in music. So we may have an entirely silent foot, and many of the standard metres of English verse depend on this; there is in fact a silent foot at the end of the second and fourth lines of *If all the world was apple pie*, as you can tell by beating out the time while saying it. In spontaneous dialogue, speakers and listeners can maintain the tempo across at least two feet of total silence; and the silent beat also plays a part in grammar, in making a contrast in meanings (see Chapter 7, Section 7.4.1.2).

Below the foot on the phonological rank scale (see Section 1.1.1) is the **syllable**; a foot consists of one or more syllables (unless the foot consists of only an Ictus realized by a silent beat). The syllable is the fundamental unit of an articulatory gesture (see, e.g., Catford, 1977: 88–91; Fujimura & Erickson, 1997: 98–99; and in auditory terms it can be characterized as being organized around a peak of high sonority). Like the foot, the syllable is also a structured unit; it consists of Onset + Rhyme. These are, in turn, realized by consonant and vowel phonemes. The Onset is realized as initial consonant or consonant cluster, or zero. The Rhyme is, naturally, the part of the syllable that is involved in rhyme: vowel plus following consonant or consonant cluster if any. We shall say a little more about the organization of syllables later, in Section 1.2.3; but here it is important to note that the realization of the phonemes serving as Onset + Rhyme will depend on the rhythmic nature of connected speech. The syllable is 'elastic' so that it can accommodate itself to the rhythmic requirements of the foot. This is the property of the syllable that enables the foot to function as a rhythmic unit: given a constant tempo, the more syllables there are in a foot, the shorter those syllables have to be.

1.2.2 Intonation: the tone group

The foot, then, is a 'rhythm group': it is the unit that manages the rhythm of spoken English. It is a constituent in the phonological structure. But it is not the highest constituent; there is a more extended sound pattern constructed out of a succession of feet. If you listen again to the apple pie quatrain you will hear a clear melodic pattern emerging, probably corresponding fairly closely to the line: one line of verse, one 'snatch' or line of melody. The name for this systematic melodic variation in language is **intonation**, the melodic line is an **intonation contour**, or more shortly a **tone contour**; and the snatch is called a **tone group** (Halliday, 1967a; Halliday & Greaves, 2008; Elmenoufy, 1988; Tench, 1990, 1996; Wells, 2006; other names used are 'tone unit' and 'intonation unit').

If we are given a text in writing, there will always be various possible ways of intoning it, each with a somewhat different meaning (cf. Davies, 1986); but generally one or a small number of these possible intonation patterns will stand out as more natural and more likely. If we say the apple-pie verse, we probably start the first three lines with a fairly high-pitched note on *all*, descend step by step, and then end with a slight but noticeable rising

pitch on the last syllable in the line: *pie*, *ink*, *cheese*. (Note that the melody is constructed by the strong syllables; the weak ones fit in with them in the place of least effort – here the *if* and the *and* at the beginning of the line would be on a neutral, medium pitch.) The third line might also have a little rise on *trees* to go with the internal rhyme. The last line, on the other hand, would probably have a highish level pitch on *should*, a similar but marginally lower one on *have*, and a clearly marked falling pitch on the final word *drink*, with a total effect of an overall movement from high to low.

The tone group is the phonological constituent that ranks above the foot: each tone group consists of a whole number of feet – one or more. This relationship whereby syllables are grouped into rhythm groups, or feet, and feet into tone groups, is the way the sounds of English are organized into larger patterns. But unlike writing, which is captured (even if very briefly) in time, so that written units can be clearly marked off one from another, speech is fluid and kinetic: there are no clear boundaries between its constituents. So in a given passage of speech we can tell how many syllables there are, how many feet, and how many tone groups; and we can tell within limits where each one is located; but we cannot pinpoint exactly where each one begins and ends. So we determine the boundaries on theoretical grounds, making generalizations that have the greatest explanatory force. One example of such a generalization is the principle just stated, whereby each unit consists of a whole number of units of the rank next below; this means that a tone group boundary will always also be a foot boundary, and a foot boundary always also a syllable boundary, and this in turn makes it easier – makes it possible, in fact – to explain how these phonological patterns function in the making of meaning (cf. Figure 1-1 above).

While the rhythm group, or foot, is largely a timing unit (it has one or two specific functions in the grammar, but its domain of operation is principally phonological), the tone group does a great deal of work in the construal of meaning: it organizes continuous speech as a sequence of units of information (see further Chapter 3, Section 3.5). In other words, it manages the flow of the discourse. It is thanks to the tone group, as defined here, that listeners get the message: not only how it is framed into items of news but also what value the speaker attaches to each. Hence it is on the basis of the tone group that, as grammarians, we are able to analyse spoken language and 'chunk it up' into grammatical units of a particular and important kind.

The tone group is the point of entry into the system of TONE: the systematic use of melody as a grammatical resource. While the potential (and actual) variation in the pitch movement in speech is immense, each snatch in the melody represents one of a small number of systemically distinct tones. At the phonological level there are just five of these: falling, rising, 'level', falling-rising, and rising-falling. These are realized in the form of a characteristic pitch movement located on one particular foot within the tone group, the **tonic** foot. (The 'level' tone is actually realized as a low with a slight rise.) Whereas the organization of speech into a sequence of tone groups carries **textual** meaning, the choice among the different tone carries **interpersonal** meaning, via the grammatical systems of KEY (see further Section 4.4.4). For the metafunctional categories of interpersonal and textual, see below Section 1.3.5.

The tone group itself has an internal structure of the form (P ^) T: an obligatory Tonic segment, optionally preceded by a Pretonic segment. Each of these consists of at least one complete foot; the tonic foot is the first foot within the Tonic segment.

Let us return briefly to the relationship among spoken language, written prose, and written verse. The 'line', in verse, evolved as the metric analogue of the tone group: one line of verse corresponded to one tone group of natural speech. In children's nursery rhymes this correspondence is often preserved intact (this is why they are valuable in helping children learn the patterns of the language); but in adult verse of course it is not – on the contrary, it becomes an idealized motif on which endless meaningful variations can be played. And this illustrates a principal strategy whereby the meaning potential of a language comes to be extended.

We can postulate an 'initial' state where the two **variables** are fully **associated**: at this point, the 'line' is just the poetic incarnation of the tone group. Then the two become **dissociated**: the poetic line takes on a life of its own, and new meanings are construed by mapping a line into more than one tone group, or having the intonation patterns cut right across those set up by the poetic form. This 'dissociation of associated variables' is one of the main semogenic resources of a language; for a theoretical discussion, see Halliday (1991).

As it happens, the verse form helps us recognize that there is a still higher unit of organization above the tone group. Listening to the four lines of the apple-pie rhyme we can hear that they make up a sequence of interrelated tone groups: beginning with a series that are alike, all ending on a rise, and ending with one that is distinct, with its final falling movement on *drink*. The rising tones suggest non-finality, whereas the fall sounds (and in fact is) culminative: it brings the sequence to an end. Together these tone groups make up a **tone group complex** (see further Chapter 7, Section 7.6); and this, in turn, is the origin of the metric **stanza**, as a higher pattern of organization in poetry.

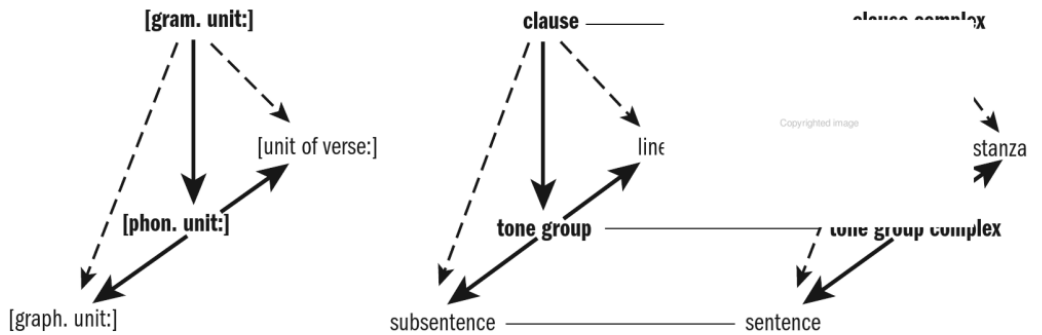

Fig. 1-4 Analogic patterning of units across content (lexicogrammar) and expression (phonology, graphology)

The explanation for all this analogic patterning lies in the grammar. When children start to speak English as their mother tongue, they soon learn to construct a unit that is a conflation of clause and tone group (later, again, they will learn to dissociate these two). This is also manifested as a unit in the rhymes they hear. When they start learning to read and write they find this unit reappearing as a (simple) sentence, with a capital letter at the beginning and a full stop or other major punctuation mark at the end; and this also turns

up as a line in written verse. Behind all these diverse entities – tone group, spoken line, sentence in writing and written line – lies the one fundamental grammatical unit, the clause.

By the same token, an analogous relationship is set up among a series of higher units: tone group complex, spoken stanza, 'compound/complex' sentence, and stanza of written verse, all of which originate as different incarnations of the grammatical **clause complex** (see Chapter 7, *passim*). And when the rhymes have been set to music, and are sung, the same patterns are reinforced over again, with the line of melody representing the clause, and the melody as a whole representing the clause complex. We could diagram these two sets of analogies as in Figure 1-4.

Then, in the very act of developing these fundamental unities, the child is also learning to pull them apart: to deconstrue the pattern, so that each of its modes of being becomes a carrier of meaning in its own right. Once a clause, for example, may be mapped either into one tone group or into two, this enhances its meaning potential in the flow of discourse; moreover, there are likely to be various places where the transition can take place. The phonological patterns (and, for a literate person, the graphological ones also) are the semogenic resources of a language; any systemic variation that they embody has the potential for making systematic distinctions in meaning – and most of these are likely to be taken up.

Here are the notational conventions for the higher units of phonology – the tone group complex, the tone group and the foot (rhythm group): see Table 1-3. Examples:

(a) /// ^ if / all the / world was / apple / pie and // all the / sea was / ink and // all the / trees were / bread and / cheese what // would we / have to / drink ///

(b) /// ^ and we've been / trying / different / places a/round the / island that / ^ em // ^ a / couple of / years a/go we // got on to this / place / called the / Surai in // east / Bali and we // just go / back there / now / every / time // ^ it / is ... ///

// oh I've / heard about / this //

/// have you / heard about it // oh ///

// friends have / been there //

/// ^ it / is the most / wonderful / wonderful / place // fabulous ///

Table 1-3 Notational conventions for higher phonological units

Figure 1-5 shows the system network for prosodic systems in English phonology. Note that the network represents the phonological resources; it does not show how these resources are exploited in the lexicogrammar. Some illustrations of this will be found in Chapters 3, 4 and 7.

Fig. 1-5 System network for prosodic systems in English phonology

1.2.3 Syllables and phonemes

Next below the foot is the **syllable**; as already remarked, a foot consists of a whole number of syllables, one or more. From an articulatory point of view, the syllable is an articulatory gesture; Catford (1977: 89) characterizes is as a '"ripple" on the surface of the initiator-power curve' – that is, the articulatory correlate of a foot. From an auditory point of view, this gesture is typically organized around a peak of sonority. All languages have something that can be called a syllable; but these somethings are far from being the same – if we compare just Russian, Japanese, Arabic and English we find great variation in how syllables are structured and how they function (for a systemic account of syllables in Mandarin, see Halliday, 1992c). In some languages, it is clear where a syllable begins and ends; but in others it may not be quite so clear – in English it is not at all obvious how to divide up a word such as *colour* or *basket* into syllables, and people dispute whether words like *chasm*, *rhythm*, *fathom* consist of one syllable or two (I once watched a game of charades dissolve into chaos as the players argued whether *comfortable* had three syllables in it or four). But the fact that English verse came to depend on counting syllables means that syllables must have been perceived as things that were able to be counted, even if there is indeterminacy in certain places. Musical settings of verse also impose a syllable pattern – which is not always the same as that required by the metre.

What is there below the syllable? English verse makes extensive use of rhyme; from that point of view, a syllable consists of two parts, the non-rhyming part, or **Onset** (which may be empty), and the **Rhyme**. This analysis is helpful in explaining the relative duration of different syllables in English, since this depends entirely on the structure of the rhyme. On the other hand, the English writing system is made up of letters, and the letters stand for smaller units of sound called **phonemes** – the individual consonants and vowels out of which both parts of the syllable are built.

The English script is not 'phonemic' if by that we understand a strict one-to-one correspondence between phonemes and letters. It never could be phonemic in this sense, because the criteria for identifying phonemes in English are internally contradictory: what are one and the same phoneme from one point of view may be two separate phonemes from another. But it clearly is phonemic in its general principle: the symbols represent consonants and vowels that contrast systemically with one each other and combine to form regular structures. Many of its symbols have more than one phonemic value; some pairs of letters ('digraphs') have to be treated as single symbols, like *th* in *thin*, *sh* in *shin*; and there are various other departures from an imaginary phonemic ideal – some of them systematic, some random. Nevertheless, speakers of English readily become aware of the phoneme as a minimal phonological unit; the fact that there is no one right answer to the question 'How many phonemes are there in English?', and there is indeterminacy where some of them begin and end (is the sound *ch* in *chin* one phoneme or two?), merely brings them into line with all the other constituents in the phonological system – syllables, feet and tone groups – and, it might be added, with most other phenomena pertaining to natural languages.

In this book we shall not need to be concerned with the detailed analysis of syllables and phonemes. For discussion of the grammar, the important part of phonology is prosody – features of intonation and rhythm. The transcription that will be needed is one that shows the intonational and rhythmic features of speech but which uses ordinary orthography for the spelling – an elaboration of the conventions introduced in the previous section.

1.3 Basic concepts for the study of language

The discussion so far has raised a number of theoretical issues, as can be seen from the variety of technical terms that have had to be used. We have referred to language (i) as text and as system, (ii) as sound, as writing and as wording, (iii) as structure – configurations of parts, and (iv) as resource – choices among alternatives. These are some of the different guises in which a language presents itself when we start to explore its grammar in functional terms: that is, from the standpoint of how it creates and expresses meaning.

At this point, we begin to need a map: some overview of language that will enable us to locate exactly where we are at any point along the route. A characteristic of the approach we are adopting here, that of systemic theory, is that it is *comprehensive*: it is concerned with language in its entirety, so that whatever is said about one aspect is to be understood always with reference to the total picture. At the same time, of course, what is being said about any one aspect also *contributes to* the total picture; but in that respect as well it is important to recognize where everything fits in. There are many reasons for adopting this systemic perspective; one is that languages evolve – they are not designed, and evolved systems cannot be explained simply as the sum of their parts. Our traditional compositional thinking about language needs to be, if not replaced by, at least complemented by a 'systems' thinking whereby we seek to understand the nature and the dynamic of a semiotic system as a whole (cf. Matthiessen & Halliday, in prep., Chapter 1, and references therein to Capra, 1996, and other proponents of systems thinking; Matthiessen, 2007a).

In the remainder of this chapter we shall present in a very summary way the critical dimensions of the kind of semiotic that language is. By 'language' we mean natural, human, adult, verbal language – natural as opposed to designed semiotics like mathematics and computer languages (cf. Halliday & Matthiessen, 1999: 29–46; O'Halloran, 2005); adult (i.e. post-infancy) as opposed to infant protolanguages (see Halliday, 1975, 2003); verbal as opposed to music, dance and other languages of art (cf. Kress & van Leeuwen, 1996; O'Toole, 1994; van Leeuwen, 1999). Of course, all these other systems share certain features with language in this specified sense; but none of them incorporates all. The dimensions, or forms of order, in a language, and the ordering principles, are set out in Table 1-4 and represented diagrammatically in Figure 1-6.

Table 1-4 The dimensions (forms of order) in language and their ordering principles

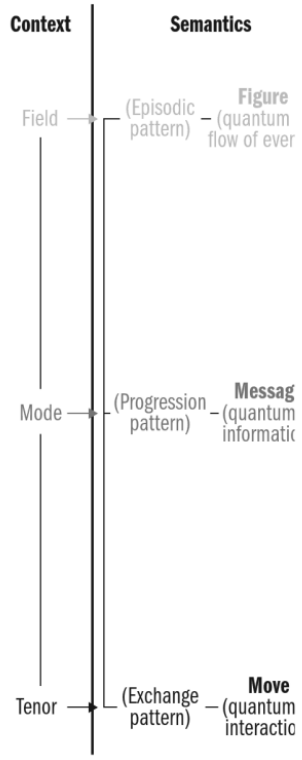

Fig. 1-6 The dimensions in language

1.3.1 Structure (syntagmatic order)

This is the compositional aspect of language, referred to in linguistic terminology as 'constituency'. The ordering principle, as defined in systemic theory, is that of **rank**: compositional layers, rather few in number, organized by the relationship of 'is a part of'. We have identified four such compositional hierarchies in English, as shown in Table 1-5.

Table 1-5 Compositional hierarchies in English

The guiding principle is that of **exhaustiveness**: thus, in the writing system, a word consists of a whole number of letters, a sub-sentence of a whole number of words, a sentence of

a whole number of sub-sentences; the number may be more than one, or just one. At the same time, as always in language, there is much indeterminacy, or room for manoeuvre: should we recognize just one layer of sub-sentences, marked off by any punctuation mark, or two – a higher layer marked off by (semi)colons, a lower one marked off by commas? This may well depend on the practice of the particular writer.

As we have seen, all these compositional hierarchies are ultimately variants of a single motif: the organization of meaning in the grammar. As the language has evolved, they have drifted apart (as will tend to happen in the history of every language); but traces of their equivalence remain (e.g. tone group : sub-sentence : line : clause). When we come to analyse the grammar, we find that the structure of each unit is an **organic configuration** such that each part has a distinctive function with respect to the whole; and that some units may form **complexes**, iterative sequences working together as a single part. Grammar is the central processing unit of language, the powerhouse where meanings are created; it is natural that the systems of sound and of writing through which these meanings are expressed should reflect the structural arrangement of the grammar. They cannot, obviously, copy the functional configurations; but they do maintain the grammatical principle that units of different rank construe patterns of different kinds. In English phonology, for example, the foot is the unit of rhythm; it is the constituent that regulates the pulse of continuous speech. In this it is distinct from other units both above it and below it: from the syllable, which organizes the articulatory sequences of vowels and consonants, and from the tone group, which organizes the pitch movement into patterns of intonation. This functional specialization among units of different rank is a feature of the structure of language as a whole.

1.3.2 System (paradigmatic order)

Structure is the syntagmatic ordering in language: patterns, or regularities, in what *goes together with* what. System, by contrast, is ordering on the other axis: patterns in what *could go instead of* what. This is the paradigmatic ordering in language (cf. Halliday, 1966a; Fawcett, 1988; Butt & Matthiessen, forthcoming).

Any set of alternatives, together with its condition of entry, constitutes a **system** in this technical sense. An example would be 'all clauses are either positive or negative', or more fully 'all clauses select in the system of POLARITY whose terms are positive and negative'; diagrammatically as in Figure 1-7. To get a more rounded picture, we attach probabilities to the two terms: 'positive, 0.9; negative, 0.1' (cf. Halliday & James, 1993).

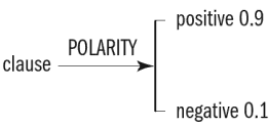

Fig. 1-7 The system of POLARITY

It will be clear that this is a more abstract representation than that of structure, since it does not depend on how the categories are expressed. Positive and negative are contrasting features of the clause, which could be made manifest in many different ways. They represent

an aspect of the **meaning potential** of the language, and they are mutually defining: 'not positive' means the same thing as 'negative', and 'not negative' means the same thing as 'positive'.

The relationship on which the system is based is 'is a kind of': a clause having the feature 'positive' is a kind of clause. Suppose we now take a further step, and say that negative clauses may be either generalized negative, like *they didn't know*, or some specific kind of negative like *they never knew* or *nobody knew*. Here we have recognized two paradigmatic contrasts, one being more refined than the other: see Figure 1-8. The relationship between these two systems is one of **delicacy**: the second one is 'more delicate than' the first. Delicacy in the system ('is a kind of a kind of ...') is the analogue of rank in the structure ('is a part of a part of ...').

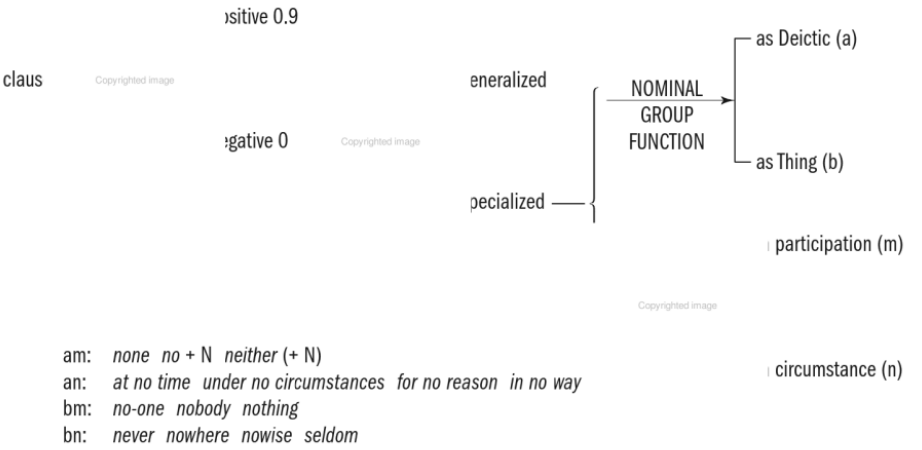

Fig. 1-8 The system of POLARITY, next step in delicacy

A text is the product of ongoing selection in a very large network of systems – a **system network**. Systemic theory gets its name from the fact that the grammar of a language is represented in the form of system networks, not as an inventory of structures. Of course, structure is an essential part of the description; but it is interpreted as the outward form taken by systemic choices, not as the defining characteristic of language. A language is a resource for making meaning, and meaning resides in systemic patterns of choice.

The way system and structure go together can be illustrated by showing a simplified version of the system network for MOOD (this will be explained in detail in Chapter 4): see Figure 1-9. This can be read as follows. A clause is either major or minor in STATUS; if major, it has a Predicator in its structure. A major clause is either indicative or imperative in MOOD; if indicative, it has a Finite (operator) and a Subject. An indicative clause is either declarative or interrogative (still in MOOD); if declarative, the Subject comes before the Finite. An interrogative clause is either yes/no type or WH-type; if yes/no type, the Finite comes before the Subject; if WH-type, it has a Wh element.

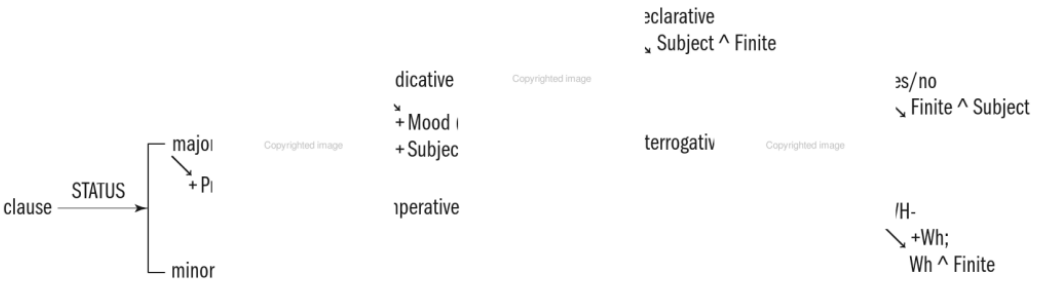

Fig. 1-9 The MOOD system network

What this means is that each system – each moment of choice – contributes to the formation of the structure. Of course, there is no suggestion here of **conscious** choice; the 'moments' are analytic steps in the grammar's construal of meaning (for the relationship between semantic choice and what goes on in the brain see Lamb, 1999). Structural operations – inserting elements, ordering elements and so on – are explained as **realizing** systemic choices. So, when we analyse a text, we show the functional organization of its structure; and we show what meaningful choices have been made, each one seen in the context of what might have been meant but was not.

When we speak of structural features as 'realizing' systemic choices, this is one manifestation of a general relationship that pervades every quarter of language. Realization derives from the fact that a language is a stratified system.

1.3.3 Stratification

We are accustomed to talking about language under different headings. School grammar books used to have chapters on pronunciation, orthography, morphology (earlier 'accidence') and syntax, with a vocabulary added at the end. This acknowledged the fact that a language is a complex semiotic system, having various levels, or **strata**. We have made the same assumption here, referring to the sound system, the writing system and the wording system, i.e. **phonology**, **orthography** (or **graphology**) and **grammar**. (We also noted, on the other hand, that grammar and vocabulary are not different strata; they are the two poles of a single continuum, properly called **lexicogrammar** (cf. Hasan, 1987). Likewise, syntax and morphology are not different strata; they are both part of grammar – the distinction evolved because in Indo-European languages the structure of words (**morphology**) tends to be strikingly different from the structure of clauses (**syntax**); but this is not a feature of languages in general.)

What does it mean to say that these are different 'strata'? In infants' protolanguage, which has as yet no grammar in it, the elements are simple signs; for example, a meaning 'give me that!' is expressed directly by a sound, like *nananana*, or maybe by a gesture of some kind. Here we have just two strata, a stratum of content and a stratum of expression (cf. Halliday, 1975, 2004).

Adult languages are more complex. For one thing, they may have two alternative modes of expression, one of sounding (i.e. speech) and one of writing. More significantly, however, they have more strata in them.

The 'content' expands into two, a **lexicogrammar** and a **semantics** (cf. Halliday, 1984a; Halliday & Matthiessen, 1999). This is what allows the meaning potential of a language to expand, more or less indefinitely. The reason for this can best be explained in terms of the functions that language serves in human lives.

We use language to make sense of our experience, and to carry out our interactions with other people. This means that the grammar has to interface with what goes on outside language: with the happenings and conditions of the world, and with the social processes we engage in. But at the same time it has to organize the construal of experience, and the enactment of social processes, so that they can be transformed into wording. The way it does this is by splitting the task into two. In step one, the interfacing part, experience and interpersonal relationships are transformed into meaning; this is the stratum of semantics. In step two, the meaning is further transformed into wording; this is the stratum of lexicogrammar. This is, of course, expressing it from the point of view of a speaker, or writer; for a listener, or reader, the steps are the other way round.

This stratification of the content plane had immense significance in the evolution of the human species – it is not an exaggeration to say that it turned *homo ...* into *homo sapiens* (cf. Halliday, 1995b; Matthiessen, 2004a). It opened up the power of language and in so doing created the modern human brain. Some sense of its consequences for the construction of knowledge will be given in Chapter 10, where we raise the question of whether learned forms of discourse, in education, science, technology and the humanities, could ever have evolved without the 'decoupling' of these two aspects of the semogenic process.

It might be asked whether an analogous stratification took place within the expression plane; and the answer would appear to be 'yes, it did', and for analogous reasons, namely separating the organizing function from the function of interfacing with the environment. Here, however, the environment is the human body, the biological resource with which sounding (or signing) is carried out. Taking sound (spoken language) as the base, the stratification is into **phonetics**, the interfacing with the body's resources for speech and for hearing, and **phonology**, the organization of speech sound into formal structures and systems (see Figure 1-10).

When we say that language is stratified in this way, we mean that this is how we have to model language if we want to explain it. A language is a series of redundancies by which we link our eco-social environment to non-random disturbances in the air (soundwaves). Each step is, of course, masterminded by the brain. The relationship among the strata – the process of linking one level of organization with another – is called **realization**.[6] Table 1-6 presents this model from the point of view of the speaker – it is hard to present it in a way that is neutral between speaking and listening. Figure 1-10 represents the stratal organization of language, and shows how the stratified linguistic system is 'embedded' in context (cf. Halliday, 1978; Halliday & Hasan, 1985; Hasan, 1999, and other contributions to Ghadessy, 1999; Martin, 1992).

---

[6] With a primary semiotic system, like the infant protolanguage (see immediately below), consisting only of content and expression, we could still use the word 'express'. But with a higher order (multi-stratal) semiotic this is no longer appropriate; we could not really say that wording 'expresses' meaning. Hence the use of a distinct technical term.

Fig. 1-10 Stratification

Table 1-6 From eco-social environment to soundwaves: speaker perspective

Language is thus organized into four strata – semantics, lexicogrammar, phonology, and phonetics. But these four strata are grouped into two stratal planes, the content plane and the expression plane. When children learn how to mean, they start with a very simple semiotic system, a **protolanguage**, usually sometime in the second half of their first year of life (see Halliday, 1973, 2003); and we hypothesize that language evolved in the same way (see Matthiessen, 2004a). This system is organized into two stratal planes, content and expression; but neither is internally stratified: content is mapped directly onto expression (vocal or gestural). This protolanguage is a child tongue rather than a mother tongue; it is not yet like the adult language spoken around young children. Children develop their protolanguages in interaction with their immediate caregivers, gradually expanding their protolinguistic meaning potentials. In doing so, they learn the principles of meaning. At some point, typically in the second year of life, they are ready to build on this experience and to begin to make the transition into the mother tongue spoken around them. This transition involves a number of fundamental changes in the linguistic system. A key

change – one that makes possible other changes – is the splitting up of each of the two stratal planes into two content strata and two expression strata. Content gradually splits into semantics and lexicogrammar, and expression gradually splits into phonology and phonetics. The realizational relationship between content and expression, more specifically between lexicogrammar and phonology is largely **conventional**, or 'arbitrary' (with certain interesting exceptions relating to prosody and to two areas of articulation, phonaesthesia and onomatopoeia). However, the realizational relationship between the two sets of content strata (semantics and lexicogrammar) and the two sets of expression strata (phonology and phonetics) is **natural** rather than conventional. Patterns of wording reflect patterns of meaning. Part of the task of a functional theory of grammar is to bring out this natural relationship between wording and meaning. The natural relationship between semantics and lexicogrammar becomes more complex and less transparent with the development of lexicogrammatical metaphor, as we shall see in Chapter 10; but the relationship is still fundamentally natural rather than arbitrary.

1.3.4 Instantiation

When we want to explain how language is organized, and how its organization relates to the function it fulfils in human life, we often find it difficult to make things clear; and this is because we are trying to maintain two perspectives at once. One perspective is that of language as system; the other perspective is that of language as text.

The concept we need here is that of **instantiation**. The **system** of a language is 'instantiated' in the form of **text**. A text may be a trivial service encounter, like ordering coffee, or it may be a momentous event in human history, like Nelson Mandela's inaugural speech; in either case, and whatever its intrinsic value, it is an instance of an underlying system, and has no meaningful existence except as such. A text in English has no semiotic standing other than by reference to the system of English (which is why it has no meaning for you if you do not know the language).

The **system** is the underlying potential of a language: its potential as a meaning-making resource.[7] This does not mean that it exists as an independent phenomenon: there are not two separate objects, language as system and language as a set of texts. The relationship between the two is analogous to that between the weather and the climate (cf. Halliday, 1992a). Climate and weather are not two different phenomena; rather, they are the same phenomenon seen from different standpoints of the observer. What we call 'climate' is weather seen from a greater depth of time – it is what is instantiated in the form of weather. The weather is the text: it is what goes on around us all the time, impacting on, and sometimes disturbing, our daily lives. The climate is the system, the potential that underlies these variable effects.

Why then do we refer to them as different things? We can see why, if we consider some recent arguments about global warming, the question is asked: 'Is this a long-term weather pattern, or is it a blip in the climate?' What this means is, can we explain global warming

[7] This use of 'system' is thus different from – although related to – its meaning as a technical term in the grammar (see Section 1.3.2 above). The system in this general sense is equivalent to the totality of all the specific systems that would figure in a comprehensive network covering every stratum.

in terms of some general theory (in this case, of climatic change), or is it just a set of similar events? An analogous question about language would be if we took a corpus of, say, writings by political scientists and asked, are these just a set of similar texts, or do they represent a sub-system of the language? The climate is the ***theory*** of the weather. As such, it does have its own separate existence – but (like all theoretical entities) it exists on the semiotic plane. It is a virtual thing. Likewise with the system of language: this is language as a virtual thing; it is not the sum of all possible texts but a theoretical entity to which we can assign certain properties and which we can invest with considerable explanatory power.

System and text are thus related through instantiation. Like the relationship between climate and weather, the relationship between system and text is a cline – the **cline of instantiation** (Figure 1-11). System and text define the two poles of the cline – that of the overall potential and that of a particular instance. Between these two poles there are intermediate patterns. These patterns can be viewed either from the system pole as

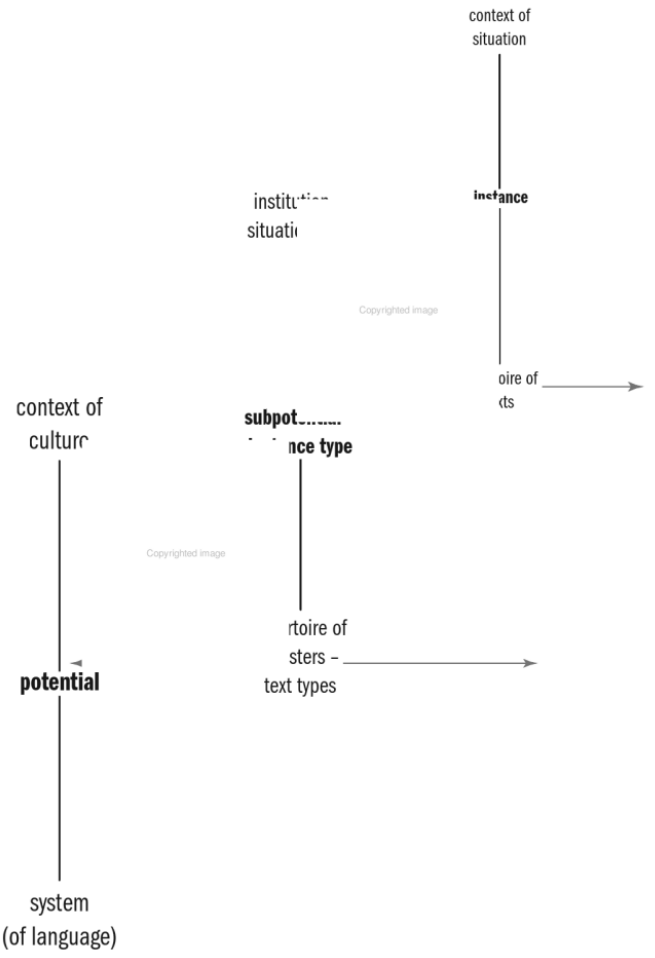

Fig. 1-11 The cline of instantiation

sub-systems or from the instance pole as instance types. If we start at the instance pole, we can study a single text, and then look for other texts that are like it according to certain criteria. When we study this sample of texts, we can identify patterns that they all share, and describe these in terms of a **text type**. By identifying a text type, we are moving along the cline of instantiation away from the text pole towards the system pole. The criteria we use when we compare the texts in our sample could, in principle, come from any of the strata of language – as long as they are systematic and explicit. However, research has shown that texts vary systematically according to contextual values: texts vary according the nature of the contexts they are used in. Thus recipes, weather forecasts, stockmarket reports, rental agreements, e-mail messages, inaugural speeches, service encounters in the local deli, news bulletins, media interviews, tutorial sessions, walking tours in a guide book, gossip during a tea-break, advertisements, bedtime stories, and all the other innumerable text types we meet in life are all ways of using language in different contexts. Looked at from the system pole of the cline of instantiation, they can be interpreted as **registers**. A register is a functional variety of language (Halliday, McIntosh & Strevens, 1964; Halliday, 1978) – the patterns of instantiation of the overall system associated with a given type of context (a **situation type**).[8] These patterns of instantiation show up quantitatively as adjustments in the systemic probabilities of language; a register can be represented as a particular setting of systemic probabilities. For example, the future tense is very much more likely to occur in weather forecasts than it is in stories (for examples of quantitative profiles of registers, see Matthiessen, 2002a, 2006a).

If we now come back to the question of stratification, we can perhaps see more clearly what it means to say that the semantic stratum is language interfacing with the non-linguistic (prototypically material) world. Most texts in adult life do not relate directly to the objects and events in their environment. Mandela's text was highly abstract, and even when he talked about *the soil of this beautiful country* and *the jacaranda trees of Pretoria* it is very unlikely that he could actually see them at the time. They were not a part of the setting in that instance. Nevertheless the meanings that are realized by these wordings, and the meanings realized by *an extraordinary human disaster* and *humanity's belief in justice* are, ultimately, construals of human experience; and when we now read or listen to that text we are understanding it as just that. Interfacing with the eco-social environment is a property of language as system; it is also, crucially, a feature of those instances through which small children come to master the system; but it is not something that is re-enacted in every text. Experience is remembered, imagined, abstracted, metaphorized and mythologized – the text has the power to create its own environment; but it has this power because of the way the system has evolved, by making meaning out of the environment as it was given.

As grammarians we have to be able to shift our perspective, observing now from the system standpoint and now from that of the text; and we have to be aware at which point we are standing at any time. This issue has been strongly foregrounded by the appearance of the computerized corpus. A corpus is a large collection of instances – of spoken and written

[8] Here the term 'register' thus refers to a functional variety of language (see e.g. Halliday, 1978; Hasan, 1973; Matthiessen, 1993b; Ghadessy, 1993; Lukin *et al.*, 2008). It has also been used in a related, but different way, to refer to the contextual values associated with such a functional variety (see Martin, 1992, and other contributions to the 'genre model' within systemic functional linguistics; cf. Matthiessen, 1993b).

texts; the corpuses now available contain enough data to give significantly new insights into the grammar of English, provided the data can be processed and interpreted. But the corpus does not write the grammar for you, any more than the data from experiments in the behaviour of light wrote Newton's *Opticks* for him; it has to be theorized. Writing a description of a grammar entails constant shunting between the perspective of the system and the perspective of the instance. We have tried in this edition to take account of the new balance that has arisen as a result of data becoming accessible to grammarians in sufficient quantity for the first time in the two and a half millennia history of the subject.

1.3.5 Metafunction

This brings us back to the question asked in Section 1.3.3: what are the basic functions of language, in relation to our ecological and social environment? We suggested two: making sense of our experience, and acting out our social relationships.

It is clear that language does – as we put it – **construe** human experience. It names things, thus construing them into categories; and then, typically, goes further and construes the categories into taxonomies, often using more names for doing so. So we have *houses* and *cottages* and *garages* and *sheds*, which are all kinds of *building*; *strolling* and *stepping* and *marching* and *pacing*, which are all kinds of *walking*; *in*, *on*, *under*, *around* as relative locations, and so on – and the fact that these differ from one language to another is a reminder that the categories are in fact construed in language (cf. Halliday & Matthiessen, 1999: Chapter 7; Caffarel, Martin & Matthiessen, 2004). More powerfully still, these elements are configured into complex grammatical patterns like *marched out of the house;* the figures can be built up into sequences related by time, cause and the like – there is no facet of human experience that cannot be transformed into meaning. In other words, language provides a **theory** of human experience, and certain of the resources of the lexicogrammar of every language are dedicated to that function. We call it the **ideational** metafunction, and distinguish it into two components, the **experiential** and the **logical** (see Chapter 5 and Chapter 7).

At the same time, whenever we use language there is always something else going on. While construing, language is always also **enacting**: enacting our personal and social relationships with the other people around us. The **clause** of the grammar is not only a figure, representing some process – some doing or happening, saying or sensing, being or having – together with its various participants and circumstances; it is also a proposition, or a proposal, whereby we inform or question, give an order or make an offer, and express our appraisal of and attitude towards whoever we are addressing and what we are talking about. This kind of meaning is more active: if the ideational function of the grammar is 'language as reflection', this is 'language as action'. We call it the **interpersonal** metafunction, to suggest that it is both interactive and personal (see Chapter 4).

This distinction between two modes of meaning is not just made from outside; when the grammar is represented systemically, it shows up as two distinct networks of systems (Halliday, 1969; cf. Martin, 1991, on intrinsic functionality). What it signifies is that (1) every message is both about something and addressing someone, and (2) these two motifs can be freely combined – by and large, they do not constrain each other. But the grammar also shows up a third component, another mode of meaning that relates to the construction of text. In a sense this can be regarded as an enabling or facilitating function, since both the others – construing experience and enacting interpersonal relations – depend on being able

to build up sequences of discourse, organizing the discursive flow, and creating cohesion and continuity as it moves along. This, too, appears as a clearly delineated motif within the grammar. We call it the **textual** metafunction (see Chapters 3 and 9).

Why this rather unwieldy term 'metafunction?' We could have called them simply 'functions'; however, there is a long tradition of talking about the functions of language in contexts where 'function' simply means purpose or way of using language, and has no significance for the analysis of language itself (cf. Halliday & Hasan, 1985: Ch. 1; Martin, 1991). But the systemic analysis shows that functionality is ***intrinsic*** to language: that is to say, the entire architecture of language is arranged along functional lines. Language is as it is because of the functions in which it has evolved in the human species. The term 'metafunction' was adopted to suggest that function was an integral component within the overall theory (Figure 1-12).

M

Fig. 1-12 Metafunction

1.4 Context, language and other semiotic systems

We have now introduced the major semiotic dimensions that define the 'architecture' of language in context (cf. Halliday, 2003: 1–29; Matthiessen, 2007a). Some of these dimensions enable us to locate lexicogrammar in relation to the other sub-systems that make up the total system of language; these are known as **global dimensions** because they

determine the overall organization of language in context: the hierarchy of stratification, the cline of instantiation, and the spectrum of metafunctions. The other dimensions enable us to characterize the internal organization of lexicogrammar and also of the other sub-systems of language, and of context; these are known as **local dimensions** because they operate locally within linguistic sub-systems. Let us summarize the semiotic dimensions of language in context under these two headings: see Table 1-7.

Table 1-7 The global and local semiotic dimensions of language in context

1.4.1 Context; language in relation to context

As we have noted above (in particular in Section 1.3.4), language operates in context. In terms of linguistic theory, we recognize this important principle by developing an 'ecological' theory of language – one in which language is always theorized, described and analysed within an environment of meanings; a given language is thus interpreted by reference to its semiotic habitat. This way of approaching language was given a considerable theoretical and empirical boost by the anthropologist Bronislaw Malinowski in the 1920s and 1930s, based initially on his extensive fieldwork in the Trobriand Islands in the 1910s; and his insights were taken up and developed within linguistic theory by J.R. Firth, and then built into a general theory of language in context by systemic functional linguists (e.g. Halliday, McIntosh & Strevens, 1964; Halliday, 1978, 1992a; Halliday & Hasan, 1985; Ghadessy, 1999; Butt & Wegener, 2007). This is the conceptualization of context that we use here.

Like language, context is extended along the cline of instantiation (Section 1.3.4) from instance to potential; and like language, it is functionally diversified (Section 1.3.5). Let us discuss these two different aspects of the organization of context, and then introduce a context-based typology of texts that we will be using throughout this book when we give examples of grammar operating in text.

As shown in Figure 1-11, context extends along the cline of instantiation (Section 1.3.4) from the overall contextual potential of a community to the contextual instances involving particular people interacting and exchanging meanings on particular occasions. The contextual potential of a community is its culture – what we call the **context of culture**,

following Malinowski. The context of culture is what the members of a community can mean in cultural terms; that is, we interpret culture as a system of higher-level meanings (see Halliday, 1978) – as an environment of meanings in which various semiotic systems operate, including language, paralanguage (gesture, facial expression, voice quality, timbre, tempo, and other systems of meaning accompanying language and expressed through the human body; cf. Thibault, 2004) and other human systems of meaning such as dance, drawing, painting and architecture (e.g. Kress & van Leeuwen, 1996; O'Toole, 1994; Martinec, 2005). Describing the cultural potential of a community is obviously a huge undertaking – one requiring the kind of commitment, support, recognition and funding given to the Human Genome Project, started in 1990. Perhaps we can imagine a vast number of Human Sememe Projects given the task of mapping out the cultural potentials of all human societies (or Human Meme Projects, to relate to Richard Dawkins' notion of meme as a cultural replicator). While no project on this scale has yet been undertaken, it is certainly possible to see the theoretical significance of such a project.

From a practical point of view, one research strategy is to move along the cline of instantiation from the potential pole towards the instance pole: for researchers, it is considerably easier to take on the task of describing a particular **cultural domain**, or **institution**, based on evidence gathered from the various **contexts of situation** operating within that institution (cf. Matthiessen, 2009c: Section 3.6). Malinowski (e.g. 1944) called institutions the 'real isolates of culture', advocating the study of institutions; and we can investigate an institution linguistically through the register that operates within it (cf. Section 1.1 above). While describing the overall potential of a culture is a daunting task, mapping out an institution by identifying and describing the different types of situation that collectively constitute the institution is a more manageable undertaking, and systemic functional contributions along these lines have been made in a number of areas including the family, education, administration, the media, and healthcare.

While there are still no comprehensive descriptions of the context of culture, the general categories of context have been known for a long time – see Halliday, McIntosh & Strevens (1964: 90–94); and they have been explored under the headings of **field**, **tenor** and **mode** (e.g. Hasan, 1973; Halliday, 1978; Halliday & Hasan, 1985; Martin, 1992). Thus any situation type can be characterized in terms of field, tenor and mode:

- **field** – what's going on in the situation: (i) the nature of the social and semiotic activity; and (ii) the domain of experience this activity relates to (the 'subject matter' or 'topic')
- **tenor** – who is taking part in the situation: (i) the roles played by those taking part in the socio-semiotic activity – (1) institutional roles, (2) status roles (power, either equal or unequal), (3) contact roles (familiarity, ranging from strangers to intimates) and (4) sociometric roles (affect, either neutral or charged, positively or negatively); and (ii) the values that the interactants imbue the domain with (either neutral or loaded, positively or negatively)
- **mode** – what role is being played by language and other semiotic systems in the situation: (i) the division of labour between semiotic activities and social ones (ranging from semiotic activities as constitutive of the situation to semiotic activities as facilitating); (ii) the division of labour between linguistic activities and other

semiotic activities; (iii) rhetorical mode: the orientation of the text towards field (e.g. informative, didactic, explanatory, explicatory) or tenor (e.g. persuasive, exhortatory, hortatory, polemic); (iv) turn: dialogic or monologic; (v) medium: written or spoken; (vi) channel: phonic or graphic.

Field, tenor and mode are thus sets of related variables, with ranges of contrasting values. Together they define a multi-dimensional semiotic space – the environment of meanings in which language, other semiotic systems and social systems operate. The combinations of field, tenor and mode values determine different uses of language – the different meanings that are at risk in a given type of situation. There are systematic correspondences between the contextual values and the meanings that are at risk in the contexts defined by these values. As Halliday (1978) suggested, field values **resonate** with ideational meanings, tenor values resonate with interpersonal meanings, and mode values resonate with textual meanings (see also Halliday & Hasan, 1985: 26)[9]. In other words, the correspondences between context and language are based on the functional organization of both orders of meaning.

The ideational, interpersonal and textual ***meanings at risk*** can be stated in terms of systems at the semantic stratum in the first instance. However, since semantics stands in a natural relation to lexicogrammar, the two being the content plane systems of language, meanings at risk can also be stated, by another stratal step, in terms of systems at the lexicogrammatical stratum as ***wordings at risk***. For example, when we consider the correlations between tenor values and terms in interpersonal systems, we should really focus on interpersonal semantic systems such as SPEECH FUNCTION in the first instance rather than on lexicogrammatical ones such as MOOD (to take an area of content from Chapter 4). Thus combinations of tenor values relating to (a) status and (b) contact correlate with different semantic strategies open to speakers for demanding goods-&-services of their listeners – for commanding their listeners. If (a) the status is unequal, with the speaker being subordinate to the listener and (b) the contact is minimal, the speaker's semantic options are very limited: it is very hard to command a stranger who is of superior status to do something; but there will be certain semantic strategies. Lexicogrammatically, these strategies will be far removed from the congruent realization of a command, a clause of the imperative mood – perhaps something like *I wonder if you would be so kind as to* ... and they will be 'dispersed' in the grammar of mood, involving not only 'imperative' clauses but also 'declarative' and 'interrogative' ones and in fact not only clauses but also combinations of clauses (see Chapter 10, Section 10.4); but semantically, they are still within range of options associated with commands. Thus accounts of 'politeness' have tended to be cast in semantic terms rather than in lexicogrammatical ones (e.g. Brown & Levinson, 1987 – an influential contribution: see Watts, 2003, for a critical review of their framework and the literature since their work; and systemic functional accounts, e.g. Bateman, 1988; Butler, 1988). Tenor is, as it were, refracted through semantics so that the lexicogrammatical resonances with tenor values are more indirect than the semantic ones.

[9] We use the term 'resonate with' because the relationship is not a one-way causal relationship, but rather a two-way realizational relationship (cf. Jay Lemke's, 1984, notion of **metaredundancy**, discussed in Halliday, 1992d). Contextual values influence linguistic choices but are also influenced by them.

Still, while field, tenor and mode resonate with semantic systems in the first instance, they do penetrate into lexicogrammar: field values put ideational wordings at risk, tenor values put interpersonal wordings at risk, and mode values put textual wordings at risk. This was in fact shown by Brown & Gilman's (1960) classic study of 'the pronouns of power and solidarity': the tenor variables of power (or 'status') and solidarity (related to 'contact', in our characterization of tenor above) resonate with different uses of the system of pronouns in various languages[10]. Indeed, the tenor variables of power and contact may be grammaticalized as part of the core interpersonal system of mood in a language, as in Japanese and Korean (see Matthiessen, Teruya & Wu, 2008, and references therein).

Field, tenor and mode variables are the basis for any attempt to develop a taxonomy of situations. At the same time, since text is language functioning in context, the field, tenor and mode variables are also the basis of any attempt to develop a taxonomy of texts operating in situations. It is certainly true that in developing a taxonomy of texts, we can adopt – we need to adopt – a trinocular perspective (see Section 1.5.1 below), matching up contextual, semantic and lexicogrammatical considerations to support the taxonomy. However, to be meaningful, a taxonomy of texts must be grounded in contextual considerations. If the taxonomy is 'on the right track', semantic and lexicogrammatical considerations will align themselves with the contextual ones.

In principle, such a taxonomy would be based on all three contextual variables – on field, tenor and mode. However, here we will present a contextual taxonomy of text that is based on field in the first instance, more specifically on the variable of socio-semiotic activity (see Matthiessen, 2006c; Matthiessen, Teruya & Lam, 2010; Teruya, 2007). We will use this taxonomy throughout the book, classifying the illustrative texts we introduce according to this taxonomy (as we have already done – e.g. Texts 1-1 through to 1-3).

So let's consider the nature of the socio-semiotic activity that constitutes a situation. In a sense, the activity that constitutes a situation is either one of behaviour or one of meaning; this is the traditional distinction between action and reflection. So we will make a basic distinction between activities of 'doing' and of 'meaning', and then further distinctions within 'meaning':

- 'doing': the situation is constituted in some form of social behaviour, involving one or more persons. Language or other semiotic systems such as gesture, gaze and facial expression may be engaged to facilitate the performance of the activity, as when language is used to coordinate a team
- 'meaning': the situation is constituted in some process of meaning. There are seven primary types:
  - 'expounding': expounding knowledge about the world – about general classes of phenomena, categorizing them or explaining them
  - 'reporting': reporting particular phenomena, chronicling the flow of events, surveying places or inventorying entities

[10] Another classic study that illuminates the relationship between tenor and lexicogrammar is Ervin-Tripp's (1972) account of terms of address in American English in relation to tenor values. The general point is that a given interpersonal system can mean in more than one way depending on the tenor values of the context in which it operates.

- 'recreating': recreating any aspect of prototypically human life imaginatively by dramatizing or narrating events
- 'sharing': sharing personal experiences and values, prototypically in private
- 'enabling': enabling some course of activity, either enabling the activity by instructing people in how to undertake it or regulating the activity by controlling people's actions
- 'recommending': recommending some course of activity, either for the sake of the speaker through promotion of some commodity or for the sake of addressee through advice
- 'exploring': exploring societal values and positions, prototypically in the public arena.

These primary types of socio-semiotic activity can be represented typologically as a system network, of course; but we can also represent them topologically, as shown in Figure 1.13. This 'pie' diagram suggests that different types shade into one another, which is indeed the case. For example, 'reporting' and 'recreating' shade into one another to provide a context for fictionalized biographies ('fake histories': see Halliday, 2010: Section 5.3), and 'recommending' and 'reporting' shade into one another to provide a context for infomercials.

Figure 1-13 represents two steps in delicacy – the eight primary types, and secondary distinctions within each type. When we reach secondary delicacy, we can begin to discern the structuring of situations belonging to the different types. And if we take one step further, we can relate tertiary types to generic structures in the literature on 'genres'. For example, situations characterized by processes of explaining can be structured in a number of different ways, but once we differentiate these explanation strategies at tertiary delicacy, for example distinguishing factorial explanations from sequential ones, we can assign the distinct structures described by Veel (1997) for explanations in school science. Similarly, there are different ways of staging narrative situations, but once we take another step and differentiate folk tales, exempla, anecdotes and other narrative strategies, we can refer to the different narrative structures that have been described in the literature (e.g. Hasan, 1984; Eggins & Slade, 1997; Martin & Rose, 1994: Chapter 2). (Clearly, the specification of the structure of situation is staged in terms of delicacy. For example, narrative situations in general share structures that embody temporal sequence, but they differ in terms of other structural elements and even with respect to what factors 'drive' the temporal sequence.)

The question of how many steps in delicacy we have to take before we can begin to discern the distinct structures of different types of situation is obviously important from the point of view of contextual description. This question is also significant from the point of view of the description of lexicogrammar. Here the issue is at what point we can begin to discern interestingly distinct uses of the resources of lexicogrammar. As a rule of thumb, we can say that this happens at the point at which the contextual structure of a situation can be fully specified. This will mean, among other things, that the elements of the structure of a situation can be investigated in terms of their distinct patterns of lexicogrammatical realization, as illustrated by e.g. Halliday (1982) and Fries (1985).

Fig. 1-13 Field – socio-semiotic process (activity) represented as a topology

The socio-semiotic activity types in Figure 1-1 combine with values from the other field variable, i.e. with values from experiential domains (areas of subject matter); and these combinations are likely to yield somewhat different patterns. For example, explanations in school science (e.g. Veel, 1997) and in school history (e.g. Coffin, 2006) are similar in various ways, but they are also different in certain respects. For example, the repertoire of explanations deployed in school history appears to be narrower; and they differ lexicogrammatically: scientific terms are construed and used very frequently in school science but hardly at all in school history (cf. Eggins, Wignell & Martin, 1993).

The socio-semiotic activity types in Figure 1-1 also combine with variables within the tenor and mode parameters of context. In terms of **tenor**, we can consider the different socio-semiotic activities in relation to different role combinations of those taking part in the activity. For example, we can imagine a series of concentric circles for different institutional role combinations as children would meet them in the course of their development as they

grow up, and then compare sets of roles for different activities, e.g. parents explaining phenomena to their children in the home versus teachers explaining phenomena to their students in primary school – or the other way around, as illustrated by the analysis in Halliday (2002b: 313–322) on an explanation offered by a five-year-old boy to his father of why the North Star stays still but the other stars don't (see Chapter 3, Text 3-7). Tenor considerations thus include the range of 'voices' taking part in the different socio-semiotic activities, including degrees of expertise and of professionalism.

In terms of **mode**, we can intersect the socio-semiotic activity types with different combinations of (a) turn: dialogic vs. monologic, and (b) medium: written vs. spoken; these combinations are represented as four different concentric circles in e.g. Matthiessen (2006c: 46), Teruya (2007), and Matthiessen, Teruya & Lam (2010: 221). The intersection of the socio-semiotic activity types with (a) turn and (b) medium makes it possible to locate the short numbered text examples we give throughout this book: see Table 1-8. Quite a few cells are empty, but most of them could in fact be filled with examples of lexicogrammatically analysed texts, and we hope that systematic overviews can be presented in future publications.

In the same way, we could use this matrix of field and mode values to locate many of the descriptions of 'genres' that have been developed over the past few decades: compare the summaries in Martin & Rose (2008), Christie & Derewianka (2008), Eggins & Slade (2005). Many of these accounts provide very good insights into the lexicogrammar 'at work' in different contexts. The accounts reveal fairly general tendencies such as the deployment of interpersonal lexis in various texts operating in 'sharing' and 'exploring' contexts, the use of grammatical metaphor within the ideational metafunction in texts operating in 'expounding' and 'exploring' contexts. They also reveal quite specific patterns such as the emergence in late adolescence of the 'relational' grammar of proving in texts operating in 'exploring' contexts (Christie & Derewianka, 2008: 222, 232).

In addition to turn and medium, mode variables also include channel, division of labour and rhetorical mode.

The division of labour between social activities and semiotic ones varies according to the activities represented in Figure 1-1. In 'doing' contexts, people exchange meanings to facilitate the social tasks they are concerned with, as when medical teams collaborate to perform surgery or when removalists talk and gesture (if their hands are free!) to coordinate difficult tasks (as in Text 4-3 in Chapter 4). In all other contexts, most of the socio-semiotic labour is semiotic in the first instance: the exchange of meanings in language and other semiotic systems is constitutive of the contexts in which it operates.

The division of labour between language and other semiotic systems covers the whole range, from contexts where language does all the semiotic work to contexts where all the semiotic work is done by some semiotic system or systems other than language. The possible combinations naturally depend on the nature of the channel, but they also vary according to field. For example, recounts in history using a graphic channel might be accompanied by timeline diagrams, but instructions in software manuals might be accompanied by flowcharts (cf. Matthiessen, 2009a).

Rhetorical mode encompasses a number of rhetorical categories concerned with the contribution of the text to the situation it operates in: informative, didactic, persuasive, exhortatory, pragmatic, and so on. However, we can relate these particular categories to

Table 1-8 Matrix of socio-semiotic activity types intersected with turn and medium. with classification of text samples used in this book (examples of 'genres' in italics)

Table 1-8 *continued*

the **orientation** of the text (i) towards the field of the situation, (ii) towards the tenor or (iii) towards some mixture of both.

(i) Orientation towards field means that the goals of the situation, or intended outcomes, are concerned with field, more specifically with the development of field, as in an 'expounding' context where the speaker's goal might be to construe a taxonomy for the addressee, a classification of some classes of phenomena. When texts operate in such situations, they tend to be organized in terms of field – in terms of the structure of the field, as when a text is organized according to the classes of a taxonomy (e.g. Text 9-12, Text 9-13). Orientation towards field is characteristic of 'expounding', 'reporting' and 'doing' contexts, and also in principle of 'enabling' contexts of the 'instructing' subtype.

(ii) Orientation towards tenor means that the goals of the situation, or intended outcomes, are concerned with tenor, more specifically with the relationship between speaker and addressee – with maintaining or changing this relationship, as when speakers try to bring their addressees closer to their own positions (e.g. the text in Table 9-20). When texts operate in such situations, they tend to be organized in terms of tenor, with a central proposition or proposal supported by text segments that provide evidence for the proposition, increasing the likelihood that the addressee will agree, or motivation for the proposal, increasing the likelihood that the addressee will comply (if the proposal is some form of command) or accept (if the proposal is some form of offer). The orientation towards tenor is thus likely to be reflected in the semantic organization of texts operating in 'recommending' and 'exploring' contexts in the use of fairly global **internal relations** – called internal conjunctive relations (see Halliday & Hasan, 1976; Martin, 1992) or internal rhetorical relations (see Mann & Matthiessen, 1991). Both evidence and motivation can be interpreted as internal versions of cause – evidence: 'I claim/you should believe that ... because ... '; motivation: 'I want you/you are obliged to ... because ...'. In general, orientation towards tenor is characteristic of 'sharing', 'recommending' and 'exploring' contexts, and also in principle of 'enabling' contexts of the 'regulatory' subtype (but see immediately below). In contrast, texts operating in contexts with an orientation towards field are much less likely to involve internal relations; instead, they are organized both globally and locally in terms of **external relations**. For the contrast between internal and external relations, see further Chapter 9, Section 9.3.2; for the orientation in the organization of text towards interpersonal or ideational meanings, see Halliday (2001).

(iii) Orientation towards both field and tenor means that the goals of the situation, or intended outcomes, are concerned with field and/or tenor. Thus the goals of 'recreating' situations may be concerned with the construal of some imaginary world, ranging from a slight variant of our own world to a world of pure fantasy; but the goals may at the same time involve moral principles embodied in tenor. In this way, utopias and dystopias are concerned with both field and tenor. The orientation towards both field and tenor is reflected in the structure of traditional folk tales or nursery tales: field is reflected in the sequence of events (initiating, sequent and final), and tenor is reflected in evaluations, which may be strung out

prosodically through the narrative and/or encapsulated in a separate 'moral' at the end of the tale (cf. Hasan, 1984). The goals of 'enabling' situations can also be said to relate to both field and tenor, but in a different way. They concern the addressee's activities in some field, but tenor comes into the picture as well because these activities are 'modulated': the addressee is either capacitated to undertake them (instruction) or required to do so (regulation). In terms of their organization, instructional texts tend to be more field-like, being organized as the sequence of steps that make up a procedure. In contrast, regulatory texts have a less clearly field-based organization; and like 'promoting' texts, they may include motivations – although typically threats of forms of punishment rather than the irresistible features of a product or service!

Channel determines the 'bandwidth' of the flow of meanings in a situation. For most of human history, the channel was only phonic, but typically with visual contact (thus also allowing for accompanying gestures, facial expressions and other forms of visual 'paralanguage'); but with the gradual emergence of writing, initially in certain city-based civilizations around five thousand years ago, graphic channels were added, and archival uses of language became possible. Technological advances have continued to enhance the potential of both phonic and graphic channels, and to enable mixed channels (cf. Halliday, 2008: 140–141). Importantly, mobile and web-based technologies (for hardware and software) have changed the possibilities of 'sharing' in rather dramatic ways, with a whole host of new options like e-mail messages, text messages, blogs, tweets and other formats associated with social media, as investigated and discussed by Macnamara (2010). As a result, the distinction between the private sphere of 'sharing' values and opinions and the 'public' sphere of exploring them has become blurred. This can be seen in places where users post reviews of commodities such as film and music, such as the Internet Movie Data Base. Reviews range from opinions that might be shared in casual conversation to analytical evaluations of the kind we would expect from expert reviewers contributing to quality newspapers. Since both 'sharing' and 'exploring' contexts are important sites for the instantiation, and (in the phylogenetic time-frame) for the evolution, of interpersonal meaning, there are likely to be interesting changes in patterns of interpersonal meaning-making – more profound than the addition of 'emoticons' to written conversation to make up for some of the loss of intonation and voice quality in spoken language. Of course, technological advances do not affect only 'sharing' and 'exploring' contexts, but also the other types of situation characterized in terms of socio-semiotic activities in Figure 1-1. However, the development of social media does indicate the extent to which companies are trying to leverage people's need for 'sharing' and orientation towards the interpersonal.

1.4.2 Semantics

Semantics is the highest stratum within language; it serves as an 'interface' between language and the environment outside language, as shown in Table 1-5. This means that semantics interfaces with context, but not only with context – it also interfaces with other systems that operate within context, viz. with the content systems of other semiotic systems and

with bio-semiotic systems such as our systems of perception and our system of bodily action (cf. Halliday & Matthiessen, 1999).

As the upper of the two content strata within language, semantics is the interface between context and lexicogrammar. Semantics transforms experience and interpersonal relationships into linguistic meaning, and lexicogrammar transforms this meaning into words, as we put it above (see Table 1-5), adopting the speaker's perspective.

The basic unit of semantics is the **text** – language functioning in context, an instance of the semantic system. A text is organized internally as patterns of logical, experiential, interpersonal and textual meaning. At the same time, it is organized externally as a unit operating in context: the structure of the context of situation that a text operates in is, as it were, projected onto the text. If the situation is one of 'meaning' in terms of the socio-semiotic activity in Figure 1-1 then the entire structure of the situation is projected onto the text. For example, in a situation of telling a traditional folk tale, the structure would be (from Hasan, 1984, but slightly simplified):

(Placement ^) Initiating Event ^ Sequent Event^{1-n} ^ Final Event (^ Finale) (° Moral)

This structure is projected onto the text operating in the situation, as illustrated in Text 5-2, Chapter 5 – and possibly also onto other accompanying semiotic processes such as a musical score. Each element, or stage, of the structure of the situation is realized by distinctive semantic patterns, as illustrated for Placement by Hasan (1984). These distinctive semantic patterns are, in turn, realized by distinctive lexicogrammatical patterns; but the patterns of wording in the lexicogrammar are always mediated by the patterns of meaning in the semantics. In Text 5-2, the beginning of the Placement is realized (via the semantics) by an 'existential' clause, followed by a 'material' clause (to characterize them in the experiential terms that we will introduce in Chapter 5):

Once, a very long time ago, there lived a man called Noah. He and his wife and his sons and their wives all worked very hard.

The 'existential' clause serves to introduce the protagonist of the tale, *a man called Noah*, as the Existent – the participant in the process of living; and the Existent is given the textual status of New information, the main point of the message (to use textual terms we will introduce in Chapter 3). The protagonist is presented in this way against the background of Time, *once, a very long time ago*; and this temporal circumstance is given the textual status of Theme, the orientation chosen for the message. The next clause still contributes to the development of the Placement: elaborating on Noah, it construes a habitual activity he and his family took part in.

The need for a placement in a traditional situation of narration thus 'trickles' down from context via the semantics to the lexicogrammar, and this need is met by the lexicogrammatical choices that we have just illustrated. However, unless texts are very short (like traffic signs), there are layers of semantic patterning between the whole text and the local units that are realized lexicogrammatically by clauses such as *Once, a very long time ago, there lived a man called Noah*. Texts have 'depth' – ordered layers of semantic patterns, ranging from the global semantic domain of the whole text to local semantic domain corresponding to

domains of lexicogrammatical patterning. This depth is reflected in traditional accounts within composition and rhetoric in notions like rhetorical paragraph and topic sentence, and linguists and other scholars concerned with the analysis of texts and the description of the systems that lie behind them have proposed various frameworks for accounting for the 'depth' of texts, including pioneering contributions from the broad tradition of tagmemic linguistics (e.g. Grimes, 1975; Beekman, Callow & Kopesec, 1981; Longacre, 1996; Longacre & Hwang, 2012; Pike, 1992).

In general, two approaches to the account of the depth of texts have emerged in various traditions: the depth of layering may be modelled in terms of a **semantic rank scale** operating with some kind of constituency structure (analogous to the lexicogrammatical and phonological rank scales discussed above), as in Longacre's work since the 1970s; or it may be modelled in terms of **internal nesting of relational organization** – along the lines of Grimes (1975) and Beekman, Callow & Kopesec (1981). Within systemic functional linguistics, we also find these two models of the depth of text – the rank-scale model with rhetorical units proposed by Cloran (1994) and the internal-nesting model derived from Rhetorical Structure Theory (RST, e.g. Matthiessen & Thompson, 1988; Matthiessen, 1992, 2002a). The two are applied to the analysis of the same text by Cloran, Stuart-Smith & Young (2007). They are not, of course, mutually exclusive; they can be interpreted as capturing different aspects of the 'depth' of texts. And as grammarians we do not have to choose between the two as long as they provide us with motivated accounts of how to relate semantics to grammar. However, at various points in our account of grammar, we will make use of the internal nesting model derived from RST since this model will enable us to explore the nature of clause complexing (Chapter 7), of the use of cohesive conjunctions (Chapter 9), of grammatical metaphor (Chapter 10), and of the source of methods of development relating to the choice of Theme (Chapter 3).

Let us try to summarize the salient features of the relationship between situation (context), text (semantics) and clause (lexicogrammar) by means of a diagram: see Figure 1.14. Globally, a text is structured according to the situation it operates in; the contextual structure is projected onto the text, and the contextual elements are realized by patterns of meaning in the text. As a semantic unit, the text consists of semantic domains of different sizes. It is likely to consist of **rhetorical paragraphs** (or **parasemes** (see Halliday, 2002d), which may or may not correspond to orthographic paragraphs in writing). In turn, these consist of **sequences** – sequences of **figures**, i.e. configurations of processes, participants involved in these and attendant circumstances. These more local domains, sequences and figures, are typically realized grammatically: sequences are realized by clause complexes, and figures by clauses. Here the grammar provides a good deal of guidance through the grammatical structure of the clause complex (Chapter 7), the clause (Chapters 3 through to 5) and its constituent parts (Chapters 6 and 8). Thus the grammar makes the local structure of the text 'tighter', more highly integrated, by constructing it not only as meaning but also as wording. However, the grammar also provides some important guidance beyond the domain of the clause complex, i.e. beyond the most extensive domain of grammatical structure. It does this by means of the resources of cohesion (Chapter 9), e.g. by means of cohesive conjunctions such as *for example*, *in addition*, *in contrast*, *therefore*, *meanwhile*, which can mark relations between sequences realized by clause complexes and also between (groups of) rhetorical paragraphs.

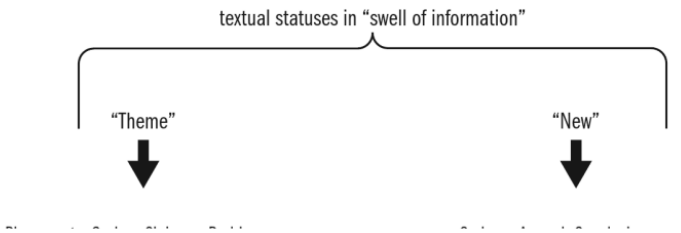

Copyrighted image

tural transitions
etorical relations)
'method of
velopment"

Fig. 1-14 Semantics as an 'interlevel' between context and lexicogrammar

In describing the structure of the text, we have foregrounded the perspective of the ideational metafunction. Sequences are construed through logical resources and figures through experiential ones; and rhetorical paragraphs and groups of paragraphs can be interpreted as being formed by logical resources – in terms of logico-semantic relations (cf. Halliday, 2001; Matthiessen, 2002a). At the same time, texts are also organized in terms of interpersonal and textual patterns of meaning. Interpersonally, a text is a series of **exchanges** between speaker and addressee – even if it is a one-sided monologue that is essentially a series of statements acknowledged silently by the addressee. These exchanges are propelled forward locally by **moves**, which are realized by clauses in their interpersonal guise (Chapter 4). Textually, a text is a flow of information, or, more accurately, waves of

information. These wave patterns extend from the whole text through rhetorical paragraphs to local waves, or **messages** – quanta of information that are realized by clauses in their textual guise, and (in spoken language) also by information units (Chapter 3).

We will discuss the complementary metafunctional perspectives on the text at various points in the book in order to illuminate the grammar, and we will return to the relationship between semantics and lexicogrammar in Chapter 10, Section 10.1.

1.4.3 Language and other semiotic systems in context

The term 'text' includes both spoken and written instances of the linguistic system. As a scientific term in linguistics, it thus differs from a common everyday sense of text as a piece of writing – and now also as a verb, in the sense of 'sending text messages' with the help of some mobile message service (as in *In what may have been a final, frantic act, Conaway texted relatives an hour later, saying they were trapped in the trunk of a car.*). In linguistics, 'text' thus means an instance of the linguistic system. However, the sense of text is being extended to other semiotic systems, and scholars refer to instances of e.g. 'visual semiotic' systems as '(visual) texts' (thus a painting would be a visual semiotic text) and they also refer to 'multimodal texts' – instances of more than one semiotic system. While this extended sense of 'text' is still hard to find in dictionaries, it has clearly been established; for example, ACARA, the Australian Curriculum and Assessment Authority glosses 'multimodal text' as 'combination of two or more communication modes (for example, print, image and spoken text, as in film or computer presentations)'[11]. And the same has happened to 'discourse' (see e.g. Kress & van Leeuwen's, 2001, pioneering contribution; researchers now talk about 'MDA': 'multimodal discourse analysis'.

Most accounts of 'multimodal text' so far have probably focused on combinations of written texts and instances of 'visual semiotic' systems. From a developmental and evolutionary point of view, it would make more sense to start with spoken texts unfolding together with instances of other **somatic semiotic systems** (i.e. other semiotic systems using some aspect of the body as their expression plane; see Matthiessen, 2009a, and cf. Thibault's, 2004, notion of the 'signifying body') before moving on to interpret and describe **exo-somatic semiotic systems**. Indeed, the protolanguages of early childhood tend to be both vocal and gestural in their expression (see Halliday, 1975, 1992d, 2004); and we can hypothesize that the same was true of protolanguages in human evolution (see Matthiessen, 2004a). Using this starting point, we could investigate the relationship (including relative timing) between choices in lexicogrammar and choices in semiotic systems other than language such as gesture, facial expression and vocal paralanguage in face-to-face conversation, building on contributions such as McNeill (2000). Here the research by James Lantolf and his team of researchers in the context of (advanced) second language learning is a source of insights and ideas. For example, he has shown that English and Spanish differ in how they construe motion through space in terms of the division of labour between lexicogrammar and gesture: certain features of motion are construed lexicogrammatically in English, but gesturally in Spanish; and the other way around. While Lantolf's framework is not derived from systemic functional linguistics, it is quite compatible (cf. contributions in Byrnes,

[11] See: http://www.australiancurriculum.edu.au/Glossary?a=E&t=multimodal+texts

2006), and it is a valuable demonstration of the importance of investigating lexicogrammar 'ecologically' instead of treating it as an 'autonomous' system: clearly, lexicogrammar and gesture evolved together as complementary systems in both English and Spanish[12]. An early systemic functional contribution to the study of language and gesture is Muntigl (2004), and the systemic functional work on language and gesture has been followed up by Hood (2011).

If we take 'text' to mean an instance of the system of a language operating in a context of situation, then we can ask: (1) how it relates to instances of other semiotic systems operating in the same context of situation, and (2) how semiotic labour is divided among these different semiotic systems – how they complement one another. Taking gesture as an example of a semiotic system operating alongside language, we can represent the basis for these two questions diagrammatically: see Figure 1-15. Both language and gesture operate in the same context of situation, and are thus coordinated within it: as spoken text unfolds, it is accompanied by gesturing. Studies like those referred to above show that speakers (and addressees) are very adept at synchronizing speech and gesture so that gesture may relate to any of the metafunctional strands of meaning that run through the spoken text, e.g. a beckoning gesture accompanying a command to the addressee such as *Come here!* (interpersonal), a pointing gesture accompanying a reference (exophoric) such as *That's huge* (textual), or a depictive gesture accompanying a description such as *It is shaped like a five-pointed star* (experiential). (2) The division of semiotic labour between language and gesture seems to vary from one language to another, as already noted; but with a given language, there will also be variation

d
la

Fig. 1-15 Gesturing accompanying linguistic meaning and wording in context

[12] This is a fundamental insight into lexicogrammar – and it is very relevant in the context of the extensive body of research on the lexicogrammar of motion, going back to Talmy's (1985) pioneering typological study. (For a recent contribution to, and review of the extensive literature in, this area, see Beavers, Levin & Thao, 2010.)

according to the nature of the context, and thus of the register. In a game of charades, or in pantomime more generally, gesture and other visually realized semiotic systems must, of course, take over entirely from language; but in contexts where speakers need their hands and arms for other purposes, as often happens in 'doing' contexts, they have to rely more on language, and even though speakers tend to gesture when they engage in a telephone conversation, their listeners do not see these gestures (unless there is a video link).

One interesting issue that relates to both questions posed above is to what extent the different semiotic systems operating in context are integrated with one another and to what extent they operate independently of one another. To explore this issue, we can posit a **cline of integration**, extending from completely integrated systems to completely independent ones (cf. Matthiessen, 2009a). An example of complete integration is grammar and intonation. In English, and in many other languages, intonation is in fact not a separate semiotic system but rather a medium of expression deployed within the interpersonal and textual systems of the language (see e.g. Halliday, 1967a; Halliday & Greaves, 2008). An example of more or less complete independence would be the use of images in early printed books in Europe: illustrations were not produced or even chosen by the author as illustrations of points in the text but were instead added by the printer as decorations.

Another interesting issue is to what extent different semiotic systems extend all the way along the cline of instantiation from the instance pole to the potential pole (cf. Figure 1-11). We can ask of any one given semiotic system how systemic it is – which clearly relates to the question of how much individual variation there is across a speech fellowship (or speech community). Language has evolved as a fully systemic semiotic system: it is possible to posit and describe the overall meaning potential for a given language, interpreting this meaning potential as an aggregate of registerial subpotentials. However, it is theoretically quite possible that certain other semiotic systems are more usefully interpreted as operating with systems located somewhere midway along the cline of instantiation; in other words, they are most usefully described in register-specific terms (cf. Halliday, 1973: Chapter 4; Matthiessen, 1990). For example, if we consider semiotic systems that have been included under the heading of 'visual semiotics', we can note how highly contextually adapted and specialized systems such as technical drawing, mass transport route cartography and press photography are; it is not immediately clear that they can all be regarded as registerial subsystems of a general visual semiotic system (cf. Bateman, 2008; Matthiessen, 2009a).

1.5 The location of grammar in language; the role of the corpus

1.5.1 Recapitulation: locating the present work on the map of language

This is not exactly a recapitulation; rather, the aim is to locate the present work in relation to the dimensions of language discussed in the previous section.

In terms of **stratification**, the book deals with lexicogrammar, the stratum of wording. If we use the familiar metaphor of vertical space, as implied in the word 'stratum', the stratum 'above' is the semantics, that 'below' is the phonology. We cannot expect to understand the grammar just by looking at it from its own level; we also look into it 'from above' and 'from below', taking a **trinocular perspective** (Halliday, 1978: 130–131; 1996). But since the view from these different angles is often conflicting, the description will inevitably be a form of compromise. All linguistic description involves such compromise; the difference between

a systemic description and one in terms of traditional school grammar is that in the school grammars the compromise was random and unprincipled whereas in a systemic grammar it is systematic and theoretically motivated. Being a 'functional grammar' means that priority is given to the view 'from above'; that is, grammar is seen as a resource for making meaning – it is a semanticky kind of grammar. But the focus of attention is still on the grammar itself.

Giving priority to the view 'from above' means that the organizing principle adopted is that of **system**: the grammar is seen as a network of interrelated meaningful choices. In other words, the dominant axis is the paradigmatic one: the fundamental components of the grammar are sets of mutually defining contrastive features (for an early statement, see Halliday, 1966a). Explaining something consists not in stating how it is structured but in showing how it is related to other things: its pattern of systemic relationships, or **agnateness** (**agnation**, a term introduced into linguistics by Gleason (1965: 199) based on Latin *agnatus* 'related on the father's side').[13]

Each system has its point of origin at a particular **rank**: clause, phrase, group and their associated complexes. Since the clause is the primary channel of grammatical energy, the first part of the book deals with systems of the clause. The second part deals with systems at other ranks; and also those of the information unit, which is the grammatical reflex of the phonological tone group. The final chapter will describe movement across the rank scale, one of the forms taken by grammatical metaphor.

Systems at every rank are located in their **metafunctional** context; this means, therefore, that every system has its address in some cell of a **metafunction-rank matrix** (see e.g. Halliday, 1970/2005: 169; 1973: 133; 1976a; 1978: 132), as shown schematically in Figure 1-16 and in more detail in Chapter 2, Table 2-8. For example, the system of MOOD, referred to above, is an interpersonal system of the clause; so it is located in the 'clause' row, 'interpersonal' column in the matrix.

Structure is analysed in functional terms, explaining the part played by each element in the organic configuration of the whole. We shall see later on that the configurational view of structure is oversimplified, if not distorted, because the way linguistic units are structured tends to vary according to metafunction (see Halliday, 1979; Martin, 1996; Matthiessen, 1988). But it is possible to reduce all types of structure to a configurational form, as a strategy for exploring the grammar.

Figure 1-16 provides a map of this general conceptual framework. It also shows the dimension of **instantiation**; and this is the route by which we return to the text. In preparing this new edition we have made considerable use of a corpus, to check the details and extend the scope of the description; and also as a source of authentic examples. Whenever we shift our perspective between text and system – between data and theory – we are moving along this instantiation cline. The **system**, as we have said, is the potential that lies behind the **text**.

[13] This was true of the earlier editions also. But there the grammar was *presented* in the form of structure, whereas in this edition we have introduced the category of system into the ongoing account.

Fig. 1-16 Matrices defined by the semiotic dimensions: the lexicogrammatical function-rank matrix in relation to the stratification-instantiation matrix

But 'text' is a complex notion. In the form in which we typically receive it, as spoken and written discourse, a text is the product of two processes combined: instantiation, and realization. The defining criterion is instantiation: text as instance. But realization comes in because what becomes accessible to us is the text **as realized** in sound or writing. We cannot directly access instances of language at higher strata – as selections in meaning, or even in wording. But it is perhaps helpful to recognize that we can produce text in this way, for ourselves, if we compose some verse or other discourse inside our heads. If you 'say it to yourself', you can get the idea of text as instance without the additional property of realization.

1.5.2 Text and the corpus

Text is the form of data used for linguistic analysis; all description of grammar is based on text. Traditionally this has been mainly 'virtual' text of the kind just described: examples made up by grammarians inside their heads to illustrate the categories of the description. The only 'real' text that was available was written text, and some notable grammarians of English, such as Otto Jespersen, made considerable use of written texts as sources of data.

In the late 1940s two inventions appeared which were to change the work of a grammarian: tape recorders, and computers. The tape recorder made it possible to capture spontaneous speech; the computer made it possible to store and access data in increasing quantities. Ten years later, when Randolph Quirk at the University of London and W. Freeman Twaddell at Brown University in Providence designed and began to implement the first corpuses of written text, they foresaw that the operation would soon become computerized. At the same time grammarians such as Halliday were recording natural speech and analysing it for intonation and rhythm (Halliday, 1963a and b, 1967a). We now have indefinitely large computerized corpuses of both written and spoken text (for recent overviews of corpus linguistics, see e.g. Cheng, 2011; McEnery & Hardie, 2012; for the use of corpora in the investigation of English, see McEnery & Gabrielatos, 2006; for the relationship to systemic functional research, see also Hunston & Thompson, 2006; Wu, 2009).

The text is typically presented in written form, on the screen or else printed as hard copy. If the original was written, its format is – or at least can be – preserved. If the original was spoken, it is usually transcribed into regular orthography; this has two drawbacks, one of omission (there is no record of intonation and rhythm) and one of commission (it is 'normalized' according to conventions designed to make it look as though it had been composed in writing) – thus for a grammarian it has rather limited value. It is still not automatically made available as speech.

The corpus is fundamental to the enterprise of theorizing language. Until now, linguistics has been like physics before 1600: having little reliable data, and no clear sense of the relationship between observation and theory. But precisely because the corpus is so important it is better to be aware of what is good about it, and also what is potentially not so good. Let us enumerate four points – three plusses and one minus – which relate particularly to our use of the corpus.

First, its data are **authentic**. This one property underlies all its other advantages. What people actually say is very different from what they think they say; and even more different from what they think they ought to say (Halliday, McIntosh & Strevens, 1964). Likewise, what people say or understand under experimental conditions is very different from what

they say or understand in real life (for example, children aged 4 to 5 were found, when they were probed, not to understand or be able to produce relative clauses and passives; whereas these appear regularly in the natural speech of children before the age of 2). The difference is less marked in writing, although it is still there. Would Jane Austen (or our own teachers in school) have acknowledged the double '*–ing*' form she used in *Mansfield Park*: '*But it would rather do her good after being stooping among the roses;*' ? (New York: Hyperion, n.d., p. 64). But it is in speech that authenticity becomes critical; and this leads us to the second point.

Second, then, its data include **spoken** language, ranging from fairly formal or at least self-monitored speech (as in interviews) to casual, spontaneous chatter. The reason this is so important is not, as people sometimes think, a sort of inverted scale of values in reaction against earlier attitudes that dismissed everyday speech as formless and incoherent; it is a more positive factor – namely, that not only is natural spoken language every whit as highly organized as writing (it is simply organized along somewhat different lines; Halliday, 1985a, 1987a) but, more significantly, it is in the most un-selfmonitored spontaneous speech that people explore and expand their meaning potential. It is here that we reach the semantic frontiers of language and get a sense of the directions in which its grammar is moving (cf. Halliday, 2002a).

There is another point that should be brought in here. Now that spontaneous speech is becoming available for study,[14] some grammarians propose to write separate grammars for it. This approach has the merit that it can highlight special features of spoken language and show that it is systematic and highly organized; but it tends to exaggerate the difference between speech and writing and to obscure the fact that they are varieties within a unitary system. Spoken and written English are both forms of English – otherwise you could not have all the mixed and intermediate forms that are evolving in electronic text. In my own work, including earlier editions of this book, I have always taken account of both, with a slight bias towards spoken language for reasons given above; I have wanted to preserve the underlying unity of the two. Either way, what matters is that spoken language can now occupy the place in linguistic scholarship that it must do if the theory is to continue to advance.

Third, the corpus makes it possible to study grammar in **quantitative** terms. It is clear by this time that grammatical systems are probabilistic in nature: that, for example, the system of POLARITY in English has to be modelled not simply as 'positive/negative' but as 'positive/negative with a certain probability attached' (which has been found to be of the order of 0.9 : 0.1).[15] Computerized parsing and pattern-matching is now reaching the point where quantitative studies can be undertaken of a number of primary systems in the grammar, using samples large enough to permit comparison among different registers (where it seems likely that probabilities may be systematically reset). Not enough work has yet been undertaken along these lines for us to build it in to the total picture; but it is a high priority field for future research. (The exploration of the probabilistic nature

[14] Though now that technological obstacles have gone, legal ones have arisen. If you record surreptitiously, you lay yourself open to being sued.

[15] See Halliday & James (1993); also Halliday (1993a), Nesbitt & Plum (1988), Matthiessen (1999, 2006a).

of language has been part of systemic functional linguistics from the start – in fact, since before the start: see Halliday, 1959. After decades of resistance in formal linguistics, there is now more general acceptance of the probabilistic nature of language, thanks to advances in both 'corpus linguistics' and 'statistical natural language processing': see e.g. Bod, Hay & Jannedy, 2003.)

What then is the problematic aspect of the large-scale corpus? Linguists who specialize in corpus studies tend to refer to themselves, rather disingenuously, as 'mere data-gatherers'. We doubt whether they truly deceive themselves; they are well aware of the theoretical significance of what they are doing and what they are finding out. But they may perhaps deceive others, encouraging them to believe that there is some disjunction between data-gathering and theorizing. It is just such a dichotomy that has hassled the linguistics of the past few decades, isolating the system of language from the text as if they were two different orders of phenomena.[16] Of course, new data from the corpus will pose problems for any theory, systemic theory included – as Steve Jones has said, 'a science without difficulties is not a science at all' (Jones, 1999: 152). But such data will not contribute towards raising our understanding unless cultured by stock from within the pool of theoretical knowledge.

We emphasize this because there was a strong current of anti-theoretical ideology in late twentieth-century thinking, at least in certain intellectual domains. This was part of a self-conscious postmodern reaction against 'grand designs'; as often happens in such shifts of fashion, what starts out as a steadying correction to the course of knowledge becomes a lurch to a position more extreme than that which it was correcting. All modelling becomes micromodelling, all categories become collections of instances. We share the commitment to data and to the study of small-scale phenomena, in semiotic systems as in systems of any kind. But to banish the macro and the system from one's thinking is simply to indulge in another kind of grand design; being 'atheoretical' disguises a particular theoretical conviction which in our view is ill-judged and ill-informed (cf. Halliday & Martin, 1993: Ch. 11). We would argue for a dialectical complementarity between theory and data: complementarity because some phenomena show up best if illuminated by a general theory (i.e. from the 'system' end), others if treated as patterns within the data (i.e. from the 'instance' end) (cf. on global warming, above); dialectical because each perspective interpenetrates with and constantly redefines the other. This is the kind of thinking we have tried to adopt throughout the present work.

[16] A 'corpus-based grammar' is fine; there is no excuse now for a grammar of a well-researched language such as English not to be corpus-based. A 'corpus grammar' would seem to be a contradiction in terms, if it means a grammar emerging by itself out of the corpus. Data do not spontaneously generate theory. Some corpus specialists now favour a 'corpus-driven' approach (cf. Tognini Bonelli, 2001). In the terms described, the present grammar would, we think, qualify as corpus-driven; the difference would lie in the relative weight given to a general linguistic theory (and to the place of theory scientific praxis). We make more use of the explanatory power of a comprehensive model of language. In this connection, it is also important to emphasize that the present grammar has been tested extensively against authentic text in a way most 'corpus-based grammars' never are: it has been applied in systematic and exhaustive analysis of large volumes of text.

1.6 Theory, description and analysis

In this chapter, we have outlined the contours of language in general – what we might call the 'architecture' of human language, using a common metaphor for the organization of a system, although a term such as 'anatomy' would arguably be more appropriate since language is an evolved system rather than a designed one. We have drawn our illustrations from English since one of the aims of this book is to present a description of the grammar of English that can be applied in the analysis of spoken and written English text, as illustrated in Halliday (1985c). However, it would be equally possible, and highly desirable, to do the same for any other language. Lexicogrammatical text analysis is an important tool to be used in addressing many problems in a community in different spheres such as education, healthcare, administration, and commerce. The number of systemic functional descriptions of languages spoken around the world has been growing steadily. In Caffarel, Martin & Matthiessen (2004), linguists present summary accounts of eight languages (German; French; Telugu; Vietnamese; Mandarin Chinese; Japanese; Tagalog; Pitjantjatjara), and these are part of the basis for cross-linguistic comparison and typology. Since then, linguists have added accounts of a number of other languages (including Danish; Spanish; Arabic; Òkó; Bajjika; Cantonese), and some of them have now been published as books.

In the next chapter, we will begin to sketch a description of the lexicogrammar of English based on the general theory of language we have sketched in this chapter. But let us first pause briefly to distinguish between **analysis** and **description**, and between description and **theory**. In empirical approaches to language, all three are grounded in data – in the first instance, in spoken and written text; but they differ in their relation to data – they differ in degree of abstraction from data.

When we observe a language, we observe it as text – as a flow of speech or as (typically) discrete pieces of writing. Texts lie at the instance pole of the cline of instantiation (Section 1.3.4), and once we have observed and collected them and made them accessible to study (e.g. by transcribing spoken text), we can proceed by **analysing** them, noting patterns in these instances. (i) If we have access to an existing account of the system of the language (at the potential pole of the cline of instantiation), then we will analyse texts by relating instantial patterns in the system. In other words, we undertake the analysis of texts by means of the **description** of the system that lies behind them, as in Figure 1-9 above, identifying terms in systems and fragments of structures that are instantiated in the text. In the course of undertaking the analysis, we are likely to find gaps in the description, or even mistaken generalizations. Text analysis is a very rigorous way of testing, and thus improving, existing descriptions because everything in a given text has to be accounted for in the description (cf. Matthiessen, 2007b: 791–792). (ii) If there is no description to draw on, this means that we will gradually have to develop one based on the analysis of a representative sample of texts (a corpus; see Section 1.3.4 above, and Chapter 2, Section 2.4). In other words, **describing** a language is a process of generalizing from the analysis of textual data. The outcome of this process is a **description** of the system of the language, and we keep testing such descriptions by deploying them in continued text analysis and by applying them to different tasks such as language education or natural language processing.

Analysis and description thus operate at the outer poles of the cline of instantiation within a given language. Regions intermediate between these two poles can be approached in terms of either analysis or description: the account of a text type can be interpreted as a

generalized analysis of a sample of texts, and the account of a register can be interpreted as specialized description of the general system; but, in either case, the account will ultimately be grounded in textual data.

While a description is an account of the system of a ***particular*** language, a **theory** is an account of language in ***general***. So we have descriptions of various languages such as English, Akan and Nahuatl; but we have a theory of human language in general (see e.g. Halliday, 1992e, 1996; Matthiessen & Nesbitt, 1996). This introduction to (systemic) functional grammar is both an introduction to the general theory of grammar and to the description of the grammar of a particular language, English. The theory includes the 'architecture' of grammar – the dimensions that define the overall semiotic space of lexicogrammar, the relationships that inhere in these dimensions – and its relationship to other sub-systems of language – to semantics and to phonology (or graphology). Thus, according to systemic functional theory, lexicogrammar is diversified into a metafunctional spectrum, extended in delicacy from grammar to lexis, and ordered into a series of ranked units: see Figure 1.17. This figure shows the general 'template' according to which the lexicogrammar of any particular language will be organized.

Fig. 1-17 Theoretical dimensions of lexicogrammar: hierarchy of rank, spectrum of metafunctions and cline of delicacy

However, this theoretical template does not include details that are specific to particular languages or even to large sets of languages. For instance, the description of the system of MOOD shown in Figure 1-9 is specific to English: according to this description, there is a system of INDICATIVE TYPE, with 'indicative' as its entry condition and 'declarative' and

'interrogative' as its two terms, and the term 'declarative' is realized by the sequence of Subject followed by Finite. This description is grounded in generalizations about English data, i.e. spoken and written texts; and all descriptions must be based on empirical evidence[17]. Naturally, when we develop descriptions for other languages, we may find similarities; for example, many languages can be described systemically in terms of a contrast between 'declarative' and 'indicative', although very few languages realize 'declarative' structurally by putting Subject before Finite, in contrast with Finite before Subject in 'yes/no' interrogative clauses (cf. Teruya *et al.*, 2007). Typological generalizations are both possible and desirable, serving many purposes (cf. Matthiessen, Teruya & Wu, 2008); but they are still grounded in empirical evidence, not based on theoretical hypotheses.

The template illustrated in Figure 1-17 will accommodate innumerable possible descriptions – a number of which correspond to languages actually spoken today. In Chapter 2, we will provide an overview map or matrix of the lexicogrammar of English in terms of metafunction and rank, and in the remainder of the book (Table 2-8), we will present a more detailed description, up to a certain point in delicacy – a point that is still a far distance from lexis, and even some distance from the region intermediate between grammar and lexis.

This introduction to (systemic) functional grammar differs in various ways from other accounts – in terms of both theory and description. (i) In terms of theory, we can locate systemic functional theory of grammar within a general family of functional theories of grammar, contrasting these with formal theories of grammar (cf. Halliday, 1977; Matthiessen, 2009b). Within the family of functional theories, systemic functional theory is unique in its paradigmatic orientation (see Section 1.3.2 above; cf. Halliday, 1966a) – its orientation to grammar as system, represented by means of system networks; other functional theories are syntagmatic in their orientation. Systemic functional theory also differs from many other functional theories in its emphasis on comprehensive, text-based descriptions – descriptions that can be used in text analysis; other functional theories have tended to foreground linguistic comparison and typology based on descriptive fragments from a wide range of languages.

(ii) In terms of description, this book is of course an introduction to a systemic functional description of the grammar of English – constituting one descriptive strand evolving among other ones in systemic functional linguistics (cf. Matthiessen, 2007b). This description may be compared with other descriptions of the grammar of English that have appeared over the past 500 years or so: see, e.g., Gleason (1965), Michael (1970) and Linn (2006). These descriptions naturally vary in many ways, e.g. relationship to theory (homogenous or heterogeneous ['eclectic']), relationship to corpus (cf. Chapter 2, Section 2.4), relationship to time (diachronic vs. synchronic, or some kind of synthesis), relationship to dialectal variation (what varieties of English are included), coverage of phenomena – from grammars of very selective coverage via grammars with a registerial focus (such as grammars of spoken English) to reference grammars, and relationship to intended users – ranging from language

[17] It is important to be clear whether technical terms are theoretical or descriptive in nature; theoretical and descriptive terms are distinguished systematically in Matthiessen, Teruya & Lam (2010).

learners to professional grammarians. Reference grammars are, in principle, the most comprehensive descriptions.

In the first half of the twentieth century, there appeared a number of single-authored reference grammars largely by Danish or Dutch grammarians – the most comprehensive being Otto Jespersen's *Modern English grammar on historical principles*, in seven volumes, which were published from 1909 to 1949 – completed by Niels Haislund after Jespersen passed away in 1943 during the Nazi occupation of Denmark. In the second half of the twentieth century, English-speaking grammarians took over the task of producing reference grammars, with Jan Svartvik and Stig Johansson representing the Nordic tradition, both of them making substantial contributions to corpus-based research into English grammar: Quirk *et al.* (1972, 1985), Biber *et al.* (1999) and Huddleston & Pullum (2002). The description of English grammar presented here is not designed as a reference grammar. However, unlike the recent reference grammars – or all previous ones for that matter, this description has been designed as one that can be used in text analysis (cf. Halliday, 1985c) – a task that imposes quite stringent demands on the description. Since the first edition of IFG appeared in 1985, researchers have analysed quite a large volume of registerially varied texts using the description presented here (see e.g. Matthiessen, 2007b: 824–830).

PTER

TWO

TOWARDS A FUNCTIONAL GRAMMAR

2.1 Towards a grammatical analysis

Let us take a passage of three sentences from the transcript of Nelson Mandela's speech and start exploring its lexicogrammar:

Text 2-1: Exploring – passage from Nelson Mandela's inaugural speech

To my compatriots I have no hesitation in saying that each of us is as intimately attached to the soil of this beautiful country as are the famous jacaranda trees of Pretoria and the mimosa trees of the bushveld.

Each time one of us touches the soil of this land, we feel a sense of personal renewal. The national mood changes as the seasons change.

We are moved by a sense of joy and exhilaration when the grass turns green and the flowers bloom.

Starting at the lexical end – with the 'content words' of the vocabulary – we find names of entities (persons and things), names of processes (actions, events, etc.) and names of qualities:

(1)	*names of entities*	
(a)	common names:	
	persons	compatriots
	things, concrete, general	soil, country, trees, bushveld, land, grass, flowers
	things, concrete, specific	jacaranda, mimosa
	things, abstract	hesitation, sense, renewal, mood, seasons, joy, exhilaration
(b)	proper names:	Pretoria
(2)	*names of processes*	
(a)	doing & happening	touches, change(s), bloom

	(b) sensing & saying	saying, feel, moved
	(c) being & having	have, is, are, turns
(3)	*names of qualities*	
	evaluative	beautiful, famous
	emotive	attached, intimately, personal, national

Here persons and things are named by nouns, qualities by adjectives (one in the form of an adverb: *intimately*) and processes by verbs. Verb, noun and adjective are grammatical **classes** – classes of word.

Word classes can be viewed 'from above' – that is, semantically: verbs typically refer to processes, nouns to entities and adjectives to qualities (of entities or of processes). They can also be viewed 'from round about', at their own level, in terms of the relations into which they enter (cf. Halliday, 1963c): paradigmatic relations (the options that are open to them) and syntagmatic relations (the company they keep). On either of these two axes we can establish relationships of a lexical kind (collocations and sets) and of a grammatical kind (structures and systems). Here are some we find construed in this text.

(i) syntagmatic/lexical [collocation]

change ... mood, season

grass ... flower, green

flower ... bloom

move ... sense ... joy, exhilaration

soil ... land, country

tree ... jacaranda, mimosa

say ... hesitation

country ... beautiful

The measure of collocation is the degree to which the probability of a word (lexical item) increases given the presence of a certain other word (the **node**) within a specified range (the **span**). This can be measured in the corpus. Thus, the word *season* occurs 1,000 times in the corpus of ten million words: this would give it a certain overall probability of occurrence. But it might be found that, given the node *change* and a span of ± 4 (that is, four words on either side), the probability of *season* occurring went up by a significant extent. This would mean that, if you hear or read the word *change* in a text, you have a heightened expectancy that the word *season* may be somewhere in the immediate neighbourhood. It is a significant feature of the meaning of the word *season* that it collocates with the word *change*. Such conditioning effects can, of course, be measured in both directions: both the increased probability of *season* in the environment of *change*, and the increased probability of *change* in the environment *season*.[1]

[1] The effects will be different if the words have notably different overall frequencies in the language: for example, a rare word *jacaranda* collocating with a common word *tree*.

Collocational patterns of this kind contribute significantly to the unfolding meaning of a text[2].

(ii) syntagmatic/grammatical [structure]

The words in these (or any other) sequences can be assigned to grammatical classes. Given that *the* is a determiner and *of* a preposition, in the first example we have determiner + adjective + noun + noun + preposition + noun. Such a sequence of classes is called a **syntagm** (e.g. Halliday, 1966a). However, this tells us very little about how it is organized or what it means. The significance of such a syntagm is that here it is the realization of a **structure**: an organic configuration of elements, which we can analyse in functional terms. *trees* denotes the category of entity being referred to; we designate its function as Thing. *jacaranda* denotes the class within this general category; it functions here as Classifier. *the* has a pointing out function, known as Deictic: it signals that some particular member(s) of this class is or are being referred to. *famous* is one of a special set of adjectives that occur straight after the Deictic, still contributing to the pointing function; we call these simply Post-Deictic. We then have to wait until after the Thing to find out which jacaranda trees are meant: it is those 'of Pretoria', with *of Pretoria* functioning as Qualifier. (We will leave out the analysis of its own internal structure for the moment.) So we can analyse as:

the	famous	jacaranda	trees	of Pretoria
Deictic	Post-Deictic	Classifier	Thing	Qualifier

There is a similar structure in *the mimosa trees of the bushveld*, except that it has no Post-Deictic, and the Qualifier is 'common' not 'proper'; and *the soil of this beautiful country* is again comparable – here there is no Classifier, while, on the other hand, the Qualifier is rather longer (showing that it can in fact contain a fully expanded nominal group). Each of the three nominal groups has the same functional outline, *the* + Thing + Qualifier with *of;* each happens to contain six words, but, more relevantly, in Mandela's speech each one contains (from the first Ictus) three feet:

[2] The notion of collocation was first introduced by J.R. Firth (1957) (but note Hoey, 2005), and gained wide acceptance, particularly in work based on corpus analysis, as in the Birmingham tradition, e.g. Sinclair (1987, 1991), Coulthard (1993), Hoey (2005) and Cheng *et al.* (2009). For further systemic functional accounts of collocation, see e.g. Halliday (1966b), Halliday & Hasan (1976: Section 6.4), Benson & Greaves (1992), Gledhill (2000), Tucker (2007) and Matthiessen (2009b); Matthiessen (1995a) relates collocational patterns to structural configurations such as Process + Medium, Process + Range, Process + Degree; Thing + Epithet (for a corpus-based study of Process + Degree, see Matthiessen, 2009b).

This combination of grammatical (functional) parallelism with phonological (rhythmic) parallelism foregrounds the grammatical pun in the word *attach* (*be attached to*, (1) mental process 'love', (2, 3) 'be rooted in' – which then sets one to rethinking (1) in terms of roots). The analysis points up how the interaction among patterns at different strata plays a significant part in the construction of meaning.

(iii) paradigmatic/lexical [the lexical set]

Paradigmatically, lexical items function in sets having shared semantic features and common patterns of collocation. Thus *tree*, *flower*, *grass* share the feature of being generic names of plants; the corpus might show that they have in common a tendency to collocate with names of colours, various forms of the item *grow*, and so on. Such sets are typically fuzzy, with doubtful or part-time members (for example, *bush*, *blossom*).

Word association tests carried out many years ago showed that people associated words along both axes: if asked what other words sprang to mind when they heard *tree* they would come up both with words that were related syntagmatically, like *green* and *grow*, and with words that were related paradigmatically like *grass* and *bush*. Of course, many pairs of words are related in both ways, like *tree* and *branch*; they enter into a systemic contrast (a *branch* is part of a *tree*), but they also collocate, as in *climbed up the branch of a tree*.

Typically, the semantic features that link the members of a lexical set are those of synonymy or antonymy, hyponymy and meronymy: that is, they are words that are alike or opposed in meaning, words that are subtypes of the same type (cohyponyms: *oak*, *palm*, *pine* ... as kinds of *tree*), or parts of the same whole (comeronyms: *branch*, *root*, *trunk* ... as parts of *tree*). Thus in the text by Nelson Mandela, in the environment of *trees*, *jacaranda* and *mimosa* suggest other flowering varieties of tree (and also resonate with *flower*); *personal* and *national* encompass other contexts of intermediate scope like *regional* and *familial*. *Pretoria* and *the bushveld* are parts of (*this beautiful*) *country* – its capital city and its open countryside; and they likewise suggest all the other parts. *joy* is related to *happiness*, *gladness* and *pleasure*, and also to its opposites *sadness* and *distress*. These 'absent' items do not need to be mentioned; they are part of the meaning of the items that are there in the text, virtually present once the relevant vectors have been established. Figure 2-1 illustrates the lexical relations set up within this passage.

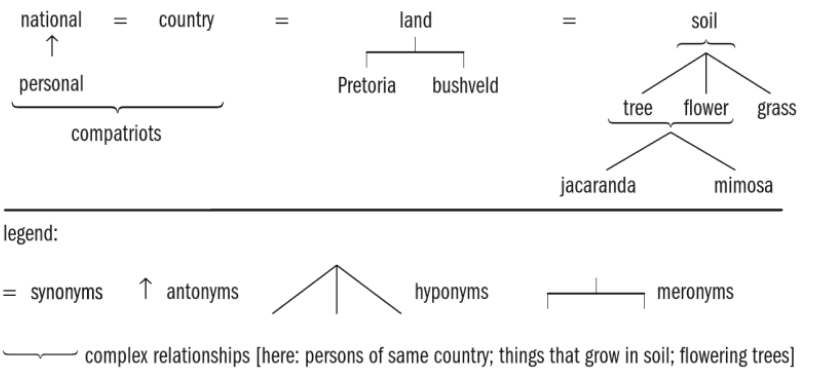

Fig. 2-1 Some lexical patterns in the Mandela extract (Note shift in sense of *land*: (1) = country, (2) = soil)

(iv) paradigmatic/grammatical [the grammatical system]

As discussed in Chapter 1, grammatical categories are organized in systems. For example, there is a system of PERSON, based (in English, as in most other languages) on the opposition of 'you-&-me' versus 'everyone (and perhaps everything) else', and then on that of 'you' as opposed to 'me' (see Figure 2-2):

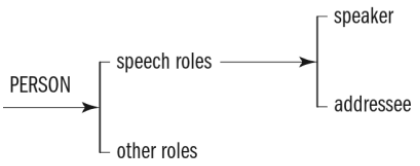

Fig. 2-2 The system of PERSON

This intersects with a system of NUMBER, opposing 'one' to 'more than one' (see Figure 2-3).

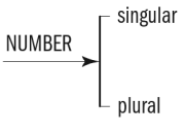

Fig. 2-3 The system of NUMBER

The way these two combine varies among different languages; for the version found in (modern standard) English, see Table 2-1 and, for further detail, Chapter 6.

Table 2-1 Systems of PERSON and NUMBER intersected

'First person plural' can mean 'more than one speaker', like a congregation in a religious assembly; more usually it means 'speaker + other person(s)'. In this case, Mandela is the only speaker (hence *I*); but the 'other persons' are made explicit as *my compatriots*, a semantically complex expression in which the meaning of 'possessive' *my* is defined by (the *com-* in) the word *compatriots* 'those who belong to the same nation'. Thus *my compatriots* means 'those who belong to the same nation as I do'; these are 'others' making up *we*. Here another grammatical system intersects (see Figure 2-4), contrasting collective *we* with distributive *each (one) of us*. In the environment of *compatriots* this means 'each person who belongs to the same nation as I do', and so this system grammaticalizes the opposition that is lexicalized as *national/personal*. Thus, through the resources of his lexicogrammar, Mandela construes each South African as an individual while at the same time linking them all inclusively with himself and with each other.

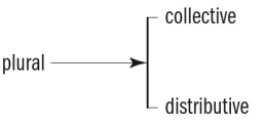

Fig. 2-4 The system of PLURALITY

Mandela is, of course, making a very large number of grammatical choices with every clause in his text. All of them contribute to the total meaning; certain ones among them may stand out as creating patterns that resonate strongly with the context of the occasion. As another example of this, we may note the way in which, in the course of these three paragraphs, Mandela combines selections in the system of process type in the clause (system of TRANSITIVITY; see Chapter 5) with the selections in the system of expansion in the clause complex (system of LOGICAL-SEMANTIC RELATIONS; see Chapter 7). These can be summarized in tabular form as shown in Table 2-2.

Table 2-2 Combinations of selections in process type and in expansion in Mandela's Inaugural

In the material clauses, the central participant is some part of the environment: *trees, soil, flowers* and the abstract *seasons*. In the mental clauses, it is *we/us* and the abstract *national mood*. The two sets of processes are then linked by a relationship of time or comparison 'as the same time as/in the same way as', with the abstract ones combining these two motifs. The overall effect of these highly patterned selections from within the grammar is to create a powerful sense of identity among individual, nation and physical environment, through a network of semantic relationships that we might represent diagrammatically as in Figure 2-5.

This brief commentary on a few sentences from Nelson Mandela's speech may serve to illustrate the lexicogrammar 'at work' creating meaning in the form of text (for further discussion of texts by Mandela, see Martin, 1999). We gain an insight into this process by placing some features of the wording, both lexical items and grammatical categories, in their syntagmatic and paradigmatic contexts in the system of the language. As the text unfolds, patterns emerge, some of which acquire added value through resonating with other patterns in the text or in the context of situation. The text itself is an instance; the resonance is possible because behind it lies the potential that informs every choice made by the speaker or writer, and in terms of which these choices are interpreted by listeners and readers. We refer to this ongoing creation of meaning in the unfolding of text as **logogenesis** (Halliday & Matthiessen, 1999: 18; cf. also Matthiessen, 2002b); this concept will become salient especially when we come to consider the relationship between different varieties

Fig. 2-5 Semantic relationships in Mandela text construed by grammatical systems

of text (for example, spoken and written) in later sections of the book (for a review, see Chapter 9, Section 9.1).

2.2 The lexicogrammar cline

We have stressed the ***unity*** of lexis and grammar, as the two poles of a single cline, or continuum; cf. Figure 2-6.

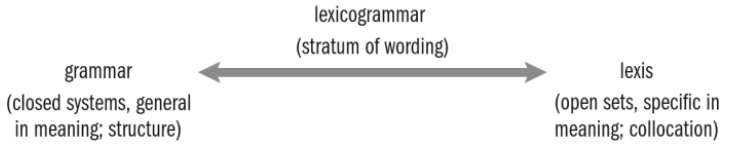

Fig. 2-6 The lexicogrammar cline

Because the two ends of the continuum are organized differently, when it came to describing them different techniques evolved: dictionary and thesaurus for lexis (Halliday *et al.*, 2004: Part one; Landau, 1989: Ch. 2; McArthur, 1986), and the 'grammar book' (typically, in the European tradition, syntactic constructions and morphological paradigms) for the grammar (e.g. Michael, 1970, and Linn, 2006). Either of these techniques may be extended all the way along the cline – but with diminishing returns: what the dictionary has to say about 'grammatical words' like *the* and *to* and *if*[3] is not very helpful; and while we can quite helpfully describe lexical items in terms of systems of features, the level of generality achieved is low. It is helpful because it shows how, when you choose a word, you are selecting among certain sets of contrasting features: for example, in Figure 2-7 the features that are 'lexicalized' (realized in words) are those of the degree of forcefulness of the order, the kind of authority behind it, and the positive or negative loading; lexical verbs

[3] i.e. words functioning as the direct realization of terms in grammatical systems.

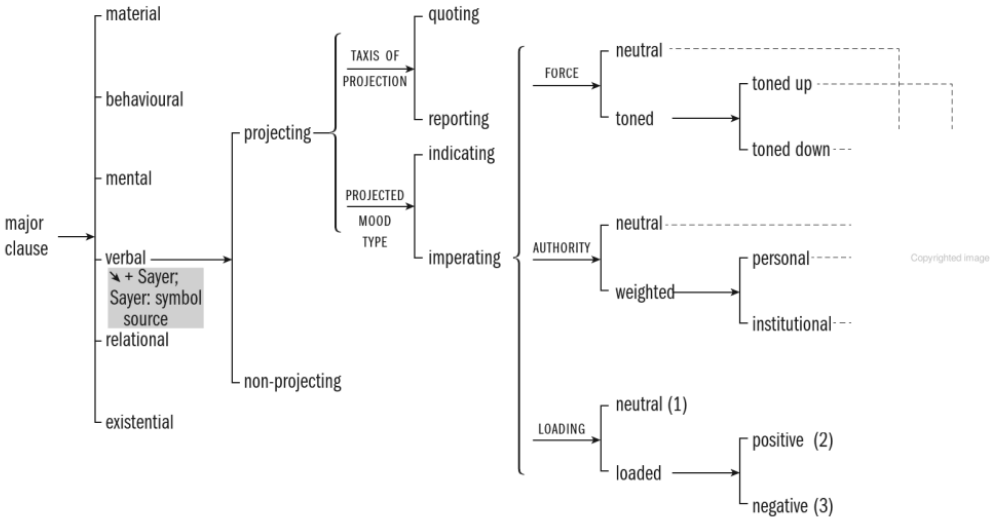

Fig. 2-7 Extension of the grammar of transitivity in delicacy towards lexis: a fragment of the lexicogrammar of the verbal process type (subtype 'imperating': see Table 5-25)

differentiated by these systems are given in Table 2-3[4]. Note that these are not definitive descriptions of these words; they are introduced here to show the principle of 'lexis as delicate grammar'. For other examples, see Hasan (1985a, 1987), Tucker (1998, 2007), and Halliday & Matthiessen (1999) and cf. Matthiessen (forthc.); and for the complementarity of grammar and lexis, see Halliday (2008: Chapter 2).

Table 2-3 Some verbs of saying differentiated by the delicate verbal process type systems in Figure 2-7

[4] To be more precise: these systems differentiate **senses** of lexical verbs. Such senses are typically given separate subentries in dictionaries. For example, the relevant sense of *tell* is 'order, instruct, or advise' rather than any of the many other senses of this verb such as 'narrate or relate (a tale or story)'. High-frequency lexical items always have more than one sense, typically many. Verb senses are likely to correspond to different process types or subtypes within one process type (cf. Matthiessen, forthc.).

The example we have just given is taken from the lexicogrammar of the clause, more specifically from the experiential system of PROCESS TYPE (see Chapter 5), which is located at the grammatical pole of the lexicogrammatical cline: taking one of the terms in this system, that of 'verbal' processes, we have illustrated how to extend the account in delicacy to the point where we can begin to discern lexical contrasts of the type exemplified in Table 2-3.

But if grammar and lexis are interpreted as the endpoints of a continuum, what lies in between them, around the middle? It is here that we locate those items that, on the paradigmatic axis, enter into series which could be regarded from both angles of vision: either, in a grammatical perspective, as rather large and fuzzy closed systems or, in a lexical perspective, as somewhat determinate and limited open sets. This would include, in English, things like prepositions, temporal and other specialized adverbs, and conjunctions of various kinds. Since this is a grammar book, we shall be treating these semi-grammatical patterns grammatically, although there will be space to give them only limited coverage.

But what sort of syntagmatic patterns would we locate around the midpoint of the cline? These are patterns that lie somewhere in between structures and collocations, having some of the properties of both (cf. Tucker, 2007). Consider patterns of wording such as the following:

take + pride/pleasure/delight + in + ... -ing

to make + things/matters + (even +) worse/more ...

waste/squander/spend + time/energy/money + on/in + ...-ing

Patterns of this mixed or intermediate type were brought to attention in the context of teaching English as a foreign language, notably in the work of A.S. Hornby (e.g. Hornby, 1954); and they have been investigated in linguistics within the framework of **construction grammar** (e.g. Kay & Fillmore, 1999; Fillmore, 2002). But it is thanks to the corpus that they can now be accessed and studied in detail; see especially Hunston & Francis (2000), also Partington (1998) – note Hunston & Francis' subtitle '[Pattern Grammar:] a corpus-driven approach to the **lexical grammar** of English' (our emphasis).

Patterns of this kind raise two distinct questions (which, however, sometimes get mixed up): How do we find them? How do we describe them? The answer to the first is provided by the corpus; patterns of this kind are difficult to retrieve by introspection, which is in any case unreliable, and require more reading and listening than linguists have usually been able or willing to devote; but they can be brought into visibility fairly easily by the use of standard concordancing techniques.

The answer to the second will, of course, depend on the grammatical theory. In systemic terms they appear as moderately delicate choices in the grammar, typically in transitivity and its related systems, having complex realizations involving both grammatical and lexical selections: see Tucker (2007). They are, on the whole, at a greater degree of delicacy than we shall be dealing with in this book; but several such patterns are brought in to the discussion (see Chapters 5, 7, 10) in order to show where they are located – their place in the overall map – and how they are approached from the grammar – how they are theorized as part of the total system of a language. We shall also see some of the problems they

present, especially the complex manner in which they may be realized (e.g. a pattern that is selected as a system of the clause but realized disjointedly throughout an entire clause complex/nexus).

Hunston & Francis (2000) refer to their object of study as 'lexical grammar', suggesting (appropriately) that these are mixed patterns which they are viewing in a lexical perspective. Such patterns have also been explored under the heading of 'phraseology' (e.g. Greaves, 2009; Cheng *et al.*, 2009), and Tucker (2007) outlines a systemic functional account of this intermediate region of the lexicogrammatical continuum.

In our grammatical perspective they appear, as we have said, as medium delicacy grammar – the most delicate grammar being lexis itself. If we maintain the grammarian's viewpoint all the way across the cline, lexis will be defined as grammar extended to the point of maximum delicacy. It would take at least a hundred volumes of the present size to extend the description of the grammar up to that point for any substantial portion of the vocabulary of English; and, as we have noted, the returns diminish the farther one proceeds. Nevertheless, there are certain contexts in which such a systemic interpretation of lexis can be valuable (cf. Hasan, 1985a, 1987; Tucker, 1998); and even in a general account of the grammar it is important to maintain a comprehensive picture that will show the relation between choice of words (lexical items) and choice of grammatical categories – especially in view of the complementarity between these two (see next section).

2.3 Grammaticalization

Consider a series of expressions:

didn't succeed ... nowise (or in no way) succeeded ... failed

These all contain a semantic feature 'negative'; but it is construed in different ways – at different points along the lexicogrammatical cline. In *failed* the 'negative' is **lexicalized**; in *didn't succeed* it is grammaticalized; while in *nowise succeeded* it is semi-grammaticalized – it is construed somewhere around the mid-point, with *nowise* occurring in an extended system including *never, at no time/seldom, rarely/hardly ever/nowhere/...*

These expressions also contain a semantic feature 'past time'. But this is **grammaticalized** in all three, with the v^d form (the past tense form) of the verb (*failed, succeeded, did*; see Chapter 6).

Imagine a 'language' in which all meanings were construed lexically. There would be different words for every lexical item in every grammatical category; not just positive/negative, as in *succeed/fail, trust/distrust* (with words meaning 'not see', 'not go', 'not know', 'not ...') but also for every tense and person of every verb, singular and plural of every noun (like *person/people*), comparison of every adjective (like *good/better*), and so on. Such a 'language' would need billions of different words; in other words, it would be impossible – impossible to learn as a system and impossible to process as text, except on so limited a scale that functionally it would not be a language at all.

It is a necessary condition of language that some meanings should be grammaticalized. These need not be the same meanings in every language; and in fact they are not – categories of place, time, size, value, number, sex are quite variable in this respect. In any case there are, as we have seen, differing degrees of grammaticalization; it is not an all-or-nothing

phenomenon. Nevertheless, some meanings (polarity is one of the most clear-cut) always seem to be grammaticalized, and many others are extremely likely to be so.

If a meaning is 'grammaticalized', this means that it is organized in the language (i) as a closed system of mutually exclusive terms, (ii) associated with some general category, and (iii) displaying proportionality throughout. For example: (i) positive/negative, singular/plural, past/present/future; (ii) a feature of all clauses, all count nouns, all verbs; (iii) [clause$_1$] positive : negative :: [clause$_2$] positive : negative :: [clause$_3$] positive : negative :: ... ; and likewise throughout. These are the three properties – **closure**, **generality** and **proportionality** – that characterize a grammatical system.

Grammaticalization is not dependent on how the categories are realized. They may be realized in a variety of ways: a change in the form, articulatory or prosodic, of some word or words; an addition of some element, to a word, a group or a clause; a change in the order of words, groups, or clauses. The realization may not be the same for all categories or in all environments; but it will be systematic in some way in the majority of cases, enough to establish and maintain the proportionality – with only a minority of 'exceptions' (which are likely to include some of the more frequent items). For example: past tense in English is usually realized as +(*e*)*d*, in so-called 'weak' verbs; second, there may be a change of vowel, in 'strong' verbs; some have change of vowel plus final -*d*/*t* (e.g. *think/thought*, *do/did*, *mean/meant*); and a tiny number are lexicalized (*am/was*, *are/were*; *go/went*). The varied nature of these realizations has no effect on the proportionality, which is exact: the relationship of 'past' to 'present' remains constant throughout.[5]

A systemic grammar is one that is organized around this concept of grammaticalization, whereby meaning is construed in networks of interrelated contrasts.[6] The term 'grammaticalization' itself, however, is problematic; it foregrounds the sense of 'process' – something being ***turned into*** a grammatical system, and this obscures the point that it is the inherent nature of language to be organized in grammatical systems. Nevertheless, we can recognize grammaticalization as a process taking place in time – in fact, in three distinct dimensions of time (Halliday & Matthiessen, 1999: 17–18). (i) We can see it in **ontogenetic time** when we observe children's early language development, which is built around the creation of proto-grammatical and then grammatical systems (Halliday, 1975, 2004). (ii) We cannot observe it directly in **phylogenetic time**, the evolution of human language; but we can track examples in the history of particular languages (for example, secondary tenses and the passive voice in English; see Chapter 6; Strang, 1970). (iii) We can see it in **logogenetic time**, the unfolding of discourse, when a passage of some extent – a clause or more – is recapitulated in a single word or group (see Chapter 10). So when we talk of the 'system' of language, as the underlying potential that is instantiated in the form

---

[5] This does not mean that the meaning of each category remains constant across all ***grammatical*** environments. We shall see in Chapter 5, for example, that the meaning of present tense differs between material and mental processes. But this has nothing to do with how the tenses are realized.

[6] We use 'contrast' rather than the Saussurean term 'opposition' simply because the latter suggests that all contrasts are binary. Reducing systems of more than two terms to sets of binary oppositions, which is always possible as a formal operation, can be arbitrary and misleading semantically. For example, the English tense system past/present/future patterns in some respects as past/non-past, in some as present/non-present and in some as future/non-future (but in most respects as a system of three terms). Any form of binary representation of such a three-term system is equally arbitrary.