

The Linguistic Analysis of Jokes

This book starts from three observations. First, the use of humour is a complex, puzzling, and idiosyncratically human form of behaviour (and hence is of scientific interest). Second, there is currently no theory of how humour works. Third, one useful step towards a theory of humour is to analyse humorous items in precise detail, in order to understand their mechanisms.

The author begins by considering how to study jokes rigorously: the assumptions to make, the guidelines to follow, and the pitfalls to avoid. A critique of other work on humour is also provided. This introduces some important concepts, and also demonstrates the lack of agreement about what a theory of humour should look like. The language devices used in various jokes, such as puns or humour based on misinterpretation, are analysed in detail. The central part of the book develops, and demonstrates, proposals for how best to analyse the workings of simple jokes. Finally, the author makes some general suggestions about the language devices that seem to be central to the construction of jokes.

The Linguistic Analysis of Jokes will be invaluable for researchers and advanced students of humour research, linguistics, and cognitive science.

Graeme Ritchie has been carrying out research in artificial intelligence and computational linguistics since 1973, investigating topics such as morphology, parsing, semantics, and creativity. In recent years, he has helped to pioneer the computer modelling of verbal humour. He is a senior lecturer in the School of Informatics at the University of Edinburgh.

Routledge Studies in Linguistics

1. Polari – the Lost Language of Gay Men

Paul Baker

2. The Linguistic Analysis of Jokes

Graeme Ritchie

The Linguistic Analysis of Jokes

Graeme Ritchie



LONDON AND NEW YORK

First published 2004
by Routledge
11 New Fetter Lane, London EC4P 4EE
Simultaneously published in the USA and Canada
by Routledge
29 West 35th Street, New York, NY 10001

Routledge is an imprint of the Taylor and Francis Group
This edition published in the Taylor & Francis e-Library, 2004.

© 2004 Graeme Ritchie

All rights reserved. No part of this book may be reprinted or reproduced or utilized in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

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
A catalog record for this book has been requested

ISBN 0-203-40695-8 Master e-book ISBN

ISBN 0-203-35064-2 (OEB Format)
ISBN 0-415-30983-2 (Print Edition)

To Regina

This page intentionally left blank.

Table of Contents

-

[Cover](#)



[The Linguistic Analysis of Jokes](#)

-

[Contents](#)

-

[Acknowledgements](#)



[1 Introduction](#)



[2 Assumptions and methodology](#)



[3 Linguistic preliminaries](#)



[4 Incongruity and its resolution](#)



[5 Two models of incongruity-resolution](#)



[6 The General Theory of Verbal Humour](#)



[7 Joke similarity and identity](#)



[8 Manipulating interpretations](#)



[9 The structure of puns](#)



[10 Some computational studies](#)



[11 Pragmatic and discourse issues](#)



[12 Speculations on joke structure](#)



[13 Future directions](#)



[Appendix A Formal definitions](#)

-

[Appendix B Examples](#)

-

[Glossary](#)

-

[References](#)

-

[Index](#)

This page intentionally left blank.

Acknowledgements

This book was written while I was on research leave from the University of Edinburgh. Most of the work was carried out during two spells in the Language Technology Group, Division of Information and Communication Sciences, Macquarie University, Sydney. Further writing took place in the Alfa Informatica group, University of Amsterdam, and the Department of Computing, City University, London. I would like to offer my thanks to my hosts for these visits: Robert Dale in Sydney, Henk Zeevat and Remko Scha in Amsterdam, Geraint Wiggins in London.

Thanks are also due to Greg Michaelson for his helpful comments on a draft of this book, and to Kim Binsted, without whom none of this would have happened.

I am grateful to the Leverhulme Trust for supporting this work with a research fellowship during the academic year 2001–2002.

Graeme Ritchie
Edinburgh, 2003

This page intentionally left blank.

1

Introduction

We set out the main aims of this book, and outline what its chapters contain.

1.1

Humour – a complex problem

The use of humour is a complex and intriguing aspect of human behaviour. Nash goes as far as to claim that ‘humour is a specifying characteristic of humanity’, comparable in this respect to ‘the power of speech, the mathematical gift, the gripping thumb’ (Nash 1985:1). Humour is present throughout social conventions and cultural artefacts, and the use of humour is highly valued in interactions between people. Despite this apparent importance, there is, as Raskin has observed, currently no major theory of humour, in the sense of something which ‘explains *what is funny, why it is funny, how it is funny, when it is funny, and to whom it is funny*’ (Raskin 1998:3, italics in original). This statement might seem to conflict with the literature on humour, where there are many proposals for dealing with humour, often claiming to be theories (e.g. Veatch (1998), Berger (1998), Latta (1999)). However, most of these ‘theories’ rarely define their basic terms formally, and are insufficiently developed to make precise falsifiable predictions. That is, they are at best interesting informal discussions, but are not formal theories or models. Nevertheless, we shall follow common practice and continue to refer to these ideas as ‘theories’.

In contrast, the newer disciplines of artificial intelligence and cognitive science deal in formal, computationally testable models of human behaviour, and have investigated many complex phenomena, such as visual perception, physical mobility, language use, and reasoning. However, they have barely touched on humour. Two major encyclopaedias in these disciplines, Shapiro (1992) and Wilson and Kidd (1999), have no sections or index entries for ‘humo(u)r’.

This book tries to bring together these strands of work – humour research and artificial intelligence – by studying the kinds of humorous phenomena which philosophers and others have already discussed, but applying as far

as possible the methodology of artificial intelligence to the task. There is also an intellectual parallel with modern generative linguistics, as will be explained in [Chapter 2](#).

1.2

What this book is about

Raskin's informal gloss (above) of what a theory would have to provide touches on the fact that humour has many aspects. Ultimately, research into humour must answer many disparate questions, such as:

- how did humour evolve as a behaviour?
- what are the social functions of humour?
- how does humour vary between cultures?
- how does the perception of humour stimulate physiological (and particularly neurological) reactions?
- what are the interrelations between humour, mood, and personality?
- what, if anything, is there in common between the types of entity (situations, events, pictures, texts, etc.) which are viewed as humorous?

It would be unrealistic, therefore, for this book to claim to provide a complete theory of humour: that is a vast task which will take many decades, perhaps centuries, to complete. We will not even consider most of the profound questions listed above. Instead we will make a contribution to the last of these questions: *the nature of those things that are considered humorous*.

We will narrow the task down even further by restricting our attention to humour conveyed in language – *verbally expressed humour* – and also by considering only *jokes*, rather than more general types of humour. These simplifications will be discussed more fully in [Chapter 2](#).

We will not be putting forward a complete theory even of the limited subarea of humour that we have chosen as our focus. Work on humour has, in our opinion, suffered from premature universal theories, in which the detailed analysis of data and hence the accumulation of evidence has been rather neglected. Authors tend to either eschew theoretical aims altogether (e.g. Alexander (1997)), or present a grand theory, usually stating it in a few sentences of ordinary language (e.g. Latta (1999)). In contrast, we will proceed by examining small classes of jokes and attempting to find generalizations. We summarize our findings in [Chapter 12](#), but these are presented as preliminary empirical results, not as a general predictive theory.

To be more exact, we will:

- advocate an approach to humour theory based on detailed description of particular subclasses of humorous phenomena;

- illustrate how various classes of jokes can be analysed relatively formally, focussing particularly on the *linguistic mechanisms* involved;
- provide a critical review of selected previous work;
- survey the work on computer models of humour, drawing some general conclusions about this subarea;
- offer some tentative suggestions about the main information factors in simple jokes.

This book can be seen as comprising three main sections: clarification of methodology (Chapters [2](#) and [3](#)), discussion of some past work in the field of humour research ([Chapters 4](#) to 7), and analysis of the (linguistic) structure of certain classes of joke ([Chapters 8](#) to 11). We also include some speculation, in [Chapter 12](#), on what a general account of joke structure might look like.

Much of our discussion, particularly in the earlier chapters, goes back to first principles in order to argue through theoretical assumptions and basic working practices. This is because these foundational issues have not, we believe, been fully explored, and later debates will be confused unless we are clear about our framework.

1.3 Related work

Given the existence of many texts surveying the history of work on humour (see [Section 1.4](#) below), we shall not attempt a comprehensive literature review here. However, in order to give a rough idea of the territory to be explored it is useful to mention some previous work which seems to have similar aims.

The most direct historical predecessor of the work reported in this volume is the computer modelling of simple verbal jokes, which will be reviewed in [Chapter 10](#), but we are aiming for greater generality and more theoretical abstraction, with less concern for the actual implementation of software.

In some respects, the closest work to ours is Attardo and Raskin's Semantic Script-based Theory of Humour/General Theory of Verbal Humour (SSTH/GTVH), which will be discussed at length in [Chapter 6](#). That research is similar in that it focusses on verbally expressed humour, and sets out to develop an abstract formal account. One difference is that Attardo and Raskin are proposing a general theory, whereas we are concentrating more on a methodology for starting to develop such a theory, together with preliminary proposals for some of the concepts which might go to make up the theory. That is, Attardo and Raskin are engaged in 'top-down' theorizing, whereas the current text focusses on 'bottom-up' analysis (in the terminology of Binsted and Ritchie (1996)).

Hetzron presents suggestions which are close in spirit and aims to those developed here, although he explains that he has in mind literary and rhetorical theory as his precedents (rather than formal linguistics and artificial intelligence as here). He argues that ‘breaking down “incongruity” and the like into more concrete categories is of necessity’ (Hetzron 1991:69). Probably as a result of using different disciplines as orientation points, Hetzron’s proposals are less formal than ours.

As noted above, there are many books and articles putting forward views about humour, often with an emphasis on the verbally expressed variety. Generally, such research discusses various intuitions about aspects of jokes with neither precise definitions nor the presentation of any abstract theory. Although we do not dispute that some of these intuitions may be plausible, and the discussions sometimes persuasive, our aim is to frame statements more formally and to offer more detailed arguments to support our theoretical constructs.

These brief remarks do not exhaust our commentary on past research. In subsequent chapters we shall discuss other work where it is relevant to the topic under consideration.

1.4

Reviews, surveys, and collections

Although we are not offering a survey of humour research, we will mention some texts that may be of particular interest, many of which contain good overviews and reviews of past work. The list is in chronological order.

Goldstein and McGhee (1972) bring together work by a number of authors on a variety of psychological topics, and present a bibliography and survey of trends within psychological research into humour. Chapman and Foot (1976) are the editors of another multi-author work, covering a very wide range of psychologically oriented topics. It should not be confused with Chapman and Foot’s 1977 publication which is a large conference proceedings containing many rather short papers on humour, from various disciplines. Wilson (1979) reviews much of past work on humour, as well as presenting a few experiments exploring the presentation of jokes. McGhee and Goldstein (1983) have chapters, mainly of a review nature, by various authors, on a variety of topics, from *linguistics and humor to the role of arousal and hemispheric lateralization in humor*. Raskin (1985) – a text whose importance will be discussed in later chapters – starts with a survey of humour research which covers the classic positions. Morreall (1987) is a collection of seminal readings in the philosophy of language, from Plato through Freud up to the 1980s, each chapter having a brief introduction from the editor. Chiaro (1992) is a brief, readable introduction to various forms of verbally expressed humour, with little detailed analysis. Attardo (1994), which starts with a good historical review and discussion of various aspects of verbally expressed humour, merits a special mention, as there

will be copious references to that text in our succeeding chapters. Ross (1998) is a textbook which tries to get readers to hone and articulate their intuitions about humour by following simple exercises.

1.5 Overview of contents

It may help to give a brief overview of the contents of later chapters. [Chapter 2](#) sets out our methodological aims and assumptions, and [Chapter 3](#) discusses the relationship between verbally expressed humour and conventional linguistic constructs. [Chapter 4](#) examines some frequently suggested ideas about verbally expressed humour, and [Chapter 5](#) examines two particular processing-oriented models based on those ideas. [Chapter 6](#) provides a critique of a very influential proposal about verbally expressed humour, namely the General Theory of Verbal Humour. [Chapter 7](#) briefly considers whether it is feasible to define notions of ‘similarity’ or ‘identity’ for jokes. [Chapter 8](#) examines some ways that information can be presented in jokes, and [Chapter 9](#) examines the area of puns, proposing some formally distinct subclasses. [Chapter 10](#) reviews various pieces of computer-implemented work in the area of simple puns and acronyms. [Chapter 11](#) explores some humorous phenomena which, while closely related to the issues discussed in earlier chapters, are more to do with the pragmatic aspects of language. [Chapter 12](#) sums up the earlier material, and offers some speculations on the components of verbally expressed humour, while [Chapter 13](#) suggests how our proposals could be developed further.

Although we advocate the injection of formality into humour research (see [Section 2.5](#)), we have confined our mathematical definitions to the appendices. [Appendix A](#) gives formal definitions for our assumed linguistic model (see also [Chapter 3](#)) and for some of the joke-presentation schemes discussed in [Chapter 8](#). It should be possible to follow the less formal discussions in the main chapters independently. [Appendix B](#) lists all the example jokes (and non-jokes) used in the main chapters. The appendices are followed by a glossary of technical terms and abbreviations.

Data (usually jokes) have been taken from various sources: joke books, Internet web pages of jokes (between December 2001 and March 2002), examples cited in academic articles, jokes heard on radio, TV, film or stage and recall of jokes heard in everyday situations. The latter are not given a citation. There are also *constructed items*, where a text has been created to illustrate a particular point.

Where we wish to show an example in phonetic form, we will use a textual representation based as closely as possible on the International Phonetic Alphabet, with the symbols between forward slashes; for example, /r^h/ for the sound of *red*.

2

Assumptions and methodology

We discuss the guidelines we shall adopt in this research.

2.1 Motivation

The exploration here takes it as axiomatic that it is worthwhile to try to understand how humour operates, since – as remarked in [Chapter 1](#) – the use, comprehension, and appreciation of humour form a complex and unexplained aspect of human behaviour. That is, the aim here is conventionally scientific: we want to illuminate the very large question of what humour is and how it works.

This is in contrast to any engineering motives which others might have (see, for example, [Section 2.13](#) below). The goal here is not to construct practical artefacts, such as computer programs, related to humour (although a good theory of humour would facilitate such construction – see [Section 2.12](#) for more discussion).

From the point of view of understanding human intelligence more generally, it is also desirable, particularly within the long-term aims of artificial intelligence, to study humour. If we ignore humour, we are omitting a central aspect of human activity from our programme to understand intelligence. If we could develop a full and detailed theory of how humour works, it is highly likely that this would yield interesting insights into human behaviour and thinking. Indeed, it is implausible to suppose that a theory of humour could be developed prior to, and in isolation from, broader theories of human activity. It is not so much that we should build a humour theory and then see how it fits in with intelligence; rather, we can use the attempt to theorize about humour to help to widen and deepen our study of human intelligence. There is no reason why we should exclude this one class of behaviour from our data.

2.2

Classifying theories of humour

As remarked in [Section 1.1](#), there are many proposals which claim to be theories of humour, from various disciplines. Humour is a extensive phenomenon, with many and varied manifestations – Gruner (1978), for example, lists Exaggeration, Incongruity, Surprise, Slapstick, The Absurd, Human Predicaments, Ridicule, Defiance, Violence, and Verbal Humour – and most theories of humour attempt to cover all varieties. Also, the content of these theories is surprisingly varied, to the extent that reviewers of the field cannot even agree how best to classify them.

Perhaps the most common classification is tripartite: *relief* or *release* theories, *superiority* or *aggression* theories, and *incongruity* theories (Morreall 1983, Raskin 1985). Attardo (1994:47) gives a similar clustering: *cognitive* (Incongruity and Contrast), *social* (Hostility, Aggression, Superiority, Triumph, Derision, Disparagement) and *psychoanalytical* (Release, Sublimation, Liberation, Economy). However there are other classifications. Wilson (1979) lists Relief, Conflict, Incongruity, Dualistic, Gestalt, Piagetian, Mastery, Freud's, and Apter and Smith's. Keith-Speigel (1972) classes theories into Biological/Instinct/Evolution, Superiority, Incongruity, Surprise, Ambivalence, Release/Relief, Configurational, and Psychoanalytical. Some of these finer classifications could plausibly be mapped on to the three broader categories, but the existence of these varied taxonomies emphasizes the lack of consensus not only about what the theory should be but also about exactly what existing theories say and how they differ.

What is noticeable about the many and varied discussions of humour is that they are not all considering the same aspect of humour. For example, those who consider aggression (or superiority) as the unifying essence of humour (see Fave *et al.* (1976) for some discussion) have little to say about why certain presentations of an idea are funny while other presentations, with equally aggressive content, are not. Minsky (1986:274–81) makes some preliminary remarks about how humour could be viewed from the artificial intelligence/cognitive science perspective, refining Freud's notion that humour is driven by our mental 'censors' which control inappropriate thoughts or feelings. Minsky suggests that jokes are illustrative examples of faulty logic, which enable humans to refine their reasoning censors in a relatively painless way. His remarks are thought-provoking, but say little about the internal structure of actual witticisms.

Hence several of these proposals could simultaneously be true. It could be that the driving force for humour is hostility, but various verbal devices, such as ambiguity, can be used to that end. As Raskin (1985:40) observes, there is no *a priori* reason why all aspects of every example of humour should be explicable in terms of a single principle. For example, the aspects of humour discussed in the current volume (namely, the language mechanisms used in jokes) are logically independent of the proposal that humour

is based upon superiority (Gruner 1978, 1997) – the role of aggression has no impact on the analyses and observations made in later chapters.

2.3 The need for description

If one wishes to develop a universal theory of humour, or even of verbally expressed humour, there are various ways that one could tackle this long-term goal. One approach, which could be termed *universalist*, is to devise an extremely general theory which is intended to cover all examples within the chosen area(s) of humour. It would then be possible to apply this general theory to particular instances of humour, or whole genres, demonstrating its aptness or perhaps falsifying it. This is an intellectually elegant and philosophically satisfying way of tackling the problem. Its drawback is the sheer difficulty of devising such an all-embracing theory, starting from the beginning. Although theory construction always has to begin with the data, the raw data consist of a vast and highly varied array of jokes, funny stories, riddles, puns, epigrams, slapstick scenes, cartoons, and so on. To cast a theoretical roof over such a broad church without reasonably detailed descriptions of the workings of the individual types of humour is very ambitious. Even if, as here, we restrict the scope of enquiry to just verbally expressed humour, the range of phenomena to be considered is unmanageably large. It would be very easy for a would-be general theory to fail either in its generality (i.e. not cover all the possibilities) or in its accuracy (i.e. cover some data incorrectly) or even in being a theory (by being so vague as to not make any meaningful predictions that could be falsified).

An alternative approach, which could be called *descriptive*, would be to examine particular types of humour thoroughly, developing detailed descriptions of the workings of the chosen genres of humour. The disadvantage of taking this route is that there could be a tendency for the various descriptions of humorous genres to be excessively tailored to the particular data, so that there are no commonalities between the accounts of different classes of joke. Without a unifying theory to suggest basic constructs, the researcher could be tempted to devise *ad hoc* mechanisms or formalisms for each case study. On the other hand, the advantage would be that the descriptions would genuinely reflect the data, and would be sufficiently detailed to provide falsifiable predictions. Further, more general, theorizing could be based on these descriptions, since the theoretician would be starting from analysed data displaying some relevant structure, rather than raw data (i.e. the unanalysed texts).

There is no right answer to the choice between universalism and descriptivism. In this book, we adopt the descriptivist line. This is partly because it is the only route that holds out any hope in the short term of computer testing (see [Section 2.12](#)), but also because the risks (of over-specialized

Table 2.1 An analogy between linguistics and humour research

<i>Linguistics</i>	<i>Humour research</i>
Strings	Grammatical texts
Grammaticality	Being a joke
Sentence type	Subclass of jokes
Structural description	Description of a joke
Grammar rule(s)	Definition for a joke class
Theory of grammar	Theory of jokes
Theory of language	Theory of humour

descriptions) are outweighed, we believe, by the benefits. These include not only falsifiability, but also the provision of richer material for further theorizing. This latter point is important. An analogy can be made here with the drive within Chomskyan linguistics to evolve a universal theory of language, or at least of grammar (see [Table 2.1](#)). In generative grammar, linguists take strings of words as the basic data, and attempt to formulate rules which will separate the grammatical (well-formed) strings from the ungrammatical. In the analogy here, joke-theorists take grammatical texts as the basic data, and try to devise definitions which will distinguish those texts which are jokes from texts which are not (see Sections [2.9](#) and [3.5](#) for further discussion). Although we will not be stating our formal definitions as generative rules, the role that the definitions play in this analogy is comparable to grammar rules, or sets of grammar rules, within syntactic theory.

The formal statements labelled, in [Table 2.1](#), ‘Description of a joke’ and ‘Definition for a joke class’ constitute the *abstract syntax* of jokes. That is, the definitions will deal in certain abstract entities which are proposed as underlying the actual texts of the jokes, and will express the ways that these abstract entities can combine or be related. These definitions will characterize what counts as a joke, but at a higher level than a grammar of the actual words involved.

Once we have developed a full description of the ways that jokes can be constructed (see Sections [2.11](#) and [2.14](#) below), a ‘theory of jokes’ will have been constituted. This will not, however, be a complete theory of humour, any more than a theory of grammar is a full theory of language and its use. However, much as a theory of grammar is a building block for a theory of language, a theory of jokes becomes part of, or informs work on, a theory of humour.

In linguistics, a theoretician’s generalizations about the nature of language are *not* based on, or phrased in terms of, simple unanalysed sentences (sequences of words). The putative theoretical laws are about structures, typically syntactic, which are themselves a result of analysing

actual sentences. The universal statements refer to more specific theoretical entities which have been posited to make sense of the way that individual languages, or areas of languages, operate. Generative linguistics was able to assume centuries of previous work, including a great deal of detailed grammatical description. The work of grammarians who formulate rules for the syntax of individual languages is not irrelevant to the grander endeavour of stating a universal theory of language; rather, grammarians are carrying out exploratory groundwork which allows further theorizing to happen.

2.4

Essentialism versus multiple components

Attardo (1994:1–2) discusses the notion of an *essentialist* theory of humour, namely the view that there is some ‘essence’ that is present in all humorous phenomena, and whose presence is a necessary and a sufficient condition for a text (or situation, or picture, etc.) to be humorous. This is in contrast to a viewpoint which holds that humorous entities are a ‘family’ or *prototypically defined* set, so there may be no one characteristic that they all have in common, although every member of the ‘family’ has something in common with some other member (cf. Wittgenstein (1967), Rosch and Mervis (1975)).

Most of the theories mentioned in [Section 2.2](#) are essentialist, in the sense that they argue for some particular ‘magic ingredient’ in all humour. Often, the evidence that an author cites in support of the chosen ingredient is at least plausible, which may suggest that many of these factors are indeed involved, though perhaps not uniquely, in humour. We adopt an open mind on what ingredients are relevant to humour, and how they may interact.

Working assumption: humour may be caused by a range of factors, and the relationships between these factors are yet to be determined.

It is important to distinguish logically between the notion of a characteristic being necessary and sufficient for some form of humour, and a characteristic being universally present. If one follows the descriptivist path outlined above, and studies small classes of humorous texts in detail, then it may be quite appropriate to seek some properties which are necessary and sufficient for membership of that *class* of humorous entities, while not subscribing to the belief that there are properties which are necessary and sufficient for humour in general. In other words, it is possible to apply a localized variant of essentialism, in which necessity/sufficiency are taken as the measures of success, but where the scope of the enquiry does not cover the whole of humour. This is the view adopted here.

2.5

The need for formality

Humour has been studied for a long time, originally by philosophers, subsequently by literary critics and essayists, and more recently by psychologists, sociologists and linguists. This long heritage of reflection and discussion has produced an accumulation of ideas, proposals and claims about the nature of humour, but, as noted in [Chapter 1](#), the vast majority of these ideas have been stated informally. What has not emerged from this background is a precise, detailed, rigorous theory of even a relatively limited subarea of humour.

Arguably, the most rigorous and precise work on humour has taken place within psychology, where there are many experimental results on topics such as the measurement of types of sense of humour (Ruch 1996), or the physiological reflexes of humour (e.g. Langevin and Day (1972), Godkewitsch (1976), Svebak (1977)). While these are undeniably scientific and have discovered interesting patterns, they have not yet led to an overall theory of humour, and certainly do not directly illuminate the area of interest here: the exact nature of texts which are deemed humorous.

It is necessary to consider what we mean by 'precise, detailed, rigorous' in this context. The two conceptual frames of reference for the discussion here are *generative linguistics* and *artificial intelligence*.

Generative linguistics (as founded by Chomsky in the 1950s) takes the view that language, a quintessentially human trait, can be examined with the same scientific precision and formality that traditional science would apply to physical phenomena. The aim of linguistics, from this perspective, is to develop a theory of the nature of language, but to do this by studying specific linguistic facts in detail. The descriptions that a linguist builds up should be sufficiently well defined that it would be possible for a machine to follow the grammatical rules (Chomsky 1975).

Artificial intelligence attempts to model human behaviour of various forms. Although the behaviour is usually characterized as 'intelligent', this term has to be defined more widely than the lay person might initially expect. For example, artificial intelligence has expended considerable effort in devising methods by which visual information (such as the eye supplies to the brain) could be interpreted in terms of actual world objects, a task which is not normally regarded as 'intelligent' when performed by humans. In developing these models, the behaviour in question must be analysed in some detail, and the rules or structures which constitute the model must be well defined and suitably detailed. A common yardstick for the degree of formality and precision needed is that it should be possible to write a computer program which would carry out the desired behaviour by following the model. The proposal for the historic Dartmouth Conference in 1956, often taken as the founding event in artificial intelligence, stated the conjecture:

that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.

(Charniak and McDermott 1985:11)

(Compare the criterion of precision attributed above to Chomsky).

The vast bulk of work on humour does not adhere to these standards, often because it falls within some other discipline.

One of the main motivations for this book is the need to import into humour studies some of the methodological guidelines and attitudes from generative linguistics (at least its declared principles, if not its practices, perhaps), and from artificial intelligence (preferably from the 'neater' end (Hayes 1987:183, Russell and Norvig 1995:25, fn. 17)). If we want to develop a sound theory of humour, we have to be more formal and more detailed in what we do. Computer modelling may be interesting in its own right (see [Section 2.12](#) for some discussion) but here it acts as a standard for judging the precision and formality of work on humour: could one write a program based on the analyses?

More generally, a formal theory of verbally expressed humour should meet certain criteria:

- (a) All terms should be clearly defined, so that there are only a relatively few primitive (undefined) notions and the relationships between all the concepts are made explicit and precise.
- (b) In seeking to explain empirical phenomena (e.g. jokes), theoretical statements should be framed in terms of the abstract formal constructs of the theory, not in terms of the entities being studied.
- (c) It should be developed in sufficient detail that it is not merely a renaming of existing informal insights.
- (d) It should be clear what would count as evidence for associating any abstract concepts with particular texts; that is, assignment of analyses to texts should not be a wholly subjective matter.
- (e) The theory should have a clear relationship to theories of language, in the sense that it should show the relationship between humorous and non-humorous texts, in terms at least of structure and possibly also of processing.
- (f) The theory should be falsifiable, in the sense that it is conceivable that there could be data which do not conform with the theory, and it should be made clear what would constitute such a counter-example.

Most of these guidelines are fairly conventional, but – as we will demonstrate in later chapters – they do not appear to be routinely followed within humour research. This is understandable given the early stage of development of the area, but these principles represent a standard towards which a scientific approach should aim.

2.6

Verbally expressed humour

This book is about verbally expressed humour (VEH), which has also been referred to as *verbalized* humour: ‘jokes which are expressed by means of a linguistic system (or its derivatives, like writing)’ (Attardo 1994:96). This term is used in contrast to the widely used phrase *verbal humour*, in order to emphasize that the humour need not be inherently concerned with verbal matters, in contrast to the way that a pun, for example, is essentially verbal. Although the term ‘verbal humour’ is used by some authors (e.g. Raskin (1985)) to mean what we are calling VEH, it is more commonly used to refer more narrowly to humour which crucially depends on the language elements, with the typical diagnostic being that verbal humour cannot be directly translated (see [Section 2.7](#) below); examples are (1) and (2). ([Appendix B](#) lists all the example items used in these chapters.)

(1) What do you get if you cross the Atlantic with the *Titanic*?

About half-way.

(Metcalf 1994:25)

(2) My brother’s a naval surgeon.

Wow, they do specialise nowadays, don’t they?

(Metcalf 1994:108)

Although some of the examples considered in our later chapters will be what are often known as ‘funny stories’ – very short narratives with amusing and usually surprising endings – verbally expressed humour (VEH) also includes puns (e.g. (3)), humorous riddles (e.g. (4)), amusing epigrams (e.g. (5)), etc., and even the retelling of amusing incidents or cartoons in cases where the retelling itself functions humorously.

(3) What do you get if you cross a sheep with a kangaroo?

A woolly jumper.

(4) What’s black and dangerous and sits in a tree?

A crow with a machine gun.

(5) Fashion is a form of ugliness so intolerable that we have to alter it every six months.

(Oscar Wilde cited by Metcalf (1994:79))

Occasionally we shall make reference to examples of humour in other media, purely in order to illuminate our discussion of VEH, but we are not attempting to analyse humour conveyed in other ways (e.g. visually).

It may seem cavalier to discard a wide range of data at the start, but the whole problem of humour is so large and complex that it will make the research more manageable if we can simplify matters slightly. However, verbally expressed humour manifests a wide range of humorous phenomena, so we have not simplified the issue to a trivial level.

Working assumption: verbally expressed humour embodies a sufficiently rich and varied set of humorous phenomena that it is worth studying.

2.7

Linguistic and propositional jokes

Verbal humour is usually contrasted with *referential* jokes (Attardo 1994: passim). Verbal humour is wholly dependent on the language used, and is regarded as not translatable into another language, whereas in referential humour, the amusing substance is in what the joke says – the events described, the characters, the situations – rather than the details of the language used. As Attardo explains, numerous authors from Cicero onwards have acknowledged this general distinction, although it is hard to be certain that they would all have drawn the boundary in exactly the same place (Attardo (2001:23) seems to equate puns and verbal jokes, for example).

Our exploration of jokes in later chapters does not accept this distinction as an axiom. In fact, we will develop analyses which group certain traditionally ‘verbal’ jokes along with certain ‘referential’ ones. The critical examples are jokes involving misinterpretation caused by linguistic ambiguity. This is because, as will be discussed in [Chapter 8](#), such mis-interpretations appear to constitute a major class of joke, regardless of whether the confusion is based on linguistic factors or not.

Nevertheless, we do recognize a dichotomy of this general nature, but would draw the line in a slightly different way. If we accept the ‘translatability’ criterion as defining the verbal/referential boundary (which some authors might not), then that yardstick is relatively atheoretical, and can be applied to a joke without assuming any particular analysis of the joke. What we offer as a counterpart is a condition which *does* rely on the proposed analysis of a joke. Firstly, we classify joke *classes*, then jokes. We define a *linguistic* class of joke as a class whose definition explicitly imposes conditions on the *linguistic form* of the joke, in phonetic, orthographic, or syntactic terms. For example, the various types of puns we shall review in [Chapter 9](#) are definable at least partly in terms of phonetic similarity. Hence the definitions we will propose ([Section 9.3.6](#) and [Appendix A.2.2](#)) define a linguistic class of joke, as it imposes conditions on the phonetic form of the jokes. In contrast, a *propositional* class of joke has a definition which explicitly imposes conditions on the propositional content (meaning)

of the joke, but not its linguistic form. The jokes discussed in [Chapter 8](#) fall under this heading.

If we accept these definitions, we can then label a joke as linguistic or propositional according to which of these classes it falls into. This classification is sufficiently different from the usual verbal/referential one that it would be misleading to adopt those existing labels. However, the linguistic/ propositional split does not replace or compete with the verbal/referential dichotomy: it is simply another way of classifying the data, a classification which may or may not be helpful in organizing discussion.

2.8 Studying jokes

The focus in this work is on *jokes*. It is not possible at this stage to give a strict formal definition of what we mean by a joke, as such definitions require a full theoretical framework (i.e. the eventual outcome of our programme of research). What we can give is an informal, pre-theoretic outline of the type of textual item (written or spoken – see [Section 3.2](#) for discussion) we have in mind.

By joke, we mean a relatively short text which, for a given cultural group, is recognizable as having, as its primary purpose, the production of an amused reaction in its reader/hearer, and which is typically repeatable in wide range of contexts. We take it as evidence that a text is a joke if it appears in a published form explicitly labelled as being a joke (e.g. a joke book, a web-site of jokes, examples in an academic paper on jokes), or if we have experienced it being delivered in circumstances which imply that others regard it as a joke.

Jokes are methodologically convenient for the following reasons:

- There is a ready source of attested examples, since collections of jokes abound, and there seems to be relatively little dispute about the status of such items as jokes (see [Section 2.9](#)).
- Jokes are relatively self-contained and are typically re-used in a wide range of settings. This is not to deny that there are certain requirements before a joke is appropriate in a particular context (see [Section 3.6](#)). However, we believe that these constraints between joke and setting are probably less, and more easily described, than the rich connections involved in general amusing incidents that happen in everyday life.
- Jokes are small, which renders them more manageable for the analyst. Also, if we should succeed in developing a range of formal descriptions of types of jokes, such data will be a potentially useful resource for psychological experiments on human reaction to humour (see [Section 2.11](#) below), where presenting a subject with a lengthy text might be inconvenient in practical terms.

Attardo and Chabanne (1992) provide a more general discussion of the concept of a 'joke'.

2.9 Jokehood and funniness

Crucial to this investigation is a distinction – not always made in humour research – between *being a joke* and *being funny*. One issue which humour research (particularly the narrower, more detailed studies which this volume exemplifies) must address is the fact that that different people tend to find different stimuli funny, and that whether an audience finds a joke funny will depend on all manner of factors: audience mood, timing of the joke-telling, cultural biases, etc. These sceptical observations are certainly true about the question of 'being funny', but they do not necessarily apply to the question of 'being a joke'. A central assumption of the work here is that, within a given culture or society, there will be a general recognition of what constitutes a joke, and that even when a person finds a joke not funny, that does not mean that they will contest that it is, in some intuitive sense, a joke.

Working assumption: within a particular cultural community, there will be broad agreement about which texts are and are not jokes.

It would be interesting to set up an experiment in which subjects (all from the same cultural background, broadly speaking) are asked to classify a set of texts into jokes and non-jokes. A statistical measure of their agreement could then be computed. Our working assumption is equivalent to a hypothesis that the agreement would be high.

What qualifies as a joke can be expected to vary between cultures or societies, and we are not at present seeking a cross-cultural or universal characterization of jokehood. As our philosophy is to start with small-scale detailed studies and work from there, confining attention to a narrow cultural scene (roughly, twentieth century Anglo-American) is not a severe limitation.

The programme of work advocated and initiated in these chapters is concerned *not* with funniness directly, but with jokehood. We believe joke-hood to be a more basic, and more manageable, concept, which will – if elucidated thoroughly – help in the subsequent study of funniness.

A distinction is sometimes made between *joke comprehension* and *joke appreciation*. The former is the process of recognizing jokehood, the latter involves funniness.

2.10 Declarative and procedural accounts

When analysing a particular joke, a linguist may well have to make use of intuitions about the text's salient features. This introspection is sometimes presented informally in terms of the steps that go on when the joke is interpreted: a *procedural* account. For example:

This puts the hearer in the situation of having to decide ...

(Attardo 1994:129)

In the first stage, the perceiver finds his expectations about the text disconfirmed ...

(Suls 1972:82)

See also Attardo *et al.* (1994:28–30).

This raises an important point about the linguistic description of humorous effects in jokes. It is intuitively clear that many jokes rely on information being revealed in a certain order (not prematurely), and authors have remarked on the use of surprise or expectation in jokes. Hence, any formal account of the linguistic mechanisms of jokes must, if it is to allow an insightful account of such humorous effects, make reference to the *sequence of presentation* of semantic content. This means that, even if we deliberately specify as little as possible about our assumed linguistic model (see [Section 3.1](#)), what we specify will have to take into account left-to-right processing of textual content. In this respect it will be different from most linguistic theories, since the normal approach in linguistics is to define abstract structures (syntax trees, logical forms, etc.), and leave left-to-right processing as a separate issue, to be left to computational linguists or others (*dynamic syntax* (Kempson *et al.* 2001) is an exception to this convention).

There are also some subtler points to note. If an account of the workings of a type of joke is stated solely in terms of an algorithm for scanning the joke, with no indication of what structures or conditions allow this algorithm to operate, then we have no characterization of what items qualify as jokes. If these more 'declarative' details are not stated, then the only definition we have of the joke type (or more generally of a humorous stimulus) is that it is any item which behaves in a particular manner under the given procedure. There is no accompanying characterization of what sorts of artefacts give rise to this processing. This is not disastrous, but it is an incomplete picture. Also, (quasi-)algorithms, such as those of Attardo and of Suls, purport to trace the hearer's mental steps *on first hearing the joke*. As Suls (1972:82) comments, the joke 'is found to be funny as a result of' his proposed process. (See [Section 5.3](#) for more discussion of Suls' model, which epitomizes the use of processing as the means of stating

a theory of jokes.) This does not tell us whether something is a joke in general: what it tells us is how the text could be perceived as funny, which is slightly different. For example, it could be that the joke-audience does not find the joke particularly funny, or even sees the punchline coming, but that does not mean that the audience does not recognize the text as being a joke. Hence, from our point of view (see [Section 2.9](#)), purely procedural accounts (such as that of Suls) appear to be addressing *funniness* (or joke appreciation) rather than *jokehood* (or joke comprehension).

In [Section 3.9](#) we will sketch our assumed model of text processing, and in [Chapter 8](#) we will discuss how various types of joke could be described in those terms. However, we will also show how these processing commentaries are related to the structure of the text.

2.11

Linguistic analysis and the theory of humour

Latta (1999) attacks the notion of a stimulus-side theory of humour; that is, a theory which states that the essential factors in humour are intrinsic to the stimulus (as opposed to being in the audience's reaction to the stimulus). He has in mind much of the work on incongruity theories of humour (see [Chapter 4](#)), which focusses on the question of which stimuli qualify as amusing. It might seem that the approach we will be developing is intended to produce a 'stimulus-side' theory, since we are concentrating on the detailed analysis of humorous stimuli. This is not the case. It is important to distinguish two notions: on the one hand, the overall content and structure of a complete theory of humour, and on the other hand, the various routes one might follow in building up such a theory. Our emphasis on humorous artefacts is part of the latter, a strategy for moving towards a theory: we are arguing that one feasible starting point is the examination of humorous stimuli. This does not mean that we see these analyses, on their own, as forming a theory of humour. Instead, these analyses should provide richer empirical data from which theories can be derived. Even a humour theory which escapes Latta's accusation of being 'stimulus-side' will, if it is comprehensive, have to say something about why certain stimuli are deemed to be humorous, even if it does this by showing how these stimuli initiate certain cognitive processes (Latta's favoured form of explanation). If the theory is to be empirical, tests will have to be carried out to show the relationship between stimuli and responses. In such a broader scheme, knowing the nature of relevant stimuli will be helpful. If a theory defines humour solely in terms of some mental processes, without indicating exactly which stimuli would cause such mental processes, then there is not much that one can do to corroborate or falsify such a theory.

It is perhaps instructive to sketch how a wide-ranging and detailed account of the structure of jokes might fit into a broader and deeper

theory. That is, let us imagine that the programme of work being proposed here has been completed: where would that then contribute to humour theory in the large? The first point to make is that even a successfully created set of joke-class definitions might well not appear in a humour theory: they are intended to be a contribution to the *process* of arriving at a theory, not part of the theory itself. Suppose that we have formulated a set of precise joke-class definitions that covers virtually all of some large and coherent set of jokes (e.g. jokes commonly available in British culture in the late twentieth century), and that these definitions are as integrated and systematic as possible, in the sense that a consistent set of abstract concepts is used throughout, with as much generalization as the data will support. Also, these analyses do not exist in isolation: we have to assume that we also have a comprehensive theory of linguistic structure and language processing, and that the joke-class definitions are stated, where appropriate, in the terminology of that (psycho)linguistic theory. Even that would be just a preliminary step, allowing research to move on to a similarly thorough analysis of the *non-linguistic* aspects of jokes, determining the patterns of content which seem to occur in jokes, and developing a fuller notion of joke class. In such a (futuristic) situation, a joke-class definition will have definite implications for how a joke of that class will be cognitively processed. That would mean that we had some stimuli (jokes) which were attested to be humorous paired with some patterns of cognitive processing which our psycholinguistic theory told us each stimulus should provoke. From there, we could see what regularities (if any) there are in the mental processes, and – we hope – move towards a cognitive theory of the humour phenomenon. Notice that the joke-class definitions have merely fulfilled a mediating role, allowing us to connect raw data (joke texts) to a relevant level of response (cognitive processing).

Why, it might be asked, do we need joke-class definitions in order to follow this theory-building path? Could we not just apply our (presumed) theory of language processing to jokes as texts, and consider the implications of that? Theoretically, that would be a valid route. It would involve examining for regularities and generalizations the cognitive processing predictions for a wide range of texts. In practice, this might have limitations: if the joke texts were not categorized in any way, but just considered as an undifferentiated set, the regularities in their associated cognitive effects might be quite hard to find. This raises an important if rather banal point:

Working assumption: different types of joke may create their effects in different ways.

It would then follow that sorting out the data in an orderly way, before we attempt to correlate mental effects with textual stimuli, should make

the work more manageable, and we are less likely to be swamped by a varied mix of unrelated data. Also, if we were to subject the predicted processing effects to psycholinguistic experiment, a mix of joke-types might lead to a mix of effects which appeared random or arbitrary. Indeed, the very design of the experiments would be difficult in the absence of a systematic description of the salient features of each joke: what would it be relevant to measure?

To sum up: although we are proposing our programme of research as worthy of the investment of effort, we are not claiming that it constitutes the construction of a theory of humour (although it might be regarded as a *descriptive theory of jokes*). Instead, it is intended as a sifting of the data, to facilitate subsequent theory-building. Moreover, the concentration on linguistic aspects (and related presentational factors) is itself just a ground-clearing tactic to allow other components of jokes to be seen more clearly.

2.12 Testing and computers

Computational modelling (as in artificial intelligence and cognitive science) is a major inspiration and guiding framework for the arguments, discussions and proposals here, even though actual computational models play only a background role. As discussed in [Section 2.5](#), it would be desirable if a model of humorous texts could, at least in principle, be tested by constructing a computer implementation of some kind. This leads to the need for a systematic methodology for evaluating such implementations.

Once again, some inspiration comes from the generative approach to linguistic theory. In that paradigm, the researcher posits symbolic rules which would give rise to well-formed sentences. Then a check is made to see whether the strings or structures characterized by these rules are indeed properly constructed sentences, as judged by the intuition of native speakers of the language. Computationally, this can be implemented by a grammar-tester program, which allows the user to apply complex sets of rules in order to see their effect in terms of sentences (e.g. Friedman (1971, 1972)).

Similarly, we could posit abstract rules for jokes, and consider whether the predicted outputs from these rules are indeed jokes. Again a computer program could be used to assist in this process, and to ensure that the rules were followed precisely. The next issue is the evaluation of the output of the program. Whereas linguists are able to be relatively informal about assessing the grammaticality of their outputs, jokehood is perhaps a more slippery concept, meriting more careful assessment, for example by arranging some controlled appraisal of the jokes built by the computer (i.e. by the rules, which embody the theory). We will return to this issue in [Section 10.9](#).

It is important to distinguish between building a joke-generating program simply to test out a formalized theory, as suggested above, and attempting to create an ‘intelligent agent’ which ‘has a sense of humour’. The former need only exercise a set of rules to see what jokes are generated, as in a grammar tester. The latter (depending on how the task is defined) might be assessed according to similar criteria to those that are applied to a human joke-teller: creating novel jokes, being particularly funny, selecting jokes which are appropriate to a context of use, tailoring jokes to a particular audience. This is altogether a harder task.

An alternative method of testing theoretical proposals, instead of building a joke-generator, would be to create a joke-understander, a program which takes as input a possible joke and gives as output a rating of whether it is a joke, or of its funniness. For some theories of humour, this might be much more appropriate. For example, it is quite common, as noted in [Section 2.10](#), for informal analyses of joke mechanisms to be phrased in dynamic terms, setting out what happens as the various parts of the text are conveyed (see, for example, the model discussed in [Section 5.3](#)). If such an account were to be formalized in terms of processing steps, a joke-understander would be the natural implementation to test the model.

Where the theory is stated more statically or declaratively, and simply specifies particular relations which must hold between symbolic entities, then a generator and an understander might be equally natural test computations. However, the understander may in practice be more awkward. Leaving aside the need for various supporting facilities (such as a natural language parser), there is the issue of creating suitable test-data. The ultimate goal of a theory of humour may be to isolate factors which are both necessary and sufficient for something to be humorous. Nevertheless, until the ideal theory is constructed in complete detail, we will always be working with partial or imperfect theories. We shall be positing rules which cover some (often small) subset of humorous phenomena. In such a situation, we will be claiming that our rules are *sufficient* for humour, but not *necessary* (see also [Section 2.14](#)). To test such a hypothesis, a joke-generator is convenient, since all its output conforms to the rules, and we can use humans as joke-evaluators, a task they are well suited to. If every output item is a joke, our claim of sufficiency is met, and every non-joke output counts as evidence against the sufficiency claim. On the other hand, how would we assess the verdicts of a joke-understander? If it classed an input joke as a non-joke on the grounds that the item did not conform to its rules, all we could conclude is that our rules did not cover all possible jokes, but that would not be a surprise. If it classed a non-joke input as a joke, this would refute the sufficiency of the rules, but this would require more intricate testing in which we would have to find potential inputs which are not jokes (according to humans) but which might be characterized by the system’s rules as being funny.

Until we require a joke-understander in its own right (see [Section 2.13](#)), a joke-generator is likely to be a more directly manageable test mechanism for declaratively stated symbolic rule systems.

2.13 Practical applications

As mentioned earlier, the motivation for this book is science rather than engineering. Nevertheless, given the computational influence already acknowledged, it is worth digressing briefly to consider some of the issues involved in technological uses of humour.

What might computational models of humour be used for? If we had a workable model of humour, there are two obvious ways of implementing it: as a humour generator, or as a humour understander.

A computer generator of jokes could be of use to those who need a steady supply of jokes, particularly if quality was less important than quantity (e.g. those firms who, in Britain at least, insert riddles or other simple jokes into Christmas crackers). At present, computer systems are not in a position to outperform human joke-creators, but for very simple jokes aimed at children, it is conceivable that such a system could be workable in a few years (see the systems reviewed in [Chapter 10](#)). Writing more complex humour, such as comedy sketches, is much further away.

A joke-appreciating program is of less obvious use. It would have to be remarkably good before we would consider using it as a form of quality control for any human joke writing.

Let us consider both joke-generation and joke-appreciation from the standpoint of human-computer interaction. Could computationally tractable theories of humour help in the building of better user-interfaces or more usable systems such as robots? If we are to cooperate with robots at work, or have intelligent agents as our constant advisors, perhaps some humour would make interactions more pleasant. Binsted (1995) argues that a computer system could be made more congenial by judicious use of humour generation in a user-interface. She suggests various situations in which a humorous remark from the system could ease the interactions – errors, poor system performance, offers of hints, requests for clarification – and speculates that certain styles of humour would be appropriate – self-deprecatory remarks, observations about the situation, use of amusing phrases. However, she acknowledges that there are difficulties in such attempts, since humour – particularly imperfectly created examples – can easily irritate rather than relax the user. Stock (1996) also argues for the desirability of injecting humour into various sorts of computer applications, including education and entertainment, although he does not offer detailed proposals for how this would work. Takizawa *et al.* (1996) comment that ‘Kitagaki developed a pun generator to be applied to a human-friendly

Japanese word processor', citing Kitagaki (1990, 1993). It has also been suggested that machine-assisted joke generation could help children or second-language learners explore language (McKay 2002, O'Mara *et al.* 2002).

Morkes *et al.* (1999) describe some experiments which demonstrate very convincingly that users find it more pleasant to make use of a computer system which employs appropriate humour. Loehr (1996) reports a preliminary experiment in putting a joke-generating module (the JAPE program (Binsted 1996)) within a user interface, concluding that it is difficult to arrange automatic generation of a joke which is *relevant* to what the user is trying to do. (In the Morkes *et al.* experiment, the test data consisted of carefully pre-constructed and directly relevant remarks, not automatically generated jokes.) In considering a system like Loehr's, it is not clear that the behaviour crucially relies on jokes being computer-generated, as opposed to merely accessed online. Unless the process of generation is subtly influenced by the context, in particular the user's input, then the interface might as well produce jokes from a stored list, either computer-generated or taken from some other source. That is, if the relevance of a computer-generated joke is measured in some crude way such as keyword matching (as in Loehr's implementation), there are no benefits, and possibly some disadvantages, in having the joke generated on the spot rather than retrieved from a suitably cross-indexed database of jokes. A perfectly intelligent joke-generator would change this argument, as it would be able to create new jokes which were directly pertinent to the user's situation, in a way that could not be simulated by a pre-computed list. We are a long way from knowing how to design such a generator.

There are not many proposals to have a user-interface *understand* jokes, but Nijholt (2002) argues that this would be a useful ability for a lifelike software 'agent' to have, in order to facilitate natural interactions with the human user. As with many advanced facilities, an imperfect version of such a facility (which is all that will be achievable in the near future) is likely to be worse than none. The goal of creating a robot that is sufficiently 'human' to use humour in a way that makes sense or appears amusing (other than inadvertently) is a long-term one.

For practical attempts at user-interfacing, the issue of evaluation should at least be relatively straightforward. Any interface which had supposedly been enhanced with humour could be compared in some controlled way with a non-humorous interface, and rated for ease of use, etc. In an engineering context, if the enhancements achieved their desired aim, they would be a success.

With these remarks, we will leave practical computational applications of humour theory aside.

2.14

A descriptive framework

Building on the discussions so far, we can outline a way of working towards a comprehensive description of verbally expressed humour. The approach starts by describing particular classes of jokes, at a level of detail that should reveal pertinent factors within the linguistic structure of the jokes.

The aim is that the humour analyst should set out the crucial aspects of the joke class, using whatever linguistic concepts (phonetic, syntactic, semantic, etc.) are needed (see [Section 3.9](#) and [Appendix A.1](#)). Such an analysis does not in itself *explain* the ‘funniness’ of the jokes; that is the task of an eventual full theory of humour, which would state what it is about the joke-class definitions that renders them humorous. That is, the statement of the essential attributes of the joke class will display certain relationships between elements of the jokes (e.g. the meanings and sounds of its parts). A (future) theory of VEH would then indicate, at a suitable level of abstraction, whether (or why) these relationships amount to humour or not. In this way, a theory of VEH does not itself have to itemize low-level details of how jokes are constructed. The important characteristics of the humorous texts are summarized by the joke-class definitions, and the humour theory then adjudicates upon these more abstract definitions. This is comparable to the distinction, in linguistics, between a grammatical description of a language (which says which structures are well formed) and a universal theory of grammar (which explains why grammars have to take certain forms).

In an earlier version (Ritchie 2000) we proposed a notation for writing succinct structural descriptions of jokes, but that attempt at formalization was premature. We do not yet know exactly what the components of a joke-description should be (see [Chapter 12](#)).

The central idea in providing a formal definition of a class of jokes is that the analyst must specify what *abstract objects* are involved, what *types* of objects these are (e.g. meanings, word-strings), what *properties* these objects have and what *relations* there are between them. That is, the aim is simply to say what is present, at some suitable level of detail. To do this, we shall have to posit certain abstract concepts as primitive elements in our description. These may be relatively simple (e.g. phonetic segments roughly corresponding to phonemes) or extremely complex (e.g. contexts in which a joke may be uttered). (See [Section 3.9](#) and [Appendix A.1](#) for some proposed basic concepts.) If the joke-class definitions are to be connected to the data, or are to support theorizing in the manner intended, the primitives must take on more meaning at some stage. There are broadly two ways in which they could do this, relative to a theory of humour: *theory-external* and *theory-internal*. A primitive is given a theory-external definition when some other theoretical system provides the details. For example, for some basic concepts (e.g. phonetic segments, syntactic structures)

linguistic theory could be expected to supply definitions. A theory-internal definition is needed where the humorous effect of a text will depend on the exact definition of the primitive in a way which could not be predicted by any theory from another discipline (e.g. linguistics, psychology). For example, the SSTH ([Chapter 6](#)) proposes that a form of conflict between meanings leads to humour, but that only certain types of conflict are funny; this is an example of a theory-internal definition.

This means that a joke-class definition shows explicitly only some of the factors that make texts humorous. The remaining factors (which may be highly significant) are embodied in the theory-internal definitions of some of the concepts employed. In our analyses, we shall not attempt to give precise, detailed or formal definitions of all our primitives. For those where a theory-external definition is appropriate, we shall assume this can be given; for the more substantive primitives (from a humour viewpoint), this task is left for further research. In this way, establishing a set of primitives suitable for describing jokes also decomposes the task (of explaining humour) into subsidiary problems, namely, providing the theory-internal definitions. In this sense, those primitives are only provisionally ‘primitive’; in the longer term, humour theory should render them non-primitive by offering full and precise definitions. Even where theory-internal definitions are needed, the hope would be that at some stage the humorous concepts would be related to notions outside humour theory, maybe after several levels of intermediate definitions. If this were not the case, if our central concepts were all completely defined within the theory of humour, we would have given a more detailed *description* of the humour mechanisms, but we would not have *explained* the humorous effect. Explanation is achieved when a phenomenon is systematically described in terms of other *independently given* concepts.

The framework advocated here is *not* a theory of verbally expressed humour. It does not lay down principles for what is or is not humorous. Instead, it is a methodological approach which is intended to push forward the search for a theory of VEH, by describing the data more precisely and at an interesting level of abstraction. However, it is virtually impossible for a statement or proposal to be completely free of theoretical content, whether deliberate or unintended. Even a very general suggestion for how to tackle the problem, as here, makes certain implicit assumptions which constitute at least minimal empirical claims.

Many of the working assumptions for the descriptive methodology advocated here will be summarized in [Section 3.9](#). They involve, for example, the existence of texts that can be segmented into subparts, and of meanings which are associated with texts. To that extent, the approach makes some very sketchy claims about humorous texts, but many of these assumptions would be shared by much research into VEH; hence, they are not highly controversial. (Of course, this is different from being *correct*.)

It is also important to acknowledge that the joke-class definitions we will offer in later chapters are still very vague. They are no more than first attempts to sketch the relevant factors in the various jokes, in a way which is slightly more precise than past practice, and which could be developed further in detail and precision.

What we will have, if this approach is followed, is a provisional or preliminary descriptive theory of jokes, built upwards from the data, but not yet claiming to cover all classes of VEH (let alone all humour). If it does not cover all humour, then it cannot claim to present necessary and sufficient ingredients for humour. However, within the proto-theory there would be descriptions of (sub) classes of joke, and each such description might well (ideally, should) define necessary *and* sufficient conditions for membership of that subclass. The overall proto-theory would then state a *sufficient* but not *necessary* condition for being a joke, namely that a text fall into one of these defined classes. That is, the provisional theory would be of the logical form:

T is a joke if $P_1(T)$ or $P_2(T)$... or $P_n(T)$

where each P_i is the definition of one class of joke.

Within each subclass, the restricted local theory would be relatively detailed, and so would allow falsifiable predictions (within that limited class of jokes). Hence by surrendering, at least for the moment, the claim of universality, we can gain some of the benefits of a detailed and relatively formal theory, albeit on a small range of data. In principle, if we were sure that the set $\{P_1, P_2, \dots, P_n\}$ was complete, then the 'if' in the above statement could be strengthened to 'if and only if', but that stage is a long way off.

This could be a manageable way to tackle the enormous problem of creating a theory of the whole of VEH, and it is the framework which we shall be following here.

It might appear that what is being proposed here is yet another taxonomy, redundantly piled upon the many classifications offered in the literature (Attardo (1994) reviews several). Our approach should differ from most existing taxonomies in the following ways:

- rather than just grouping jokes into labelled classes, we shall provide definitions of the joke classes in terms of what we propose as their internal mechanisms;
- these definitions will, as far as possible, be given at a level of formality which most earlier analyses have lacked;
- a clear interface with conventional theoretical notions from linguistics and logic will be specified, so that the relation to mainstream language theories is clearer.

Our aims differ from some researchers who are adamantly not interested in theory:

the following chapters ... do not presume to present a 'theory' of English verbal humour. It is the author's deeply felt conviction that in no possible sense whatsoever of the term is such a 'theory' either desirable or attainable.

(Alexander 1997:18)

The anti-theory approach is also exemplified by the completely analysis-free presentation in Redfern (1984), as Raskin (1987) makes very clear. We are definitely not of that persuasion. Detailed description and analysis are not ends in themselves, but offer a route to a more adequate theory of jokes, and thence of humour.

3

Linguistic preliminaries

We set out what assumptions we will make about linguistic concepts, and how analysis of jokes is related to traditional constructs in linguistics and in language-processing.

3.1

The need for linguistic concepts

One methodological issue in studying verbally expressed humour is the need for linguistic notions such as ‘word’, ‘sentence’, ‘type of speech act’, etc. (Here ‘linguistic’ will be taken in a fairly broad sense, to include pragmatics as well as morphology, syntax, semantics, etc.). It is important that analyses of humour make sound assumptions about linguistic structure, but humour analysts should not have to develop a complete theory of language before proceeding to their central interest. This would be less of a problem if there were a single, accepted, fully detailed theory of language, as we could then adopt this as the foundation or conceptual framework for stating our analyses of humour. Unfortunately, there is no consensus even about what a theory of language should look like, and there is certainly no received theory that we could rely upon. In rare cases, a humour researcher may opt for a particular linguistic framework within which to express his humour theories, although the clearest example of this – Raskin (1985) – involves a linguistic theory devised by the same researcher.

Computer programs which manipulate information to humorous effect (see [Chapter 10](#)) typically include *ad hoc* language processing mechanisms in order to achieve their results. This is a practical necessity. If we wish to implement and test software now, we cannot wait for some linguistic theory to become universally accepted as the truth. We have to write program modules which will ensure that the relevant pieces of text are put together in the desired manner. The disadvantage of this approach is that it may be hard to disentangle the humour-specific concepts from the hand-crafted language processing mechanisms, particularly as the latter may not look like standard linguistic constructs.

The methodological approach that we have adopted is something of a compromise between these two positions. We have analysed various joke classes, and have tried to set down the regularities that we found. In doing this, we have assumed whatever linguistic notions seemed to be needed in order to state our generalizations precisely. These proposed linguistic notions were not invented whimsically or arbitrarily: they were based closely on ideas which are relatively commonplace with linguistics and logic. What we did not attempt to do is compile a complete linguistic theory. Our aim was simply to specify what facilities a linguistic theory, or more specifically a model of language processing (see [Section 2.10](#)), should provide to humour theory.

Our assumed set of linguistic concepts is summarized briefly in [Section 3.9](#) and defined more formally in [Appendix A.1](#). This set of definitions is very far from being a comprehensive model of language, and it is not intended as such. It can be thought of more as a *specification of an interface* between humour theory and language processing theory, in the spirit of what is known, within computer science, as *data abstraction* (Abelson and Sussman 1985: Ch. 2). It lists those technical terms for which a linguistic theory would have to supply definitions, if our humour analyses were to make sense. (This is slightly different from Attardo (2001: Ch. 3), who seems to be proposing a substantive model of text processing.)

The list of required linguistic definitions ([Section 3.9](#) and [Appendix A.1](#)) are assumed by our humour analyses (i.e. the definitions are logically prior to the analyses), but this does not reflect the chronology of the work. The definitions were to a large extent an *outcome* of carrying out the humour analyses, rather than a starting point, since we refined our ideas about what constructs were necessary by analysing the data (jokes). [Appendix A.1](#) may be consulted as the various analyses are presented in later chapters.

3.2 Speech and writing

One basic issue to be clarified is whether we are dealing with spoken or written language, or both. It is obvious that each of these modes are commonly used to convey jokes, but the consequences for any attempt at formalization must be made explicit. In considering the verbal mode of a joke, it is important to distinguish these separate issues (as Dienhart (1999: Section 6) does not):

- (a) Which mode(s) of linguistic communication (spoken or written) allow the humorous effect to be achieved?
- (b) Which mode(s) of linguistic communication (spoken or written) allow the humorous effect to be achieved most effectively?

- (c) In stating abstractly the humour-relevant relationships within a joke, which levels (spoken or written) are essential?
- (d) Which linguistic knowledge does the joke rely upon?

When a pun is included in a joke book, it is conveyed in writing, but the humour-relevant relationships are defined on the phonetic form of the words (see [Chapter 9](#) for fuller discussion), and hence the written version is dependent on the reader knowing how to pronounce the written forms.

- (6) Do you believe in clubs for young people? Only when kindness fails.

(Attardo 1994:97)

On the other hand, the joke (6) – which appears in some form or another in various articles, including Pepicello and Weisberg (1983) and Shultz (1976) – could be conveyed equally well phonetically or in writing, as its central ambiguity is present in both spoken and written forms. No pronunciation knowledge is required, and the humour-relevant relations could be defined equally well on the phonetic or orthographic representations. The same will be true for the (large) class of jokes where there is no language device involved (*referential* jokes – see [Section 2.7](#)). This also means that we should, to be rigorous, distinguish between a particular telling (instance) of a joke, and the joke itself. The same joke could be told in written form on one occasion, and in spoken form on another. In the interests of simplicity, we have not stated our analyses in terms of a joke-telling or joke-instance. Instead, we have analysed jokes in the abstract, assuming a particular medium. This can be viewed as abstracting across all tellings which use that medium (i.e. making a *type/token* distinction (Cherry 1970)).

Consider a joke which hinges on an ambiguity of the spoken form which is not present in the corresponding written form, such as some of the examples in Dienhart (1999: [Section 5.3](#)), or (7). (As noted earlier, text between forward slashes, such as /r'ʌ/, is a written approximation to a phonetic form, using, as far as possible, the International Phonetic Alphabet.)

- (7) What is black and white and /r'ʌ/ all over?

A newspaper.

In (7), either orthographic form (*read* or *red*) destroys the essential ambiguity, and so it is arguable that this joke cannot be told in orthographic form; it certainly cannot be used *effectively* in that medium. Also, the humour-relevant relationships are statable perspicuously in terms of the phonetic form, but much less neatly in terms of a written version.

As far as linguistic knowledge is involved, the joke relies on knowledge of pronunciation. Catherine Watson (personal communication) has pointed out that there is a similar issue involving intonation. In certain ambiguity-based jokes, the two readings correspond to different intonations. Hence, if the joke-teller chooses a particular prosodic contour, this rules out one of the possible interpretations of the text. Example (8) may be such a joke.

(8) Why do birds fly south in winter?

It's too far to walk.

(9) Jean Harlow kept calling Margot Asquith by her first name, or kept trying to: she pronounced it *Margot*. Finally Margot set her right. 'No, no, Jean. The *t* is silent as in Harlow.'

(Matthews 1974)

It is clear that the original form of (9) was spoken, but it is unclear what written form is most suitable for reporting this joke, even though the cited version settles for one particular way of conveying the final word. This joke relies on the hearer's *knowledge* of spelling (as do some of the riddles in Dienhart (1999: Section 6)), but the phonetic form, where there is no spelling, is most appropriate to convey this wording. Either phonetic or orthographic form would suffice for a variant in which the punchline ended in ... *as in your surname* instead of ... *as in Harlow*. That variant both avoids the issue of representation and demonstrates how the joke differs from jokes like (1) and (2): although it relies on linguistic knowledge, it is not crucially dependent on the exact linguistic form of the delivery (it is propositional rather than linguistic, in the terminology of [Section 2.7](#)). On the other hand, Cerf (1946) quotes a poem which mocks the multiple pronunciations of *-ough* in English, by rhyming words such as *bough* and *cough*, or *trough* and the invented *scough*. Crystal (1995) quotes a poem about a pterodactyl in which several words customarily starting with *t* are spelt instead with *pt*. For both these examples, the appropriate medium of communication is the written form, but they rely upon knowledge of pronunciation for their effect. Such items are unusual, and are of a slightly different genre from the types of verbally expressed humour that we will be focussing on.

(10) There are 10 kinds of people in the world – those who understand binary numbers and those who don't.

Although example (10) is definitely a verbally expressed joke, it is hard to deliver purely phonetically, since either choice of pronunciation for the symbol *10* (/tʰɒ/ or /tʊ/) loses the essential information about the text. The relevant knowledge concerns binary numerals, where *10* encodes the

number two. As with spelling/pronunciation based jokes, the knowledge concerns the encoding method used in the text, but it could also be seen as extra-linguistic real-world knowledge (about binary arithmetic). Although knowledge of pronunciation and spelling is so closely involved with language that we might regard it as qualitatively different from world knowledge, it is not obvious that this is the case from the viewpoint of its role in joke comprehension.

Crystal also offers (11) as a joke which is hard to convey in spoken form:

(11) TOO MUCH SEX makes you shortsighted.

(Crystal 1995:406)

For this example, we need to go beyond having an orthographic level of representation, and to include some level where typeface (font size) information is available. There is a phonetic counterpart to (11), namely (12).

(12) Too much sex [*voice drops to a whisper*] makes you hard of hearing.

This, conversely, does not work well in written form, except insofar as it can simulate the phonetic delivery by relying on knowledge of stage-directions, as here.

Crystal gives (13) as a joke which is easier to deliver phonetically, although an attempt can be made to simulate the effect typographically:

(13) Airline passenger: Where does this door go to-o-o-o...

(Crystal 1995)

Surprisingly, Crystal also says that (14) works only in the written mode.

(14) What did one sheep say to the other?

I love ewe.

(Crystal 1995)

However, there are many puns like this which are passed around orally, relying on the hearer's ability to spot a potential homophone in the spoken text. Crystal may be confusing some of the questions listed above: what knowledge (of words, etc.) is needed to interpret the joke, in which mode can the joke be conveyed, and in which mode can the joke be conveyed *most effectively*. The spelling *ewe* may make the joke clearer to the recipient – consider the difficulty of comprehending a written version using the spelling *you* (similar comments apply to some of the examples in Dienhart (1999: Section 5.5)).

Notice the contrast between (14) and (7). Because the ambiguous word or phrase occurs in the punchline of (14), the ambiguity need not

be maintained, and indeed must be seen by the audience (hence the spelling may influence the *effectiveness* of the telling). On the other hand, in (7) it is essential that the ambiguity in the question misleads the audience; hence, the possible means of delivery are limited.

Providing we have a well-defined set of linguistic relationships (e.g. between written and spoken forms of words), there may be a choice regarding which linguistic level we take to be the canonical form of the joke text. The very fact that we are concentrating on verbally expressed humour is already a simplification, since there are jokes which could be conveyed either pictorially or linguistically, widening further the issue of which communicative mode is to be assumed in an analysis.

For these reasons, we will not stipulate whether verbally expressed humour is in the form of spoken or written texts. In an analysis of some particular joke, or class of jokes, it may be appropriate to specify that a particular mode will be assumed as the basic level. That is, if the natural generalizations about certain forms of puns are phonetic, then assuming a phonetic level as the representation of the text may be a more insightful descriptive approach, with the written form being used merely as a secondary representation which can be used to convey the more apt phonetic form. In other cases, it may be sufficient to adopt one of the modes while indicating that either level is equally effective.

Hence, when proposing a detailed formal analysis of some class of jokes, one essential step is to specify the set of atomic items which, when strung together into a linear sequence, will constitute the humorous item. In keeping with the discussion above, this first level of units may be phonetic or orthographic, or indeed any other suitable level that the analyst posits (e.g. Braille, Morse code). This symbol system out of which the linear form of humorous items (jokes) are to be composed could be called the *communicative alphabet*. Each joke will, in its most basic form, consist of a *sequence* of elements of this alphabet. If we wish to model the spoken transmission of jokes, for example, then this alphabet could be phonemes or similar phonetic/phonological segments.

The communicative alphabet is canonical only in the sense that it is the most *basic* level of representation – it does not purport to capture central properties in any definitive way. For example, if a joke is conveyed in writing or in speech, these two tellings of the joke would have to be analysed with different communicative alphabets, but this does not embody a claim that two different jokes are involved. For many jokes, it would be quite feasible (and even desirable) to develop analyses which show that the written and spoken form are in some sense ‘the same’ joke.

We shall adopt the term *text* to describe the sequence of symbols of the communicative alphabet which represent the joke at that level. This term is to be regarded as neutral, with no connotations of the joke being written rather than spoken.

As this book is in written form, all the illustrative examples will be conveyed in writing, even where this is an inappropriate medium, but we shall try to make it clear what form of the joke is intended.

3.3 Syllables, segments, and morphemes

It is reasonable to assume that we can, in framing our analyses of jokes, postulate some conventional linguistic analysis of the text, and then make statements involving concepts such as ‘words’, ‘meanings’, etc. However, one oddity of jokes is that they sometimes step beyond the conventional labellings needed for purely linguistic purposes, by clustering and segmenting the units with more freedom. For example, in (9), there is a comparison between the strings *-ot* and *-ow*. It is quite hard to say what sort of units these are, as they are defined in orthographic terms, having no syntactic or semantic content which might merit them being classed as morphemes. Although their *pronunciation* is a syllable, the whole point of the humour is that the orthographic form is different in each case, so it is hard to state the critical relationship in terms of syllables, assuming that syllables are defined phonologically. Similar remarks apply to the string *-ough* in the Cerf poem alluded to above.

A particular variant of this is the use of a substring of the text as if it were a morpheme, as in (15). (See Dienhart (1999: Section 5.5) for other examples.)

(15) He may not have been actually disgruntled, but he was certainly far from grunted.
(Wodehouse cited by Morreall (1983:70))

Examples like these make it doubtful that there could be a single well-defined level of grouping of phonetic units (e.g. phonemes) into slightly larger units which are relevant to humorous constructions. It appears that joke-making has complete freedom to group and segment without regard to any particular set of linguistic classifications. That is, the humorous use of language may require mappings between linguistic levels (e.g. the phonetic and the morphological) which are not part of any conventional linguistic analysis of the text. Hence, in order to provide an accurate and thorough analysis of a joke, it may not be sufficient to define some set of relationships between the various linguistic constituents which would already exist in a structural description (according to a normal linguistic theory) of the text. A conventional segmentation of example (15) into various levels of constituents would not show *gruntle* as a morpheme within *disgruntled* (and it could be argued that it is this distortion of linguistic normality which provides the humorous effect).

Working assumption: once an analysis of VEH has specified its communicative alphabet (phonetic or orthographic), any sequence of elements of that alphabet can be treated as a unit within analyses, unless otherwise stipulated.

3.4 Strings and structures

Much work on the linguistic aspects of humorous texts tends to describe the relevant factors in terms of words or sequences of words (strings), as opposed to more abstract linguistic structures. For example, Attardo *et al.* (1994) use as one of their central terms *connector*, which is defined as ‘any segment of text that can be given two distinct readings’; that is, the definition is stated in terms of textual strings.

Hockett (1977:259) proposes that there is a precise definition of the ‘punch’ of a joke (often known as the *punchline*). It is ‘the shortest terminal sequence, the replacement of which by suitably chosen words will transform the joke into a nonjoke’. This is pleasingly precise, and well suited to approaches which seek algorithmic definitions of grammatical forms (cf. the preoccupation in mid-twentieth century North American linguistics with ‘discovery methods’). Also, it is approved of by Attardo (2001:103) and (apparently) adopted by Attardo *et al.* (1994:31).

On closer examination, Hockett’s proposal turns out to have some counter-intuitive effects, as follows. There are many ways to turn a joke into a non-joke by small changes to wording. For example, in (6) changing *clubs* to *discipline* would remove the humour. However, Hockett defines the punchline as the *shortest terminal sequence* which is thus vulnerable. Hence it will in many jokes be the last word or phrase, since it is often possible, with ingenuity, to rewrite the end of the last sentence in a way which makes it unfunny. For example, in (6), replacing *fails* with *from parents and a rich home life is not available to develop their characters* yields a non-humorous text. What is unfortunate about this is that *fails* is now the ‘punch(line)’, not the sentence *Only when kindness fails*. This decomposition makes it awkward to define semantic or discourse relationships between the punchline and the preceding text, since the preceding text is now a sentence followed by a sentence lacking a predicate (verb phrase): *Do you ... kindness*. The situation is further complicated when we consider the textual variations, (16) and (17), which Attardo offers (to illustrate another point).

(16) Do you believe in clubs for young people?

Only when kindness fails, my friend.

(Attardo 1994:99)

(17) ‘Do you believe in clubs for young people?’ someone asked W. C. Fields. ‘Only when kindness fails,’ replied Fields.

(Attardo 1994:99)

Hockett's definition assigns punchlines of *fails, my friend* and *fails, replied Fields* to these jokes, which emphasizes the way in which relatively peripheral material may be included while semantically important content such as *Only when kindness* is excluded.

Consider (18), a much analysed example.

(18) 'Is the doctor at home?' the patient asked in his bronchial whisper. 'No,' the doctor's young and pretty wife whispered in reply. 'Come right in.'

(Raskin 1985)

Strict application of the Hockett formula defines the final word *in* as the punchline (it could be replaced with *back tomorrow*). Once again the Hockett punchline is a rather semantically empty fragment whose semantic relationship to the rest of the text is hard to state. Also, it could be argued that the overall interpretation depends on several parts of the text, including the phrase *young and pretty*, the use of *whispered*, as well as the final utterance *Come right in*. Leaving these intact while replacing *in* might spoil the joke, but it could also leave a slightly incoherent text. Hockett's definition makes no mention of creating a revised text which is not only non-humorous but coherent.

Another slightly awkward consequence of Hockett's definition is that a joke can have an empty sequence as the punchline. If it is possible to render a joke non-humorous by appending further text to the original, then strictly the shortest terminal sequence which is being replaced is the empty string. For example, adding to the final fragment of dialogue in (18) the text *and wait for him, but don't wake the baby* makes the story non-humorous, and even relatively coherent, as it accounts for the *whispered* verb. (Cf. the comment by Raskin (1985:105) that adding *He will be back soon* or *You can wait if you like* would allow the interpretation involving medical consultation.) It could be argued that the sexual innuendo is still present, largely due to the stereotyped use of the phrase *young and pretty*, but since (18) would function as a joke without that phrase, the demonstration could be based on that variant.

No alternative formal definition of 'punchline' is offered here. It is not at once obvious that a universal definition of punchline is something that is immediately helpful at the current stage of development of a theory of VEH. There may not even be a worthwhile definition which extends across all classes of joke. In the more descriptively-oriented approach advocated here, the segmentation of a joke into a punchline and any other subparts is the responsibility of the analyst who devises a structural account of that joke (or a class of jokes).

The above discussion of Hockett's proposal demonstrates that in defining the internal components and relations within a joke, excessive

concentration on surface text strings may lead to over-restrictive or misleading positions. In the chapters that follow we will not rely solely on the notion of textual strings, but will also consider syntactic and semantic constituents, in order that relevant semantic and discourse relationships can be defined between these items.

3.5 Grammatical well-formedness

In [Table 2.1](#), a theory of jokes was depicted as separating out jokes from other forms of grammatical (i.e. syntactically well-formed) texts. This embodies a premise of all the analyses that we shall present:

Working assumption: jokes are made up of grammatically well-formed sequences of words.

A counter-example to this would be a joke where the text *must* contain some ungrammaticality to achieve the humorous effect.

(19) Why did the antelope?
Nobody gnu.

(Crystal 1995, Dienhart 1999)

The wording used in (19) appears to have an ungrammatical question part, but in fact it is *grammatical*, strictly speaking, since it appears naturally in a discourse such as *Why did the gazelle die? Why did the antelope?* The sentence is just semantically or pragmatically odd as an isolated sentence. (Those over-acquainted with a certain comic song or free software may not at first realize that *gnu* has a conventional pronunciation homophonous with *knew*, and another similar to *noo*.)

This joke could have been written as in (20), which has no linguistic oddity in the question part. However, the answer is still ungrammatical, as written.

(20) Why did the antelope? Nobody gnu.

(constructed item)

On the other hand, (19) sounds completely grammatical if spoken, although the joke depends on subparts of the text being given alternative forms in a way that does not correspond to a complete syntactic analysis of the text. Crystal's claim that this example works only when written can be taken as an observation about the ease of spotting alternative readings. This is comparable to (14) earlier, where the joke can be conveyed phonetically, but one particular spelling of the punchline helps the audience to see the joke.

Our assumption that jokes are grammatically well formed does not rule out the possibility of jokes in dialect or in textual representations of nonstandard speech. These modes are themselves systematic, and a grammar (descriptive, not prescriptive) could exist for such forms of expression. If we can make the assumption of grammaticality, then the theory of humour need not contain information about grammaticality. That is, the generation of humorous text is seen, abstractly, as characterizing a *subset* of the grammatical texts which are deemed humorous. It is not concerned with characterizing from scratch the set of possible sentences of the language in question. The issues of *generative capacity/power* (Kimball 1973: Ch. 4) which sometimes concern linguists are therefore irrelevant.

The question of how much linguistic detail is needed in a humour analysis is an open and empirical question. Until we have developed some suitably detailed analyses of humorous texts, it is hard to predict what linguistic properties of texts will and will not be relevant (see [Section 3.8](#)).

3.6

Background information

It is uncontroversial to observe that understanding a particular joke may require not just knowledge of the language used to convey the joke, but also other types of knowledge (factual, cultural, social, etc.). For example, we have already observed that (10) relies on a grasp of basic binary numerals. In this respect, a joke is not different from other forms of text: understanding a text depends on familiarity with the relevant vocabulary and grammatical forms, and complete apprehension of the meaning may be dependent upon specialized knowledge. There is no evidence to suggest that jokes have a special requirement beyond that of non-humorous texts, in this respect. What causes a text to be a joke will usually depend on the meaning(s) of that text, and it is the comprehension of these meanings that in turn require specific linguistic or factual knowledge (if these two types of knowledge can be distinguished). This applies to both linguistic and propositional jokes ([Section 2.7](#)).

In non-humorous language understanding, an audience may understand a text to a varying degree, depending upon the available background knowledge. Similarly, the appreciation of a joke may depend to the extent of relevant knowledge.

(21) What goes ‘Pieces of nine! Pieces of nine!’?

A parrotty error.

Example (21) can be seen as a joke providing that the audience has the following background knowledge: *Pieces of eight!* is a stereotyped utterance by a parrot (based on a bird in R. L. Stevenson’s *Treasure Island*); there is a standard term *parity error* (in computing). However, if the audience also

knows that a parity error involves an erroneous numerical quantity, the joke may become funnier, and the knowledge that parity errors often involve numerical discrepancies of just one may further enhance the joke. (Consider a variant in which the phrase *Protty pelly* was used instead of *Pieces of nine!*; this should still qualify as a joke, albeit a less funny one for some people.)

As we are focussing on jokehood rather than funniness, we would define (21) as being a joke relative to a state of knowledge which contained the minimum to support the basic semantic relations which make it a joke, leaving aside the enhancement that further information might provide. Philosophers and linguists have drawn attention to a particular kind of required background information, namely *presupposition* (Levinson 1983, van der Sandt 1988, Raskin 1985:69–70, Attardo 2001:50–3). Roughly speaking, a sentence *S* presupposes some proposition *P* if *P* would have to be true for *S* to make sense. A classic example (Russell 1905) is the sentence *The king of France is bald* which is meaningful if there is indeed a king of France (in whatever context the sentence is uttered) of whom baldness could be asserted, whether accurately or not; in a context where there is no such entity, the sentence is hard to categorize as either true or false.

When a piece of text presupposes some proposition (potential fact), then there are various ways that the hearer of the text may react. If the hearer believes the proposition to be untrue, then the text may be deemed incomprehensible; if the proposition is believed to be true, the text may be regarded as comprehensible (even though the hearer may dispute the truth of the text itself); in the absence of information one way or another about the truth of the presupposed information, that presupposition may be accepted as true, since the speaker appears to be assuming it. For example, *My brother works in France* may lead a hearer to accept, in the absence of contrary evidence, that the speaker has a brother, a fact presupposed by the phrase *my brother*. This latter case – acceptance of presupposed material in the absence of conflicting knowledge – is known as *accommodation*.

These concepts will show up in our linguistic interface ([Section 3.9](#) and [Appendix A.1](#)) in two ways. First, our assumed models of text-interpretation and of inference will operate relative to some set of beliefs (what we will term a *situation*). Second, at each step of text-interpretation (roughly, each clause) the current situation will have to satisfy a set of *prerequisites* for the next piece of text; these are intended to cover classical presuppositions and any other requirements for discourse coherence.

3.7 Ambiguity

As will become clear in later chapters, *ambiguity* (of various sorts) is a central device in much verbally expressed humour:

Deliberate ambiguity will be shown to underlie much, if not all, of verbal humor.

(Raskin 1985: xiii)

At this preliminary stage, all we will do is propose some terminological conventions. Ambiguity occurs, loosely speaking, when a linguistic item (typically a sentence or portion of a sentence, although it could also be a discourse) has only one representation at one level (e.g. phonetically) but more than one representation at another level (e.g. semantically). These levels are assumed to be in some sort of ordered hierarchy; for example, Chomsky (1957:11) lists *phonemic*, *morphemic*, and *phrase structure*, while Winograd (1983:17) outlines a sequence of processing strata: *phonological*, *morphological*, *lexical*, *syntactic*, *semantic*, *reasoning*. The level at which the ambiguous item has more than one representation has to be higher in the usual hierarchy of linguistic strata – an item with one semantic representation but two different pronunciations is not deemed ambiguous. (As we shall see in [Chapter 9](#), some authors appear to assume a wider notion of ambiguity, based on *similarity* of phonetic strings, but this has never been defined clearly.) There is no standard usage within linguistics about labelling types of ambiguity. For example, ‘syntactic ambiguity’ seems to be named after the lowest level at which there are *multiple* representations. Pepicello and Weisberg (1983) and Pepicello and Green (1984) criticize Ben-Amos (1976) for referring to the similarity of sound between *red* and (the past tense form) *read* as ‘phonetic ambiguity’, since this names it after the level at which the representations are the *same*. They offer instead the label of ‘morphological’ ambiguity, on the grounds that the words have different morphological analyses. It is possible to imagine an analysis in which it was deemed to be *lexical* ambiguity, since different lexical items are involved. If one was considering written texts, one could argue for calling it ‘orthographic’ ambiguity. In spite of their criticisms, Pepicello and Green (1984) describe the uses of *blue* (as meaning ‘unhappy’ or as in ‘bluebird’) as ‘phonological’ ambiguity, despite the fact that the two forms are phonologically the same. The readings of *ground* (as a past participle of *grind* or a synonym for *soil*) display, they say, ‘morphological’ ambiguity. (Dienhart (1999: Section 3) gives a succinct and convincing critique of the confusions underlying some of Pepicello and Green’s classifications.) It is commonplace in linguistics to refer to the ambiguity in a sentence such as *Every man loves some woman* as ‘scope’ ambiguity, which does not really refer to a linguistic level, more to a particular aspect of the sentence. Binsted (1996) uses the term ‘pronunciation ambiguity’ to describe the case where the same spelling may be pronounced in more than one way (e.g. *wind*); this seems to be named after the *mapping* between written and phonetic levels. As Pepicello and Weisberg (1983) point out, the

term ‘semantic ambiguity’ is not very helpful, as so many of the types of ambiguity give rise to multiple meanings. It is not, as they claim, totally vacuous, as there is the further level of *pragmatic* ambiguity, at which a semantically unambiguous sentence could have different contextual effects (e.g. (24) at the end of this section).

One way to clarify this situation would be to adopt a convention that ambiguity is described in terms of *two* representational levels: one at which the item has a single representation, and one at which it has multiple representations. Thus *read* and *red* would be *phonetic-lexical* ambiguous, and also *phonetic-orthographical* or *phonetic-morphological* if these levels were present in the assumed framework. In general, if the levels are totally ordered, as they usually are in linguistics, then if an item is ‘*M-N* ambiguous’ where *M* and *N* indicate levels, it will also be ‘*M-K* ambiguous’ for any level *K* which is higher than *N*.

This could lead to cumbersome notation if uniformly applied; for example *Every man loves some woman* would be ‘unscoped-scoped’ ambiguous. In such cases, shorter mnemonics could be introduced, providing it was clear that they abbreviated a two-level characterization of the type of ambiguity.

We will adopt this terminology where helpful in our discussion of example jokes. However, our definitions of joke mechanisms do not rely on a formal definition of linguistic ambiguity, so we have not included the idea of an ordered sequence of multiple representation levels within our catalogue of linguistic facilities ([Section 3.9](#) and [Appendix A.1](#)). The notation will be used solely to make our informal discussion of data slightly more precise.

It is not difficult to find examples to show that virtually any semantic mechanism which can create ambiguity can be used in a joke. Consider (22).

(22) In New York, someone is knocked down by a car every two minutes. He’s getting pretty fed up about this.

A fairly conventional analysis of the semantics of (22) would say that the two meanings which are used to create the joke differ in the relative scopes of the quantifiers introduced by *someone* and by *every (two minutes)*; this is ‘unscoped-scoped’ ambiguity in our terminology.

(23) This was a mythical beast with the head of a lion and the body of a lion, but the body of *another* lion.

(paraphrase of a joke by Woody Allen)

Example (23) plays off the generic reading of *a lion* against the specific reading.

It is also plausible that some level of linguistic description beyond the sentence will be needed, where speaker's goals, intentions, beliefs, etc. are represented. For example, consider the joke (24).

(24) Diner: Waiter! There's a fly in my soup!

Waiter: Please don't shout so loudly – everyone will want one.

In this example, the humour is based on the waiter-character's (presumably deliberate) misconception about the nature of the speech act involved – is the diner-character complaining or expressing pleasure? This could be viewed as *semantic-pragmatic* ambiguity (see [Section 11.1](#) for fuller discussion).

3.8

The right level of abstraction

It is tempting to think that a humour-theoretic analysis of a verbally expressed joke will in general require a comprehensive linguistic analysis of the joke's structure, at all levels. While it is true that such a complete array of information should provide all the necessary information for humour analysis, it is misleading to imply that so much detail is always necessary. Much of the linguistic structure may be irrelevant to the humorous effect, and a detailed linguistic account of every aspect of the text may obscure the relevant relationships.

There are two separate questions that could be asked about the linguistic characteristics of a joke, roughly corresponding to 'means' and 'end':

- Effect: What is the end result of the text, in terms of setting up relationships between entities such as words, meanings, etc?
- Method: What linguistic devices are used to arrive at this effect?

(25) A lady went into a clothing store and asked 'May I try on that dress in the window?'

Well,' replied the sales clerk doubtfully, 'don't you think it would be better to use the dressing room?'

(Clark (1968) cited by Oaks (1994))

For example, it could be argued that what makes (25) a joke is the unexpected presentation of a hitherto hidden (and slightly risqué) interpretation of the earlier portion of the text (see [Chapters 5](#) and [8](#) for more details).

That would be a hypothesis about the *effect* which the text achieves, and the theory of humour would then be expected to have something to say about why sudden revelations of mischievous ideas are funny (see [Chapter 4](#) for a fuller discussion of such theories). This abstracts away from the fact that this effect is achieved by a particular linguistic *method*,

namely a variety of orthographic-syntactic ambiguity. It is quite possible that a very similar effect (hidden interpretation) could be achieved using some other kind of ambiguity; for example, in (6) orthographic-lexical ambiguity is used.

- (6) Do you believe in clubs for young people?
Only when kindness fails.

We, therefore, should distinguish between two different goals that a descriptive account of verbally expressed humour might have:

- to set out the essential relationships that must hold between representations (at various linguistic or knowledge levels) of a text in order for it to be a joke;
- to describe the available linguistic devices for achieving these effects.

A theory which sets out to define the central characteristics of humour should aim at least for the former, which is itself a difficult goal. The second aim (investigating methods) is also relevant if one wishes to carry out detailed predictions or testing of the more abstract ideas, or use them in (for example) some computer implementation (see Sections [2.12](#), [2.13](#)), but it is a logically separate issue.

There is no prior distinction between what counts as an effect and what is a method, as this will depend on the analysis offered. That is, humour analysts have to propose what they think are the appropriate abstract accounts of humorous phenomena (i.e. the essential effects). This implicitly assigns more concrete or more precise details to the role of methods. Although the discussion of (25) above suggested that we might want to analyse a joke in terms of ambiguity (effect) rather than itemizing the various linguistic devices for achieving ambiguity (methods), it is also possible to argue ([Chapter 8](#)) that linguistic ambiguity is itself acting as a *method* in order to allow multiple interpretations of an event (the *effect*). So the distinction is all in the eye of the analyst, and reflects the choice of which level of abstraction is the most suitable for summarizing the essential properties of the class of jokes.

There are numerous jokes which rely for their effect on phonetic-syntactic (or orthographic-syntactic) ambiguity. Oaks (1994) provides many examples from a variety of sources, including (25) and (26).

- (26) Postmaster: Here's your five-cent stamp.
Shopper (with arms full of bundles): Do I have to stick it on myself?
Postmaster: Nope. On the envelope.

(Clark (1968) cited by Oaks (1994))

If a precise account of the method of this joke were required, it would be necessary to describe the syntax trees and their parts. This particular example involves discontinuous dependencies involving relationships between non-adjacent elements. In (26), one syntactic analysis, if annotated fully, would have to show that the particle *on* is related to the verb *stick* as a verb-particle construct (similar to *stamp out*, *brush off*, etc.). This is relevant to the internal structure of the joke. Since the object phrase is prominalized as *it* (rather than *the stamp*) it must appear before the particle *on* (since *Do I have to stick it on myself?* would not be an acceptable way of saying this). If this surface ordering did not occur (e.g. *Do I have to stick on the stamp myself?*) the ambiguity, and hence the opportunity for the joke, would not arise. These details, however, are unlikely to be the concern of the humour theorist, since humour theory is more concerned, in a joke like this, with the *effect* – ambiguity – than the *method*. The creation of ambiguity can be delegated to whatever linguistic model seems helpful.

This also means that humour theorists need not agonize about more complex linguistic matters such as *unbounded dependencies*, a more severe form of discontinuous constituency which has significant implications for linguistic theorizing (Gazdar *et al.* 1985: Ch. 7). Although there will be jokes which include such constructions, their internal details are unlikely to be of interest from a humour point of view, unless we are considering *methods*.

We shall concentrate, in our analyses, on finding what we believe is the most illuminating and general definition of each class of joke that we consider; this may deal in abstract constructs such as ambiguity, or phonetic similarity. Usually, there will be a vast array of linguistic devices (methods) available for achieving these abstract properties (effects). In such situations, we shall not try to catalogue these devices, but shall assume that the relevant methods will usually be just those linguistic techniques which will achieve the required effects, with no specialization for humour. We wish to avoid obscuring our statement of the humour-relevant relationships by including – as some taxonomies of VEH do – a catalogue of routine linguistic ways in which attributes such as ambiguity or phonetic similarity may arise.

3.9

Summary of linguistic constructs

[Section 3.1](#) outlined the methodological status of the abstract linguistic notions that we will assume in our analyses of jokes, and [Appendix A.1](#) supplies a detailed formal statement. Here we shall give a very brief informal summary of these constructs, to facilitate understanding of later chapters.

Text strings can be either *phonetic* or *written*, and for each of these types there is a notion of *similarity*. Phonetic and written strings are related

systematically by mappings of *spelling* and *pronunciation*. There is a *lexicon* which links text strings to clusters of linguistic information, including a *part of speech*, such as NOUN, VERB, etc., and a *semantic item*. Text interpretation is modelled in an abstraction of left-to-right processing ([Section 2.10](#)) in which a *text-meaning mapping* consumes segments of a text string (loosely, sentences or clauses) from the start of the string, converting each in turn into a semantic representation. An ongoing representation of the meaning so far in the text is maintained, with presuppositions of each semantic chunk being checked as it becomes available. At each step of this process, there is also a form of *inference* which fleshes out the information in the text with logical implications, default assumptions, etc. Ambiguity can appear either as multiple outputs from the text-meaning mapping, or multiple results of the inference phase, and in both cases the various results are assumed to be ordered according to how 'obvious' they are in context.

The basic building blocks of the semantic representation are *concepts*, which can be *associated* with each other, and *propositions*, roughly true or false statements. The context for interpreting a text string is made up of a configuration of these, known as a *situation*, and, at each step in the text, the ongoing context for further processing is also a situation. Hence, meanings within a text can be said to 'be linked to' the current situation. There is also a primitive notion of an utterance being *contextually appropriate*: this complex (but undefined) term indicates that a statement/question/command is in some way suitable to the situation of use and is not bafflingly bizarre.

4

Incongruity and its resolution

We examine the most widely supported candidate for the role of ‘essential ingredient’ in humour.

4.1

What is incongruity?

If there is one generalization that can be extracted from the literature about humour, it is that humour involves *incongruity*. This point, with varying terminology, has been made by numerous authors: Keith-Speigel (1972) lists 24, the earliest from 1759. Morreall (1987) sees it in the writings of Hutcheson (1750) and Hartley (1810), and also in the much quoted remark of Kant that ‘Laughter is an affection arising from the sudden transformation of a strained expectation into nothing’ (Kant 1892) (see also Keith-Speigel (1972), Shultz (1976), Attardo (1994), Raskin (1985), Suls (1983), and with a slight re-wording – Wilson (1979)). Beattie is credited (e.g. by Keith-Speigel (1972), Shultz (1976), Rothbart (1976), Wilson (1979), Suls (1983), Raskin (1985), Attardo (1997)) with suggesting it even earlier:

Laughter arises from the view of two or more inconsistent, unsuitable, or incongruous parts or circumstances, considered as united in one complex object or assemblage, or as acquiring a sort of mutual relation from the peculiar manner in which the mind takes notice of them.

(Beattie (1776) quoted by Raskin (1985))

As is often observed (e.g. Keith-Speigel (1972), Shultz (1976), Wilson (1979), Suls (1983), Raskin (1985), Attardo (1994)), the idea was also proposed by Schopenhauer in 1819:

The cause of laughter in every case is simply the sudden perception of the incongruity between a concept and the real objects which have been thought through it in some relation, and laughter itself is just the expression of this incongruity.

(Schopenhauer 1883:76)

Paulos (1980) quotes Hazlitt (1819) as saying ‘The essence of the laughable is the incongruous, the disconnecting of one idea from another, or the jostling of one feeling against another’; Keith-Spiegel also cites Hazlitt as proposing incongruity, but as Morreall (1987) comments, and the quotations offered by Raskin (1985) show, Hazlitt had a wider view of factors contributing to humour. Freud (1966) also argued the case for incongruity, although this is usually overshadowed by his advocacy of a relief theory of humour.

A slightly more recent, and very influential, presentation of the incongruity approach was that of Koestler (1970), who defined the concept of *bisociation* as ‘the perceiving of a situation or idea, L , in two self-consistent but habitually incompatible frames of reference, M_1 and M_2 ’; see [Section 4.2](#) below. The intuitive and persuasive idea of incongruity remains current, appearing in work as varied as Attardo (1997) and Katz (1993).

The main problem, if we are to develop a detailed theory of verbally expressed humour, is that the notion of incongruity is not clearly defined. Schopenhauer did not set a good example, from a scientific point of view:

I shall not pause here to relate anecdotes as examples to illustrate my theory; for it is so simple and comprehensible that it does not require them, and everything ludicrous that the reader may remember is equally valuable as a proof of it.

(Schopenhauer 1883:77)

Consider the following definition, from a conventional dictionary:

incongruous: 1. out of keeping or place; inappropriate; unbecoming.
2. not harmonious in character; inconsonant; lacking harmony of parts.
3. inconsistent.

(Random House 1988)

None of the various authors cited above venture a definition which is more precise, detailed or formal than a typical dictionary entry, or which improves greatly on Beattie’s 1776 suggestion.

It is not even obvious that all the writers on this subject have exactly the same concept in mind. In addition to the quotations from Beattie, Kant and Koestler given above, we have:

the source of the ludicrous is always the paradoxical, and therefore unexpected,
subsumption of an object under a conception which in other respects is different from it ...
that is, a thing or event, which certainly can be subsumed under that conception, and
therefore thought through it, yet in another

Nerhardt then gives very brief outlines of three studies to test this hypothesis, but these are inconclusive, and he observes that 'we should exercise caution in attempting to operationalize the similarity proposed in the model'. Hence, Nerhardt's ideas are by no means a clear, workable, and empirically supported notion of incongruity-based humour. Although it is not inconceivable that these notions could be applied to some forms of verbally expressed humour, that would require a careful definition of the component concepts such as dimension and similarity in the case of jokes. That would, in turn, require some detailed analysis of data (jokes), which is exactly what we are concentrating on.

Latta (1999: Ch. 7) gives a lengthy diatribe on the general vagueness, inconsistency and incoherence of the 'definitions' of incongruity offered in the literature, focussing mainly on McGhee (1979), Paulos (1980), Shultz (1976), and Morreall (1987).

One consequence of the vagueness of terminology illustrated above is that a given suggestion may be interpreted in different ways. To Raskin, Kant advocates incongruity, but to Wilson he suggests a relief theory, while Morreall sees both in Kant's work. Bergson (1911) is seen by Wilson as a dualist, but by Keith-Spiegel and Raskin as arguing for incongruity. Hartley, another in whom Morreall sees incongruity, is classed by Keith-Spiegel as proposing a surprise theory. Maier (1932) is seen as a proponent of a *Gestalt* approach by Wilson, but an incongruity approach by Shultz.

Another aspect which has been remarked upon by many authors (e.g. Morreall (1983)) is that the incongruity must occur in a setting which is suitably secure, since if it signals or embodies danger then it will not seem amusing.

Related to this, but subtly different, is the need for the incongruity to be *licensed* in the sense of appearing in circumstances (e.g. the telling of a joke, or in a caricature) where incongruity is an accepted, playful notion. People will suspend their disbelief when hearing a story in a way that they will not when interpreting real life (cf. Rothbart (1976)).

Even considering this patchwork of informal and overlapping proposals, it is possible to discern two separate notions of incongruity, which might be labelled *static* and *dynamic*. The static form of incongruity can be a property of a particular situation or configuration of elements, or even an event if it is sufficiently brief that it is regarded as instantaneous. The dynamic form of incongruity is to do with a *sequence* of images or ideas occurring or being presented in a way which creates some effect on the hearer; that is, the oddity or incongruity does not involve a configuration of objects perceived all at the same time, but the temporal sequence of events or ideas creates the effect. Some authors seem to be focussed more on static incongruity (see, for example, the quotations above from Beattie, Schopenhauer, Clark, Keith-Spiegel, and Gruner), and others on dynamic

(see Kant, Shultz, and Nerhardt). Central to the dynamic variant is the concept of an *expectation*, created by a sequence of events or ideas.

Both static and dynamic incongruity appear within VEH. Static incongruity shows up in amusing verbal descriptions of situations or events (or even descriptions of amusing situations or ideas), or short pithy epigrams. In such cases, there is no need to establish particular expectations within the recounting. The incongruous effect may be dependent upon the vocabulary or phrasing used, since that may contribute to the sense of oddity, absurdity, or uncommonness, but there is no contribution from the passage of time in the describing or revealing of the scene.

In dynamic verbally expressed incongruity, the effect is created (or at the very least exaggerated) by the temporal sequence and manner in which the scene is described or revealed, with earlier stages of the text establishing expectations to be violated or tensions to be released (see more detailed exploration in Chapters 5 and 8).

There is a related, but logically separate, distinction to be made, between what Morreall (1983:62) calls ‘incongruity in things’ and ‘incongruity in presentation’, called here for brevity *inherent* and *presentational* respectively. That is, there may be some scenes or notions which are amusingly incongruous (inherent) or there may be those which are made funny by the way they are described or revealed (presentational). In real examples, there may be a slight blurring at the boundary of these two categories, as it could be argued that even when one perceives an idea or situation fairly directly, there is still an intervening layer of presentation. Morreall’s most straightforward example of incongruity in things – finding a bowling ball in his refrigerator – could be claimed to involve a presentation by whoever placed the ball there, but that is a different concept of presentation. The presentation, for Morreall in experiencing the event, was minimal, in that no one controlled his perception of the stimuli involved.

This informal taxonomy thus offers four subclasses of incongruity-based humour: static-inherent, static-presentational, dynamic-inherent, and dynamic-presentational. Static-inherent humour is when we perceive some situation all at once which strikes us as amusing, regardless of some method of conveying it by some third party (e.g. Morreall’s bowling ball, or the slogan ‘Gobi Desert Canoe Club’ (Attardo 1994)). Static-presentational humour is where a particularly constructed presentation of a situation is amusing, but does not require any development through time to achieve that effect (e.g. Oscar Wilde’s description of fox-hunting as the ‘unspeakable in full pursuit of the uneatable’). Dynamic-inherent humour is where we experience some sequence of events and the order of these events contributes to the amusement (e.g. Nerhardt’s weight-lifting experiment, or a slapstick scene in which the chaos builds to a climax or uses repetition). Dynamic-presentational humour includes all the humour normally referred to as incongruity-resolution ([Section 4.3](#) below).

4.2

Bisociation, frames, and scripts

The proposal of Koestler (1970), that humour is created by *bisociation* (see [Section 4.1](#) above), is a widely cited version of the incongruity theory. Even where Koestler is not referenced, the essence of his ideas recurs. As Norrick points out (citing commentary by Johnson (1976) and by Wilson (1979)): ‘many other writers both before and since Koestler (1964) have advanced bisociation theories of humour in various guises and in various names’ (Norrick 1986:227). The very prominent SSTH can be viewed as a bisociation-based approach (see [Chapter 6](#)).

If there is such a wide consensus that bisociation embodies the appropriate notion of incongruity, the question arises: why do we need to study the specific mechanisms of jokes? There are several reasons why the Koestler conjecture does not close down the debate even on incongruity-based theories of humour. We will now review three of the main reasons.

Undefined terms

There is still not a clear formal definition of Koestler’s terminology (‘frames’, ‘perceive in’, ‘habitually incompatible’) which would allow researchers to predict whether particular stimuli (e.g. specific texts) would count as manifesting bisociation or not. This is essential if bisociation theory is to become a predictive, falsifiable scientific theory, or if it is to be amenable to computer testing. Proponents of bisociation or its later versions might well argue that formal definitions have been provided, by the introduction to artificial intelligence of *frames* (Minsky 1975) and its various developments which use the term *scripts* (Schank and Abelson 1977), (Raskin 1985), concepts which are intuitively similar to a Koestler frame in more than just name. Minsky proposed the frame as a useful and psychologically plausible form of knowledge representation for a wide variety of subject matters and a range of tasks. Computer programs were written (e.g. Bobrow *et al.* (1977)) that purported to exemplify the frame-based approach. It would seem, therefore, that there must be a formal definition of a frame. In fact, there was no single accepted definition, and what the various implementations of frames (e.g. Roberts and Goldstein (1977), Bobrow and Winograd (1977)) had in common was that they were general-purpose data structures, generalizations of data types in more mundane programming languages such as Pascal. Indeed, some of the devices that were debated in the mid-1970s as ‘knowledge representation’ (e.g. default values, slot restrictions) are now routine parts of the general programming language Common Lisp (Steele 1990). A program which uses Common Lisp structures, or is written in an object-oriented programming language like Java (Arnold *et al.* 2000), could be argued to be frame-based. That is, the frame data structure can be seen to be so general that it makes very

little claim about how the knowledge is structured. The answer to ‘what is a frame?’ is ‘(virtually) anything’.

Frames versus general inference

A reasonable way to view Minsky’s central insight was that he was drawing attention to the effectiveness, and ubiquity, of knowledge which was both *specialized* (for the subject or task at hand) and *precomputed* (rather than being inferred on every occasion from general principles). There was also the use of *default information* in the absence of data to the contrary. (The justification for scripts (Schank and Abelson 1977) was similar, but emphasized sequences of events.) The choice of a data structure to represent such knowledge is something of a separate issue. In the mid-1970s, there was an immediate response to the frame-based approach from those who argued that some suitable form of logical inference would be capable of gaining the same effects as a frame representation (Hayes 1979). The central points were that the reasoning phenomena pointed to by Minsky (specialization, precomputation, defaults) are not incompatible with the use of logical inference. This case has become stronger as logics have become more sophisticated, particularly default logics which allow *nonmonotonic* or *default* reasoning (Reiter 1978, 1980), Turner (1984: Ch 5).

An analogous argument arises in the case of bisociation theories. Suppose it is proposed that a particular joke involves the bisociation of two frames, on the grounds that there is a conflict between two perspectives on the information in the text. Is this different from, or the same as, saying that the text supplies information which leads to conflicting inferences?

- (27) An ancient, wizened man accosts a lady of the night and inquires as to her rates. She replies, \$5 on the floor, \$10 on the couch and \$15 in bed.’ As he hands the hooker \$15, she remarks, Okay, once in bed,’ to which he objects, ‘No, three times on the floor.’

(Norrick 1986:239)

Commenting on (27), Norrick suggests that a schema set up by the early part of the text ‘conflicts’ with a schema supplied by the punchline. These observations could equally well be phrased as saying that the joke depends on the hearer making a default inference that the old man will be sexually weak, and that this (defeasible) inference is then overturned by subsequent information. It is not at all clear that these analyses differ in substance, and it is hard to envisage data which would distinguish them empirically. If that is the case, then using the terminology of frames is just a handy shorthand, with no substantive theoretical import. The crucial issue is: how does viewing something as frame-conflict tell us anything we would not glean from viewing it as some other form of logical clash (such

as the overriding of a default inference), or even just informally saying that there has been a clash between two perspectives?

Humour knowledge versus general knowledge

Articles advocating bisociation are not always explicit about the status of their frames with respect to general knowledge. Are frames entities which participate solely in the processing of humorous phenomena, or are they more general information structures used in other cognitive activities? Opting for frames being peculiar to the process of humour interpretation would raise certain questions. How is it that the joke-audience has a language-understanding mechanism (frames) which is used solely during jokes, and is different from the language-understanding methods used for other texts? This question is not restricted to language-understanding narrowly viewed, as sometimes the frame-clash occurs in propositions inferred from the meaning of the text; hence, this would entail a humour-specific reasoning mechanism. The postulation of a completely separate cognitive apparatus for perception and reasoning during humorous stimuli is hugely implausible. This is not to argue that we treat joke texts as if they were non-joke texts in every way; the point here is that it is bizarre to propose duplicating all our linguistic and world knowledge, and perhaps even introducing a different formal mechanism, for the processing of jokes. That position is mentioned here merely for completeness, but is almost certainly a straw man (although the suggestion by Raskin (1985:128) that having a sense of humour involves having more scripts available seems to hint at such a position). It is much more likely that proponents of bisociation assume that frame structures are suitable for general perception and cognition (as did the 1970s advocates of frames). In that case, we have to assess the Koestler formula from the point of view of *necessity* and *sufficiency* for humour. In what way do humorous bisociations differ from other alignments or clashes between frames that might occur in, for example, analogies, misunderstandings, or poetic figures of speech (metaphor, simile, etc.)? The following observations seem rather bisociation-like, but concern the semantic patterns underlying poetic devices:

two powerful but unrelated images are presented to us individually and we are forced to discover their relation.

juxtaposition seems to promise coherence and thus impels us to try to construct a coherence.

(Hobbs 1990:129)

For each of these other types of phenomena, one might expect that either they (all) involve the relating of two 'habitually incompatible' frames, or they are based on a different formal mechanism. But some misunderstandings

are humorous, others are not; some analogies may be amusing, most are not; some metaphors or similes may be amusing. If, on the other hand, Koestler and his successors are not proposing bisociation as a necessary and sufficient condition for humour, it is unclear what they are suggesting.

None of this is to suggest that the idea of bisociation has no merits, or that it is wholly incorrect. It is an interesting way of looking at humorous phenomena, being particularly convincing for humour where a collection of preconceived ideas about some situation is played off against another cluster of expectations. For example, BBC TV's *Monty Python's Flying Circus* presented a sketch depicting a football match between teams of famous philosophers, with the perspective of the commentator veering back and forth between a sporting contest and a philosophical debate.

With suitable development, bisociation might lead to a useful and insightful theory of humour. The point here is that this development is yet to be done. To offer bisociation theory as a theory of humour is to pass off as an oak tree something which is believed to be an acorn.

The most prominent bisociation theory of recent decades, Raskin (1985)'s SSTH, is one of the few which even starts to address the issues raised above. Raskin tries to make the notion of 'frame' more precise (as his version of the script), puts his structures and their associated operations forward as a general theory of language interpretation and inference, and elaborates on which kinds of frame-clash (*script opposition*) count as humorous. The SSTH therefore goes much further than work which merely asserts that particular jokes exemplify bisociation. However, we will argue in [Chapter 6](#) that the SSTH is still disappointingly vague.

4.3

Incongruity-resolution

4.3.1

Notions of resolution

One very common proposal which further explores the use of incongruity in humour is the *incongruity-resolution* (IR) theory (see Suls (1983) for discussion). This analysis states that incongruity alone is not sufficient to create humour, but this incongruity must be *resolved* (i.e. seen to be not as incongruous as first thought). This is done by a multi-stage process in which an initial incongruity is created, and then some further information causes the resolution of that incongruity. Shultz (1976) claims that this analysis is of 'immense heuristic value in accounting for vast samples of humour'. Ruch says, citing a number of writers including Shultz (1972) and Suls (1972), that there is general agreement about the existence of this two-stage structure in the process of perceiving and understanding humour' (Ruch 1992:31).

The way of creating and resolving the incongruity can be quite varied, as are the analyses presented in the literature.

(28) What is grey, has four legs, and a trunk?
A mouse on vacation.

(Rothbart 1977)

Rothbart suggests that in (28), the incongruity is presented by the question part – namely, that the question appears to have a surprisingly easy answer – and then resolved in the answer, which is also a surprise.

(29) Why did the cookie cry?
Because its mother had been a wafer so long.

(Shultz 1976)

With example (29), Shultz says that the answer is initially seen as incongruous, with *wafer* interpreted as ‘a type of cookie’, but then resolution occurs with the realization that there is an alternative interpretation, ‘away for’. Attardo (1994:128–31), in contrast, opts for ‘away for’ as the more prominent meaning, with ‘wafer’ as the alternative. Pepicello and Weisberg (1983) quote Shultz (1974) as saying that the word *wafer* actually *creates* the incongruity, and go on to claim that in an ‘informal study’ most people interpret *wafer* as ‘away for’ initially. Rothbart and Pien (1977) say that (29) has three incongruities: a cookie crying, the ‘surprisingness’ of the answer, and the cookie having a mother.

(30) Why did the elephant sit on the marshmallow? Because he didn’t want to fall into the hot chocolate.

(Rothbart and Pien 1977)

Rothbart and Pien also say that in (30), the question presents an incongruous situation, and the answer both explains (resolves) it and adds a new incongruity.

In contrast, Pepicello and Weisberg claim that the comparably structured joke (31) is simply a parody of a riddle, in which a question is posed which is impossible to answer.

(31) Why do elephants paint their toenails red?
So they can hide in cherry trees.

(Pepicello and Weisberg 1983)

Some authors seem to regard resolution as *removing* incongruity, others treat it as *explaining* incongruity. In some of these examples, the answer ‘explains’ but *reinforces* the incongruity. Attardo remarks (with respect to example (31)):

... the resolution of the incongruity is not a ‘real’ resolution: it does not get rid of the incongruity – it actually introduces more of the same

but because it has some distorted verisimilitude it is accepted playfully as a pseudo-resolution.

(Attardo 1994:144)

Suls analyses (32) in IR terms, since the answer ‘comes as a surprise’ but can be made to follow from the question with a little thought.

(32) If your son flunks out of school and is illiterate and anti-social, what can he grow up to be?
An Italian policeman.

(Suls 1977)

Attardo (1997) suggests that the concept of resolution can be identified with the *logical mechanism* of the General Theory of Verbal Humour, but this explains one undefined notion in terms of another (see [Chapter 6](#) for discussion).

Most expositions of incongruity-resolution have the incongruity occurring prior (temporally) to the resolution, with the resolution (the later event) producing the humorous effect. This sequence would not fit naturally with a joke in which an incongruity is implied by an already established line of reasoning, so that the resolution seems to precede the incongruity.

(33) Sitting on the side of the highway waiting to catch speeding drivers, a State Police Officer sees a car puttering along at 22mph. He thinks to himself, ‘This driver is just as dangerous as a speeder!’ So he turns on his lights and pulls the driver over. Approaching the car, he notices that there are five old ladies – two in the front seat and three in the back. The driver, obviously confused, says to him, ‘Officer, I don’t understand, I was doing exactly the speed limit! What seems to be the problem?’ ‘Ma’am,’ the officer replies, ‘You weren’t speeding, but you should know that driving slower than the speed limit can also be a danger to other drivers.’ ‘Slower than the speed limit? No sir, I was doing the speed limit exactly ... twenty-two miles an hour!’ the old woman says a bit proudly. The State Police Officer, trying to contain a chuckle, explains to her that ‘22’ is the route number, not the speed limit. A bit embarrassed, the woman grins and thanks the officer for pointing out her error. ‘Oh, thank you,’ she says, ‘It’s a good thing you didn’t see us on Route 119.’

(edited from www.jokes2000.com)

In (33), the punchline (the final utterance by the driver-character) does not seem to resolve anything. Instead, it implies a mildly incongruous situation (an elderly lady driving at breakneck speed) for which the logical justification (resolution) has not only already been supplied but also causes

the inference of the incongruity. It is hard to see how the resolution, in any plausible sense of the word, occurs after the incongruity.

Authors advocating the IR approach rarely make systematic comparisons between jokes and non-humorous texts. Many stories, even true accounts of actual incidents, involve incongruity in some sense, and this incongruity is often resolved, if we take this to mean ‘explained’ or ‘given a reason’. However, not all such occurrences are funny, suggesting that resolution of incongruity is not sufficient for humour.

4.3.2 *The INC-RES component*

The main empirical support for a notion of incongruity-resolution humour comes from psychological investigations of humour appreciation (Ruch 1992). Many studies have been carried out in which subjects are asked to rate a set of jokes and/or cartoons on several scales (e.g. for degrees of funniness). Statistical analysis of these results has shown, according to Ruch, that three factors consistently recur, two related to the abstract structure of the joke/cartoon (*INC-RES* and *NON*), and one describing joke content (*SEX*). Ruch’s gloss of *INC-RES* (mnemonic for incongruity-resolution) is:

Jokes and cartoons of this humor category are characterized by punchlines in which the surprising incongruity can be completely resolved. The common element in this type of humor is that the recipient first discovers an incongruity that is then fully resolvable upon consideration of information available elsewhere in the joke or cartoon.

(Ruch 1992:31)

Of *NON* (mnemonic for ‘nonsense’), Ruch comments that ‘the resolution information gives the appearance of making sense out of incongruities without actually doing so’:

... the punchline may (1) provide no resolution at all (2) provide a partial resolution (leaving an essential part of the incongruity unresolved) or (3) actually create new absurdities or incongruities.

(McGhee *et al.* (1990) quoted by Ruch (1992:32))

Ruch also states that these two components, together with the separate notion of sexual humour (*SEX*), ‘provide an exhaustive taxonomy in classifying jokes and cartoons at a general level’.

The claim that the relevant subset of jokes/cartoons involves the resolution of incongruity is an estimation made by the experimenters – the methods employed in these studies do not extend to decomposing the individual data items in a way which defines or displays the exact form

of ‘incongruity’. This is crucial from our viewpoint, as our goal is to develop a formal model of the salient components of humorous items. It would be interesting (but beyond the scope of the current volume) to take sets of items (jokes and cartoons) which are consistently classed as INC-RES or NON and try to develop a formal description which would embody the ideas glossed above, i.e. to apply the descriptive methodology we are advocating to this specific corpus of examples.

4.4 Summary

Incongruity, and incongruity-resolution, are two of the most central and frequently proposed concepts in humour theorizing, but there is no agreement about what they mean or about where they fit into a theory of humour.

In [Chapter 5](#) we will examine two particular variants of incongruity-resolution, one of which is so widely advocated that we will go on to discuss it further in [Chapter 8](#).

5

Two models of incongruity-resolution

We look in more detail at ways in which the incongruity-resolution approach could be fleshed out.

5.1

The two models

As shown in [Chapter 4](#) there is not exact agreement about what constitutes incongruity-resolution (IR). Nevertheless, there are some particular variants which are more clearly defined. We will now examine two of these, which are somewhat similar to each other. The first is the notion shared by many authors, such as Heller (1974) (as ‘non-retentional disambiguational puns’), Shultz (1976:13), Minsky (1980:10), and Paulos (1980); here it will be referred to as the *forced reinterpretation* (FR) model. The other model to be examined is the *two-stage* model of Suls (1972). (Although Suls claims that his model is also applicable to cartoons, we will consider it here only as an account of verbally expressed humour (VEH).) For both models, there is a certain basic arrangement, as follows. A joke is analysed as being in two main parts: the initial portion of text, the *set-up* (or *joke body* (Godkewitsch 1976)), and the second part, the *punchline*. The set-up creates no particular incongruity that the audience is aware of, but the punchline, at least initially, does not make immediate sense. However, a way is found to allow the punchline to be congruous (the resolution). Where the two variants differ is as follows:

- Forced reinterpretation: The set-up has two different interpretations, but one is much more obvious to the audience, who does not become aware of the other meaning. The meaning of the punchline conflicts with this obvious interpretation, but is compatible with, and even evokes, the other, hitherto hidden, meaning. The meaning of the punchline can be integrated with the hidden meaning to form a consistent interpretation which differs from the first obvious interpretation.
- Two-stage (Suls): The punchline creates incongruity, and then a *cognitive rule* must be found which enables the content of the punchline to follow naturally from the information established in the set-up.

There is no clear definition of the IR theory within the literature, but these two versions capture the essence of the idea, and most authors do not distinguish between the two, nor between them and the more general notion of IR (e.g. Attardo (1997:397)).

As has been implied above and observed by various writers, the general idea of IR may not be universally applicable to VEH, and it appears that each of the two models below fails to cover all the examples that might, informally, be argued to be IR phenomena. However, it is clear that there are at least subclasses of jokes which are very naturally accounted for, at least intuitively, by these models, so it is worthwhile examining them in more detail.

5.2

The forced reinterpretation (FR) model

5.2.1

The status of the FR model

The forced reinterpretation arrangement is very prominent as an embodiment of the IR approach, and is often equated directly with IR theory. Some authors seem to put it forward as a general account of all humour, or all verbally expressed humour. For example, the following quotations are not intended to describe some subclass of jokes, but are offered as general rules:

The comic effect arises when an alternative, non-favored and therefore non-expected interpretation is revealed, at the punchline, as the correct one.

(Dascal 1985:95)

the punch semantically reverses the sense we would expect from the build-up, and forces an unexpected sense to our attention.

(Norrick 2001:258)

The humorous effect comes from the listener's realization and acceptance that s/he has been led down the garden path ...

In humour, listeners are lured into accepting presuppositions that are later disclosed as unfounded.

(Dolitsky 1992:35)

Attardo (1994: Ch. 2) discusses the *isotopy disjunction model* (a synthesis of suggestions by various authors starting with Greimas), which seems to be a forced reinterpretation model, and Raskin's SSTH ([Chapter 6](#) below) can be seen as another variant.

5.2.2 *The basic units*

There are various entities involved in the typical FR account, including:

- SU_1 : the first (more obvious) interpretation of the set-up text
- SU_2 : the second (hidden) interpretation of the set-up text
- PL : the meaning of the punchline
- I : an interpretation formed by integrating the meaning of the punchline with SU_2 .

There are also various relationships and properties that are of interest, based on various observations made informally in the literature (the labels used here are invented as *ad hoc* mnemonics – part of the problem is that there is no standard terminology for these):

- OBVIOUSNESS: SU_1 is more likely than SU_2 to be noticed by the reader.
- CONFLICT: PL does not make sense with SU_1 .
- COMPATIBILITY: PL does make sense with SU_2 .
- CONTRAST: there is some significant difference between SU_1 and SU_2 (or possibly SU_1 and I).
(This was called COMPARISON in Ritchie (1999).)
- INAPPROPRIATENESS: I is inherently odd, eccentric or preposterous, or is taboo, in that it deals with matters not conventionally talked of openly, such as sexual or lavatorial matters, or forbidden political sentiments. These differ in terms of which norms are being flouted: those of everyday logic – leading to ABSURDITY – or those of socially acceptable discourse – leading to TABOO effects.

Notice that CONFLICT and COMPATIBILITY may in principle rely on a different meaning of 'does make sense with', and may not simply be opposites or negations of each other. CONFLICT is the symptom that the punchline meaning does not fit in, thus signalling a need for reinterpretation; COMPATIBILITY describes the relationship between punchline meaning and set-up meaning in the eventual resolution. These could, in principle, rely on different factors.

Although some or all of these relations or properties are sometimes proposed as the *essential* ingredients in IR humour, it is possible to find texts which, although plausibly of the general form of an FR joke, display only some of these relationships, yet which are humorous. For example, (25) relies on a (hidden) ambiguity regarding the structure of the initial question, and this is disambiguated in a surprising way by the answer; that is, it falls within the FR model.

- (25) A lady went into a clothing store and asked ‘May I try on that dress in the window?’
Well,’ replied the sales clerk doubtfully, ‘don’t you think it would be better to use the dressing room?’

Although it could be argued that the hidden meaning is ABSURD, there is no evidence that it additionally has some significant CONTRAST with the more obvious meaning (other than just being different from it). Similar remarks apply to (26).

- (26) Postmaster: Here’s your five-cent stamp.
Shopper (with arms full of bundles): Do I have to stick it on myself?
Postmaster: Nope. On the envelope.

Hence, not all of these relations or properties are necessary to create humour, even within the FR class of jokes.

The extent to which the hidden interpretation is INAPPROPRIATE (or has other suitable properties) may vary between jokes, as may the degree of CONTRAST between the interpretations. It is also conceivable that the sharpness of the CONFLICT between punchline and the first interpretation may vary, as may the OBVIOUSNESS of that interpretation. A more detailed version of the FR model may involve a complex disjunction of conditions, where each condition is itself a conjunction of thresholds for certain properties. For example, it may be that a text is humorous if either its CONTRAST level is above a certain threshold or the INAPPROPRIATENESS of the hidden interpretation reaches some other minimum level.

5.2.3

Previous variants compared

All discussions of FR-style jokes admit, as pre-theoretic (or informally theoretic) entities, the items listed in [Section 5.2.2](#) above: SU_1 , SU_2 , PL , I . It is therefore reasonable to ask how such theoretical commentaries can be stated as properties of, or as relationships between, these basic entities. The lack of standard terminology means that it is sometimes hard to be sure what types of entity a proposed relation is between, and whether one author’s incongruity is comparable to another’s. The five relations/ properties used here (CONFLICT, etc.) are chosen as some commonly occurring notions, to organize this discussion.

Confusingly, several of these constructs (particularly CONFLICT, CONTRAST, and INAPPROPRIATENESS) are referred to as ‘incongruity’, thereby obscuring the relationships. Some of these apparent differences of opinion may be more an artefact of the lack of precise terminology or definitions, in that the CONFLICT between punchline and set-up may be conflated with, or confused with, CONTRAST between the revised interpretation and

the initial interpretation. Shultz says ‘the incongruity consists in the relation between the last line, or punchline, and the part that precedes the last line’ (Shultz 1976:13) – i.e. CONFLICT. Of example (30), Rothbart and Pien (1977) observe that the resolution involves an elephant sitting in a cup of hot chocolate, and refer to this as ‘an incongruous situation’ (i.e. ABSURDITY), and they suggest that this residual ‘incongruity’ is a potential problem for IR theory, in that the oddity has not been eliminated (‘resolved’). Alexander (1997:16) says that a crucial part of joke perception may be ‘an incongruity between the punchline and what comes before’ – i.e. CONFLICT. Freud (1966)’s proposals can be viewed as opting for (a form of) INAPPROPRIATENESS as constituting incongruity.

Raskin’s SSTH (see [Chapter 6](#)) has a variant of CONTRAST as the central factor, namely script opposition. The SSTH also seems to package INAPPROPRIATENESS into script opposition, by regarding this (or other properties of the second interpretation) as relative to the first interpretation, rather than as inherent properties of the less obvious meaning. Giora (1991) proposes that the main mechanism in FR-style jokes is that the punchline provides a *marked increase in informativeness*, and that the interpretations differ in their ‘markedness’. The notion of markedness is not formally defined, but this seems to be tackling CONFLICT, OBVIOUSNESS, and possibly CONTRAST. De Palma and Weiner (1992) discuss a notion of *accessibility* for the meanings of words, thus addressing OBVIOUSNESS. Attardo (1997) suggests that SSTH’s ‘script opposition’ could be defined in terms of Giora’s and De Palma and Weiner’s concepts, thus replacing a description of CONTRAST with an amalgam of accounts of CONFLICT and OBVIOUSNESS.

5.2.4

Building on the FR model

The FR account does not particularly illuminate one of the central concepts of IR theory (and of much theorizing about humour): *what kind of incongruity is funny?* What the model does do is strip away some of the aspects of how this particular genre of joke *conveys* the incongruity. In this way, it reduces the research problem to a set of component subproblems. Most of the abstract concepts outlined above depend upon knowledge about language and (via linguistic meaning) the world, and research into language processing will illuminate these parts of the model. The residual subproblems are therefore:

- OBVIOUSNESS: What makes one potential interpretation more obvious than another?
- CONFLICT: In what way(s) must a piece of text be related to preceding discourse in order to stimulate a search for another interpretation?

- COMPATIBILITY: What does it mean for a punchline to fit in with ('be resolved with') a set-up interpretation?
- CONTRAST: What does it mean for two interpretations to differ in an amusing way (as opposed to merely being different)?
- INAPPROPRIATENESS: What factors make an interpretation inherently more amusing?

Of these, those concerning OBVIOUSNESS, CONFLICT, and COMPATIBILITY are to some extent more general research issues in linguistics and language processing, and the simplest initial assumption to make is that no special mechanisms are needed which are peculiar to humour.

However, jokes may (in common with non-humorous stories) require a certain semantic or pragmatic licence, so that the audience can 'suspend disbelief': jokes can feature entities which would qualify as semantically ill-formed in some descriptions of the world, such as talking kangaroos or walking cauliflowers. These semantic oddities are not necessarily a source of CONFLICT.

Definitions of several of the necessary concepts might be forthcoming from other disciplines: research into the interpretation and assimilation of sentences within a discourse (CONFLICT, COMPATIBILITY); semantic research, studies of belief systems and of social attitudes (CONTRAST, INAPPROPRIATENESS); sociology, psychology (INAPPROPRIATENESS); psycholinguistic and computational research into parsing and semantic interpretation of texts, particularly ambiguity and its resolution (OBVIOUSNESS). If all these building-blocks could be defined independently of humour theory, then defining the FR model of humour would reduce to the following: what combinations of these factors result in humour?

On the other hand, if it is necessary to have specific definitions of some of these concepts (for example, if only certain types of CONTRAST result in humour), then the humour-theory account of FR jokes would also include these definitions (cf. [Section 2.14](#)). If we accept that something like the FR model applies to a wide range of jokes, then defining CONTRAST and INAPPROPRIATENESS, and determining the overall combination of these factors necessary and sufficient for humour, are central to humour theory.

This is not to say that the FR model does not, in its present state, contribute to our understanding of humour. Many authors have agreed that the humorous effect usually depends upon the incongruity being brought to the attention of the audience abruptly (e.g. Katz (1993, 1996)), and the FR mechanism is one way to do this. Hence, it offers an answer to that particular subproblem within a theory of humour.

In [Chapter 8](#) we will examine the FR approach in more detail.

5.3 Suls' two-stage model

5.3.1 *Outline of the processing*

Suls (1972) outlines an IR-style processing model, which could be summarized, from the viewpoint of the audience, thus:

- as text is read, make predictions
- while no conflict with predictions, keep going
- if input conflicts with predictions:
 - if not ending – puzzlement
 - if it's the ending, try to resolve:
 - * no rule found – puzzlement
 - * cognitive rule found – humour.

This differs from the FR model in that Suls does not demand that any ambiguity be present in the set-up, which allows him to cover examples such as (32), (34), and (35), where it is hard to argue that there is any ambiguity in the set-up.

(34) O'Riley was on trial for armed robbery. The jury came out and announced, 'Not guilty.'
Wonderful,' said O'Riley, 'does that mean I can keep the money?'

(Suls 1972)

(35) Fat Ethel sat down at the lunch counter and ordered a whole fruit cake. 'Shall I cut it into four or eight pieces?' asked the waitress. 'Four,' said Ethel, 'I'm on a diet.'

(Suls 1972)

On the other hand, in examples where there is clear use of ambiguity – as (6), (25), (26) – Suls must postulate a cognitive rule which simulates the effect of reinterpreting the set-up. (There is a puzzling passage on page 96 of Suls (1972), where Suls seems to imply, in discussing work by Kagan (1967) with infants, that his model differs from Kagan's in that his system does indeed reinterpret the set-up, contrary to both his algorithmic diagram and the implications of his examples.)

Suls gives only a broad outline of his model, not much more detailed than the summary given above of the FR model, and his use of the term 'incongruity' is not very precise, although it seems to correspond, in the informal terminology of [Section 5.2](#), to CONFLICT:

the perceiver finds his expectations about the text disconfirmed by the ending of the joke
In other words, the recipient encounters an incongruity – the punch line. (p. 82)

humor derives from experiencing a sudden incongruity which is then made congruous.
(p. 82)

the ending of the joke [our (34)] is unexpected and incongruous ... but can be so interpreted as to make sense. (p. 90)

Incongruity of the joke's ending refers to how much the punch line violates the recipient's expectations. (p. 92)

(Suls 1972)

5.3.2 *Prediction*

A further factor which is often discussed in the context of IR-style jokes (and elsewhere) is surprise, or *violated expectations*, but the term 'expectation' also lacks a precise definition. In some cases, it may just mean the 'expectation' that the more OBVIOUS interpretation is the correct one.

Rather than simply allowing for a CONFLICT relation between punchline and set-up, Suls suggests that in (34) there should be a CONFLICT with some *prediction*: 'It is then predicted that he will say "Does that mean I can go now?"'. However, there is no evidence that there is some specific prediction: Suls does not show that this remark rather than any other (e.g. 'Thank you') is predicted. Suls may feel the need to follow this line of reasoning owing to his definition of incongruity (see above).

Shultz uses a similar definition to Suls, implying that this captures the intentions of Schopenhauer (1819), Kant (1892), Maier (1932) and Koestler (1970): 'Incongruity is usually defined as a conflict between what is expected and what actually occurs in the joke' (Shultz 1976:12). (Notice that this is not identical to the gloss of 'incongruity' given by Shultz elsewhere in his article – see quotation in [Section 5.2.3](#) above.)

It may therefore be necessary to consider two possible subtypes of the CONFLICT relation: one sets the punchline against the (more OBVIOUS meaning of the) set-up, and the other sets the punchline against a predicted interpretation (derived from the set-up).

Since Suls attributes the humorous effect to the discovery of a cognitive rule that explains an apparent incongruity, the clash with a predicted ending must be, in our terms, CONFLICT rather than CONTRAST; that is, it serves to detect the punchline, but not to produce the humorous effect. Shultz, on the other hand, seems to be postulating CONTRAST (a humour-creating discrepancy) with a predicted ending. We will return to the issue of prediction in [Section 8.4](#).

5.3.3 *Questions arising from the two-stage model*

Suls quotes approvingly from authors who observe that perception (including language understanding) is an active process, so presumably he

would accept that even in a smooth-flowing congruous text, the perceiver may be using cognitive rules to assimilate each sentence in turn. A cognitive rule is, after all, a very general construct: ‘a logical proposition, a definition, or a fact of experience’ (Suls 1972:82). However, Suls gives no indication of what would push a potential punchline over the threshold between normal rule-based comprehension and incongruity-driven reasoning. Moreover, he gives no discussion of the distinction between a punchline simply not making sense at first (a misunderstanding) and being humorously incongruous. He makes no mention of anything corresponding to our INAPPROPRIATE and – because he is not postulating an ambiguous set-up – CONTRAST is not applicable. Suls also does not explore the possibility that the nature of the resolution (i.e. the content of the cognitive rule) may contribute to the humour, although he suggests that the complexity of the reasoning needed may be a factor. His claim is that an incongruous punchline that can be found to make sense by finding a suitable cognitive rule is both necessary and sufficient for humour.

Like the FR model outlined earlier, this model leaves the hardest parts still to be done, and is more of a sketch of a ‘delivery’ method than an explanation of incongruity in humour.

The residual questions from Suls’ framework are:

- How surprising must a portion of text be to count as a punchline?
- What is a ‘cognitive rule’?
- What types of cognitive rule give rise to a humorous resolution rather than merely the patching up of a misunderstanding?

The last of these questions could be seen as being the central issue of much of humour research. It is quite common for writers on humour to observe that jokes often rely on a ‘local logic’ or ‘quasi-logic’ (Attardo 1994: Section 4.0.2): some distorted form of reasoning which is close to, or analogous with, sound reasoning, but which is sufficiently bizarre to produce humour. Inspecting Suls’ examples (particularly (34) and (35) above), it is plausible that the humorous effect relies heavily on some faulty form of reasoning.

5.4

Comparing the two models

At first sight, the two variants discussed above seem similar, but closer inspection shows interesting differences:

- The two models cover, or attempt to cover, different subclasses of joke. The FR model requires an ambiguous set-up, whereas the two-stage model makes no mention of ambiguity, and may even be unable to handle it.

- The FR model decomposes the humorous effect into slightly simpler concepts, and so at least starts to address ‘incongruity’.
- The two-stage model relies on some (undefined) form of ‘humorous logic’, and so leaves the difficult problem of ‘incongruity’ relatively untouched.

Another important observation is that there are still further jokes which have been postulated to be incongruity-resolution jokes ((29), (30)), but which are not covered by either of these approaches. (Although (29) involves ambiguity, the ambiguity is not in the set-up; see [Chapter 9](#)). That is, the two ‘models’ discussed in this chapter describe two *subclasses* of IR humour.

Suls’ two-stage model is presented entirely procedurally: he gives an algorithm for joke-interpretation, not a structural definition of the class of joke. Accounts of the FR model in the literature vary, but it is quite common for them also to be procedurally stated; see discussion in [Section 2.10](#).

5.5 Summary

Suls’ model does not seem to have been taken up in the three decades since its publication. The FR model recurs in proposals within the literature, although terminology may vary, and it also forms the core of the highly influential SSTH (see [Chapter 6](#)). We shall therefore discuss it further in [Chapter 8](#).

6

The General Theory of Verbal Humour

We examine and criticize an influential set of proposals regarding verbally expressed humour.

6.1

Introduction

The Semantic Script-based Theory of Humour, SSTH, (Raskin 1985) has been an extremely influential and widely cited contribution to the search for a theory of verbally expressed humour. The SSTH led to the General Theory of Verbal Humour, GTVH (Attardo and Raskin 1991), which is still being refined (Attardo 1997, 2001, Attardo *et al.* 2002, Attardo 2002). The aim of this chapter is to examine both the SSTH and the GTVH, and to point out some ways in which they are flawed. Much of this criticism will be methodological, questioning whether these are in fact substantive theories of humour. Empirical justification will be less central to the discussion, but, as we will argue, this is because these theories seem to make very few testable claims.

Our emphasis on the SSTH/GTVH is prompted by the fact that these ideas have become very prominent, and are treated by many writers as forming the received theory of (verbally expressed) humour:

[The SSTH] has been the most influential work in humour research in its recent history. The GTVH has become widely accepted across the various disciplines involved in humour studies as a theoretical basis for research. It has gained, despite its faults and lacks, a canonical status such that it is nearly impossible to find published humour research that does not make reference to it in some way.

(Rutter 1997)

The SSTH/GTVH is also one of the few attempts to approach verbally expressed humour in a systematic and theoretical fashion, and as such is to be welcomed. Attardo (2002) has claimed very strongly that the SSTH/GTVH is a formal, formalizable and falsifiable theory. We will argue that this is not the case.

6.2 The SSTH

Although the Semantic Script-based Theory of Humour (SSTH) was originally claimed not to be an incongruity-resolution theory (Raskin 1985, Attardo and Raskin 1991), this classification has since been accepted (Attardo 1997, 2001).

The SSTH states that the meaning of the text of a joke can be represented as a script (or an arrangement of scripts), where a script is a structured configuration of knowledge about some situation or activity (see [Section 4.2](#)). A joke consists of an initial portion where there are two possible interpretations, so that two possible (configurations of) scripts can be associated with that text. However, one of these possible interpretations is more obvious and is the one naturally perceived by the reader/ hearer, with the other meaning passing unnoticed initially. The final portion of the joke draws this other interpretation to the reader/hearer's notice, suddenly and in a potentially surprising manner. This is an instance of the FR model reviewed in [Chapter 5](#) (see the many citations Giora (1988) gives to support just this point).

The main distinctive component of the SSTH (compared to the traditional FR model) is that the two scripts (i.e. the interpretations) must, if the text is a joke, be *opposed* in some way (*local antonyms*). There are three very abstract oppositions (actual/non-actual, normal/abnormal, possible/impossible) and these are then instantiated as more specific oppositions (e.g. sex/no sex, good/bad). That is, 'oppositions' fall into a hierarchy of varying degrees of generality. In terms of our discussion in [Chapter 5](#), Raskin is opting for a version of CONTRAST as the critical humour-creating ingredient within an FR framework. Leaving aside some unclear proposals in Attardo (1997), the most recent account of script opposition is offered by Attardo *et al.* (2002) and Attardo (2002). A script is deemed to be a structure in which certain parts (intuitively, those more central to its meaning) would be *foregrounded* (or *in focus*). Two scripts are then opposed if they differ in some part of their foregrounded material. The GTVH is then a development of the SSTH.

6.3 The GTVH

The central idea of the General Theory of Verbal Humour is that a joke depends on contributions from six different *knowledge resources* (KRs). The KRs are:

- Script opposition (SO): the central idea of the SSTH, that a text must be interpretable as two opposing' scripts, is also at the heart of the GTVH.

- Logical mechanism (LM): the *logical mechanism* ‘accounts for the way in which the two senses (scripts) in the joke are brought together’ (Ruch *et al.* 1993:125). A joke may have one or more associated LMs.
- Situation (SI): the *situation* of a joke is, informally speaking, the setting, in terms of characters, objects, location, etc. described in the text.
- Target (TA): the *target* of a joke is what might be known informally as the ‘butt’; that is, the person, group or entity which is being ridiculed, attacked or presented in a negative light. This is optional, as some jokes do not have targets.
- Narrative strategy (NS): the NS is a broad classification of the linguistic structure or style for presenting the joke.
- Language (LA): this source supplies the actual linguistic units (words, etc.) used in the text.

These KRs are strictly ordered in the sequence shown in the above list, in the sense that any decision taken by a KR affects or constrains the options available to any KR lower down in the chain. Hence, the choice of SO (the topmost KR) limits what can happen in all other KRs, and the choices by LA (the lowest KR) are informed by the selections of all the other KRs.

Attardo and Raskin exemplify these different factors using variants on a ‘light bulb’ joke. They start from the joke (36):

(36) How many Poles does it take to screw in a light bulb?

Five. One to hold the light bulb and four to turn the table he’s standing on.

(Attardo and Raskin 1991)

This is characterized as having the following parameter values:

SO:	dumb
LM:	figure-ground reversal
SI:	light-bulb changing
TA:	Poles
NS:	riddle.

(Or, in Attardo (2001:73), the SO is ‘normal/abnormal;smart/dumb’, the NS is ‘question and answer’ and the LA is “‘How”, “many”, etc.’).

They compare this with other possible variants of the joke, or even different jokes, showing how these other examples can be seen as differing in, or sharing, the values for certain parameters. For example, (37) has the same values for SO, SI, TA and NS, but different LM (*false analogy*):

(37) How many Poles does it take to screw in a light bulb?

Five. One to hold the light bulb and four to look for the right screw-driver.

(Attardo and Raskin 1991)

Beyond this, the essential position regarding scripts and the creation of humorous effects is inherited directly from the SSTH.

6.4 Discussion

6.4.1 *The definition of the basic concepts*

It is not easy to extract clear definitions from the SSTH/GTVH literature for any of its central terms. It is not unreasonable for a theory to change over a period of years, but the imprecision is greater than might be accounted for by such development. We will try to assess what seems to be the latest (2002) consensus on the theory's details (Attardo *et al.* 2002, Attardo 2002), rather than attempting to chronicle all the changes that have taken place.

To some extent, all of the six GTVH parameters are under-defined, but in some cases this is more of a problem than others. TA (target) and NS (narrative strategy) are presented informally, with no precise definition, but their intuitive meanings are relatively clear (although Attardo (1998) surprisingly uses *joke* as a possible value for NS). However, the other knowledge resources are rather obscure. As SO and LM seem to be central to the humorous effect, this is particularly unfortunate. We will consider scripts, SO, and LM in turn.

Scripts

In considering the SSTH notion of script, the various metatheoretical issues raised in [Section 4.2](#) are relevant. If the SSTH is to be a substantive theory of humour, then we need to know what counts as evidence for the abstract notion of script, as opposed to some more vague and general term such as 'interpretation' or 'information'. Simply observing that text-understanding relies on stereotyped knowledge about everyday events is not sufficient. Raskin's informal analyses of jokes in terms of scripts could be re-written with the term 'interpretation' replacing that of 'scripts'.

Published accounts of SSTH scripts use such general terms (directed graphs, or slot-and-filler data structures) that a script could be any type of knowledge structure whatsoever. There is sometimes an indication (Attardo 2001:3) that every script has one node associated with a word/ morpheme whose surface form (written or spoken) appears in the text, so

that a script acts as the lexical knowledge about that word. However, in 2002 it is made clear that not all scripts are linked to lexical items.

In 2002, a script has an internal hierarchy, so that some parts are more *salient* or *foregrounded* than others, an enhancement which (when fully formalized) will give slightly more content to the concept. On the other hand, the nodes of a script can be individual letters of the alphabet (in order that the similarity between the two phrases in a pun can be described as a comparison of two scripts). This moves the notion of ‘script’ closer to ‘general data structure’.

Hence the claim that the content of a portion of a humorous text is representable using scripts says little more than that it is representable in some finite symbolic form, an assumption made by much of modern linguistics and artificial intelligence.

If the interpretation of a text corresponds to some configuration of several scripts (as seems to be implied in recent writings), it is not clear which scripts are to be opposed. This seems to need some notion of a dominant or central script within each of the two interpretations.

Script opposition

There are, in the SSTH/GTVH literature, various ways of characterizing what is meant by SO (script opposition). Raskin (1985) stated that the opposing scripts had to be local antonyms, but this expression lacked a non-circular definition. Attardo (2001:18–19) suggests that there are *semantic axes* along which items can vary. Local antonyms are items which have different values along these axes.

There is also a descriptive approach to SO, which involves sketching criteria for opposition in terms of what the two scripts *describe*; that is, their denotation in the world depicted by the text (or joke): actual/non-actual, possible/impossible, etc. Much of the writing about scripts, including Raskin’s original exposition, emphasizes the descriptive route, with the terminology of ‘actual/non-actual’, etc. The weakness of this is that there is no theoretical definition of what it means for a script to describe scenarios which are actual/non-actual, normal/abnormal, possible/impossible. Even if it were intuitively clear what it meant for a text to describe actual/non-actual contexts, the role of scripts would still be obscure. As we have indicated in [Section 2.5](#), a theory should show how its underlying abstract entities capture a pattern in the data.

A third account of SO is a formal, structural approach (Attardo *et al.* 2002, Attardo 2002), which consists of formal statements about relations between the structures which constitute the scripts. Two scripts, represented as directed graphs or as slot-filler sets, *overlap* (a necessary condition for humour) if they have parts which are ‘foregrounded’ – in the sense mentioned above – and are the same; for example, in the two script

structures posited for (18) (in [Section 3.4](#)), the VISIT node is an overlap. Opposition between two nodes is then defined as *being different* and *being foregrounded*. It appears that the foregrounded elements are the central, distinctive core of the script, in contrast to peripheral and highly general parts (e.g. that some participant is human). So two scripts are opposed if they differ in their non-peripheral elements.

As we noted, the SSTH is a variant of what we have called forced reinterpretation (FR). Hence, the 2002 GTVH seems to say that, in a text of the FR sort, if the two interpretations differ in non-peripheral content, then the text is humorous. While this is much more clearly defined than the older notion of SO, it does raise the issue of empirical adequacy – see [Section 6.4.4](#).

It is not clear how these three approaches fit together. Attardo (2001) uses local antonymy as the measure of difference between foregrounded parts of two scripts, but Attardo *et al.* (2002) and Attardo (2002) opt for simple non-identity as the notion of difference. What we have dubbed the descriptive account (actual/non-actual, etc.) is retained, apparently as a classification of types of SO but not as the central definition.

Attardo (2001) also states that he believes in the correctness of his proposals for SO in Attardo (1997), but these are completely different from those summarized above. Attardo (2001:98) alludes to the notion of a ‘humorous SO’, which implies that some script oppositions are not humorous.

Logical mechanism

The first difficulty regarding the logical mechanism (LM) is that its functionality is unclear: where is its formal role in the framework? In particular, does the assignment of a particular LM:

- describe how the two scripts are related?
- indicate what is odd about the events described in the resulting interpretation (script)?
- state how the scripts (or other information) are conveyed by the text?
- something else?

GTVH articles cite example jokes along with the name of the LM(s) involved, as in the assignment of *figure-ground reversal* to (36) in [Section 6.3](#).

(38) Who supports Gorbachev?

Oh, nobody. He is still able to walk on his own.

(Attardo and Raskin 1991)

(39) Why does a donkey eat thistles?

Because he’s an ass.

(Attardo and Raskin 1991)

Both (38) and (39) are classed as *juxtaposition* (of situations), but what that means is not spelled out; the SSTH seems to say that *any* joke will involve two scenarios (scripts) being pushed together.

The *garden-path* LM (Attardo and Raskin 1991) seems to involve the audience being led to expect one interpretation whereas another is suddenly revealed as the correct reading. This is very puzzling, as the SSTH seems to analyse *all* jokes as garden-paths, informally speaking (in the sense that there is a ‘hidden’ interpretation of the text which is suddenly revealed at the end), but the GTVH now postulates that there is also the possibility of a ‘garden-path’ LM.

Attardo emphasizes that the LM is not some static property of the text which enables resolution to occur, but the process which happens in resolution: ‘... *the LM of a joke is the resolution of the SO (incongruity)* ... LM is the resolution itself’ (Attardo 1997:409, italics in original). However, Attardo gives no definition of what a ‘resolution’ is. His sole example is a ‘semantic rule’ which causes the two conflicting (opposed) semantic items to be forced together, and it is not clear in what sense this is a ‘resolution’, nor is it easy to generalize from this very simple example. In the more canonical SSTH joke (e.g. (18) in [Section 3.4](#)), the mechanism that forces the alternative meanings together, and creates the conflict, is the fact that both meanings are plausible interpretations of the set-up part of the joke, which is very much an inherent part of the SSTH/GTVH mechanism, not one of its parameters. Typical accounts of resolution (Attardo does not cite any particular one as definitive) are informal, and do not illuminate this point. Attardo *et al.* (2002) quote the study by Paolillo (1998) which suggests LMs for each of 800 *Far Side* cartoons (Larson 1982, 1983, 1984, 1985). It is interesting that Paolillo did not feel a need to redefine the GTVH framework to cover these non-verbal examples, possibly feeling that it was obvious where the LM fitted in a theory of cartoon humour. Paolillo glosses each LM type in English, usually (but not always) stating the definition in terms of some manipulation or configuration of scripts. Some of these definitions do not appear to be specific to humour, and describe semantic or conceptual arrangements that might be conveyed by ordinary text, such as the following:

Consequence: a situation representing a consequence of some event is represented, leaving a prior series of events to be inferred.

Juxtaposition: two scripts are presented simultaneously in the same situation.

Sequence: a temporal ordering is imposed on the two scripts.

(Paolillo 1998)

Sometimes, an LM appears to be classifying what is odd or absurd about the situation described in a joke, and in other cases it seems to be classifying the method of conveying this information to the audience.

The 2002 conjecture is that ‘a significant class of LMs can be accounted for formally as mappings [between scripts]’ (Attardo *et al.* 2002), but this has not been developed very far.

While it is reasonable for more empirical work to be needed to work out which LMs actually occur in jokes, it is odd for a supposedly formal theory of humour to have a central construct which, more than a decade after its introduction, still does not have a clearly defined function (i.e. what it connects to what) within the theoretical apparatus. (Raskin (1995) speculated that the LM was not needed in the GTVH.)

A further problem is that no indication is given of how the LM is supposed to contribute to the humorous effect. The (necessary and sufficient) definition of when a text counts as a joke is retained from the SSTH, and it makes no mention of LMs. It is not explained if a text can be humorous without an LM, or if a non-humorous text can have an LM.

Overall, the LM seems to be a very versatile category in which almost any aspect of a joke can be placed.

6.4.2

Parameters without structure

The ‘knowledge resources’ in GTVH supply values for the six parameters SO, LM, SI, TA, NS, and LA. This is illustrated by the passage:

jokes (1) and (2) are both *dumb, figure-ground reversal, light-bulb changing, Polish riddles ...*

(Attardo and Raskin 1991:329)

and also:

From the point of view of the GTVH, each joke can be viewed as a six-tuple, specifying the instantiation of each KR as a parameter:

(5) Joke: LA, SI, NS, TA, SO, LM

(Ruch *et al.* 1993:126)

Some of these parameters do not seem to have any internal structure: they are atomic symbolic labels:

Each KR is a list or set of lists from which choices need to be made for use in the joke.

(Attardo and Raskin 1991:313)

the values for the LM and SO seem to be limited in number ... while the SI and LA are much more numerous.

(Ruch *et al.* 1993:126)

Hence the way that they contribute to the structure of a joke must be determined by how these symbols are used (i.e. how they are related to each other, or manipulated in some way). Symbolic parameters make sense only as input to some organized configuration or process whose behaviour they affect. *The accounts of the GTVH give no indication of what this central mechanism or framework is.* The KRs contribute the values for the six parameters, but it seems to be assumed that these parameters are enough to totally define the joke (they 'assign a unique descriptor to any joke' (Ruch *et al.* 1993:126)). This is not convincing, given that the illustrative parameter sets look simply like tuples of symbols such as *dumb*, *figure-ground reversal*, *light-bulb changing*, *Polish*, *riddle*, etc. Such tuples define a joke only as input to some joke-defining formula or procedure, since a joke is not itself a small tuple of atomic symbols.

To some extent this lack of content is related to the definition (or lack of it) of the exact contribution of the language (LA) KR:

The difference in the choice of words, syntactic constructions, and other language options, including the division of the text into sentences, will be referred to as the difference in the language. (p. 297)

the parameter of language finds itself in a unique relationship to all the others, namely, taking input from all of them. (p. 299)

(Attardo and Raskin 1991)

It is not clear whether two different jokes can ever have the same value for the LA parameter, if that parameter defines exactly the words and their syntactic structure; that is, it seems that two surface texts are identical if and only if their LA values are identical.

The LA KR produces a set of *choices* (Attardo and Raskin 1991:298, 321), and the other KRs affect the LA KR by constraining or preselecting these choices. The statement that the LA *parameter* will 'take input from' the others should perhaps say that the LA *knowledge resource* takes input from the others. However, the way in which other parameter choices affect the LA KR is not outlined. It also seems odd that the central mechanism for imposing meaning on the various parameter values should be the LA KR, which is supposed to be contributing a parameter itself, and is not presented as the master KR to which others are subordinate (unless that is how the remark about its 'unique relationship' is to be interpreted). This would mean that the real substance of the GTVH was embodied in the LA KR, which does not seem to be what Attardo and Raskin intend.

6.4.3

Relationship to theories of language

As discussed in [Chapter 3](#), it is desirable to be explicit about how any proposed theory of verbally expressed humour relates to theories of, or assumptions about, language. The position of GTVH is not as clear as it could be, in two respects.

Firstly, some of the knowledge resources listed in the GTVH seem to describe general textual characteristics rather than humorous ones. Attardo (2001:24) acknowledges that the SI parameter is indeed applicable to general texts, not just humorous ones, but it also seems plausible that any text could be classified according to its language (LA) and narrative strategy (NS). A non-humorous text might even have an intended target (TA).

For example, the question–answer pair (40) has the same syntactic structure as many simple riddles.

(40) What do you get when you cross a horse with a donkey?

A mule.

(constructed item)

It is not clear how the LA value for this example differs from that for the joke (3), but presumably they are very similar, and the NS values are most likely identical.

(3) What do you get when you cross a sheep with a kangaroo?

A woolly jumper.

The GTVH presents LA, NS, TA, and – usually – SI as if they were humour-related, thereby confusing the theoretical position. Blurring this distinction also has consequences for empirical testing of the GTVH (see [Section 7.2.2](#)).

The second in clarity concerns the adoption of a linguistic framework. As discussed in [Section 3.1](#), any theory of VEH has to make some assumptions about language, and so must have some vocabulary of linguistic categories and relationships with which to describe the data in a suitably systematic and general fashion. Although a theory of humour is not a theory of language, it must connect in a systematic way to the relevant linguistic factors.

The GTVH approach is generally to use specific linguistic tools, often devised by the SSTH/GTVH originators. Originally, the linguistic theory assumed was Raskin's semantic script theory (SST) (Raskin 1985). More recently, Attardo (2001) has set out various elaborations of this basic mechanism. While explicit postulation of the linguistic basis is very sound, it does mean that it is hard to express the SSTH/GTVH ideas in any other linguistic terms. As we have argued earlier, it can be helpful to separate

claims about humour factors from a particular linguistic theory. The SSTH offers a package deal, in which one has to accept not only the claims about humorous texts, but also the script-based approach to semantics as a general linguistic framework.

6.4.4

Falsifiability and empirical adequacy

One highly desirable property for a scientific theory of some naturally occurring phenomena is that it should make some clear predictions about what data should be present, and that theory should (via its predictions) be falsifiable. The authors claim that the SSTH is falsifiable (Raskin 1985:132) and of the GTVH say:

GTVH is falsifiable ... It can be falsified on a purely theoretical ground by discrediting one of its KRs and replacing it with another as well as by simply adding a new KR to it. Such a falsification would amount to a theoretical revision of GTVH.

(Attardo and Raskin 1991:328)

However, the lack of any criteria for what would count as a valid KR means that one could posit further KRs, as sketchily defined as the existing six, and there would be no way to arbitrate between the proposals.

The authors go on to say that empirical falsification of the GTVH could come from its predictions of joke similarity; we will consider this in more detail in [Chapter 7](#).

Doubts about falsifiability are also raised by the way in which notions such as scripts and script opposition appear to be arbitrarily flexible in order to accommodate putative counter-examples. For example, in footnote 1 in Attardo and Raskin (1991), Morreall is cited as claiming that various classes of humour (e.g. (15) in [Section 3.3](#)) are not covered by the SSTH. The authors respond that the theory does indeed handle these examples, perhaps with some minor extensions, but they do not show that any of these joke-classes conform to the arrangement outlined in the SSTH (which was presumably Morreall's point).

Raskin and Attardo do not provide systematic examples of texts where some but not all of the SSTH ingredients are present, resulting in a lack of humour, alongside very similar, but humorous, examples in which the missing ingredients are present. It would be interesting for SSTH proponents to explain how certain script pairs are opposed/unopposed, and then show that texts which brought script pairs together were funny if and only if the scripts were opposed. Such an exercise would be particularly interesting in view of the 2002 definition of script opposition (see [Section 6.4.1](#)), since it may be difficult to distinguish humorous texts from mere misunderstandings, given such a weak definition of opposition.

We would argue that it is very difficult to work out any objective predictions from the GTVH as it stands. Most of the actual humour-related content of the GTVH is from the original SSTH, which also has little predictive power. It is very difficult to imagine what could, in principle, count as evidence against the SSTH.

6.5 Summary

The GTVH consists of a classification system for jokes, built upon the SSTH description of the conditions which make a text humorous. This critique has argued the following points:

- (a) The GTVH parameters are not sufficient to define a joke without some central framework or process to which they act as parameters. Not only do the GTVH articles give no clue as to the workings of this essential part of the theory, its existence is not even hinted at.
- (b) At least four of the six GTVH classifications appear to be equally applicable to non-humorous texts, and so it is not clear why they form a central part of a theory of humour.
- (c) The relationship of the GTVH to models of language structure or use is unclear.
- (d) Most crucially, the basic concepts, particularly the central mechanisms of scripts, script opposition, and logical mechanism, are so ill-defined (or defined in different ways) that the theory has very little substance. In particular, it makes no predictions beyond the informal ones already present in other incongruity-resolution approaches.

Any theory will, at its early stages, have some underdeveloped aspects, as it is impossible to produce a complete and detailed theory all at once, let alone to summarize it in a single article. However, some of the weaknesses of the GTVH discussed here are fairly fundamental. As noted earlier, the GTVH is coming to be regarded as an accepted theory, but that makes it all the more important to subject the details to rigorous scrutiny. We will not be able to move forward unless we accept that what we have at present is not perfect. The fact that Attardo (1997, 2001, 2002) has sought to improve the GTVH is an acknowledgement of that methodological principle, but the critique offered here is more radical.

7

Joke similarity and identity

We consider how similarities between jokes might be captured formally.

7.1

Centrality of information

It seems intuitively clear that, within the content of a joke, some parts are more essential than others. For example, consider the trivial variation (42) of example (41).

(41) Why did the elephant wear red socks?

Because his green ones were being washed.

(Dienhart 1999)

(42) Why did the elephant wear blue socks?

Because his yellow ones were being washed.

(constructed item)

In these two jokes, the precise choice of colours for the socks is not relevant (although it may be arguable that other words, such as tartan or puce, might have enhanced the joke – see [Section 12.3](#)). However, it is hard to ask about the colours of socks in this way without mentioning actual colours, so some choice was necessary. It is tempting to say that (41) and (42) are ‘the same joke’, so similar are they. However, it is far from clear that jokes can be organized into tidy equivalence classes so that any two jokes are either ‘the same’ or ‘different’. For example, is (43) the same joke as (41), or a different one?

(43) Why is your dog wearing brown boots?

Because his black ones are being mended.

Although (41) and (43) could be seen as differing only in substitution of similar vocabulary items (e.g. *elephant* ¡ú *dog*), if we abstract slightly further, there is some similarity to (44).

- (44) Why do firemen wear red braces?
To keep their trousers up.

Moving further in the direction of abstraction, (44) can be seen as similar to (8), emphasizing that all these examples use the device of ambiguity of question-focus (see [Section 8.1.1](#)).

- (8) Why do birds fly south in winter?
It's too far to walk.

Also, some of the similarity that we intuitively sense in these joke texts may be, at least in part, due to ordinary linguistic similarity. That is, (44) and (8) are textually similar to the non-jokes (45) and (46) respectively.

- (45) Why do firemen wear strong helmets?
To protect their heads.

(constructed item)

- (46) Why do birds fly south in winter?
It's too cold in the north.

(constructed item)

O'Mara and Waller (2003) describe a controlled test for assessing whether someone has grasped a simple joke: the participant is presented with the joke and a set of words or phrases which could be substituted for a particular part of the joke. The participant is asked to choose a substitution which will leave the joke working in the same way as before. O'Mara and Waller point out that the joke-preserving substitution often involves a semantically similar word (e.g. *frighten* for *scare*), but sometimes a quite different phrase can 'retain the underlying structure of the joke' (p. 80).

- (47) Why did the boy take a hammer to the school?
Because it was the day they broke up.

(O'Mara and Waller 2003:80)

For example, in (47), the joke-preserving substitution for *a hammer* might be *glue*. This suggests that O'Mara and Waller regard these relatively different textual forms as, informally, corresponding to a single underlying abstract core.

There is the related issue that changing certain parts of a joke may transform it into a non-joke. It might seem that any part which can lead to such a change is more 'essential' than the more peripheral information such as the colour-names used in (41) or (42). However, some changes may alter the joke not into a non-joke, but into another similar joke (as

in (41) and (43)), or, if radical enough, into a dissimilar joke. That is, there may be content or structure in a joke which is crucial to that text being a joke, but it is not obvious which factors define a unique ‘identity’ for that joke. The actual text is unique (subject to the various idealizations we discussed in [Chapter 3](#)). Changes to the text will, by definition, create a different text, and may or may not preserve the joke status, but there may not be a definable sense in which changes can preserve (or fail to preserve) the text’s identity as ‘the same’ joke.

It might seem that we could define the essential information in a joke as those properties of the joke which instantiate whichever joke-class definition licenses the text as a joke: if text T falls into joke-class C by virtue of having properties $P_1, \dots, P_n$, then $P_1, \dots, P_n$ constitute the crucial aspects of T as a joke. However, that leaves open the level of abstraction involved in choosing our joke-class definitions. For example, the forced reinterpretation joke-class (Chapters [5](#) and [8](#)) makes only very general requirements on the form of the joke, and so it defines a broad and varied class of jokes. If all FR jokes were deemed ‘the same’, the equivalence class would be very large and would contain some intuitively very different jokes. Within that large set, subclasses could be defined (for example, according to types of ambiguity device used). However, there is no obvious and principled way of choosing, from within the hierarchy of subclasses, a level which is particularly relevant to stating the joke’s essence.

7.2

Formal measures

A formal theoretical model of jokes could offer a basis for a measure of similarity or a definition of identity for jokes. There are two ways in which this could happen – by defining some abstract ‘core’ for each joke (see [Section 7.2.1](#)) or by having a set of parameters, perhaps varying in importance (see [Section 7.2.2](#)). Such notions of ‘sameness’ or ‘similarity’ would then be specific to the form and content of the particular model, and would not be pre-theoretic concepts which could be applied to jokes in isolation from the analysis assigned by the theory to the jokes in question. That is, a definition based on a theoretical analysis might well not supply simple operational criteria such as

two jokes count as the same joke if one can be transformed into the other by a substitution which meets the following conditions ...’. Instead, one would define similarity/identity in terms of the properties of the analyses that the theory assigned to the jokes.

7.2.1

Core and variations

The JAPE computer program (see [Section 10.5](#)) can generate slightly different texts from a single schema, owing to variations in the language

generator component. Two jokes which are based on the same schema could be defined to be *schema-equivalent*. Although this would capture the informal notion of ‘having the same underlying idea’, it is not clear that it is a particularly interesting definition, or that it accurately models ordinary judgements of ‘sameness’.

Similarly, Attardo and Raskin’s GTVH ([Chapter 6](#)) would allow definitions of *SO-equivalent*, *LM-equivalent*, etc., which would offer five different equivalence measures (or six, if we include *LA-equivalent*, in which each equivalence class would – presumably – contain exactly one text). If it were accepted that, of the GTVH parameters, only LM and SO are humour factors (as is argued in [Chapter 6](#)), then we could define a notion of ‘using the same SO and LM’, which might capture the notion of ‘same abstract joke’.

Hofstadter suggests the notion of an *ur-joke*, which is a ‘skeleton’ from which many jokes can be built, all with some deep similarity:

many very different jokes can share the same ur-joke, differing merely in how they
clothe’ it – that is, what setting they place it in, how the joke is phrased, and so on ...
many jokes are built on a combination of ur-jokes ...

(Hofstadter and Gabor 1989:431)

Hofstadter offers a number of illustrative examples, all of which suggest that he sees these ‘skeletons’ as operating at an underlying level, where the twisted logic of the joke operates, rather than at any specifically linguistic level. For example, (48), (49) (as well as two cartoons) are said to use the *role-reversal* ur-joke.

(48) What did Mickey Mouse get for Christmas?
A Dan Quayle watch.

(Hofstadter and Gabor 1989)

(49) Notice at a swimming pool: we don’t swim in your toilet – please don’t pee in our pool.

(Hofstadter and Gabor 1989)

Only one actual ur-joke is sketched in English. The variants (50) and (51) are suggested as being based on the same ur-joke.

(50) A man in his fifties goes to the doctor and says, “Doctor, I’ve got a problem. You see, when I was younger I always used to get erections that I couldn’t bend with my hand. Now though, I can bend every erection I get. What I want to know is, am I getting stronger or weaker?”

(Hofstadter and Gabor 1989)

(51) God goes to the doctor and says, “Doc, I’ve got a problem. You see, I used to be able to make stones that were so heavy I couldn’t lift them. But now I *can*’t make a stone that I can’t lift. The question is, am I getting more or less omnipotent?”

(Hofstadter and Gabor 1989)

The ur-joke is roughly summarized as:

An individual interprets two tests as measuring the same desirable quality (strength, intelligence, etc.). Not only are different answers obtained on the two tests, but in fact, a high score in one test *necessitates* a low score on the other; they are merely flip sides of each other! Then at a later time, the outcomes of the same two tests are reversed, but of course this is of no more diagnostic value than the first time. So the individual remains baffled.

(Hofstadter and Gabor 1989:435)

Hofstadter and Gabor’s discussion group remarked upon the apparent lack of humour of the ur-joke itself, when presented on its own. They speculate that the act of uncovering the ur-joke (in the process of interpreting the joke) ‘is sufficiently interesting to yield some humor’.

The ur-joke idea is not developed in great detail in the cited article, and does not appear to have been developed further (although it may have influenced the GTVH – see Attardo and Raskin (1991:303), Attardo *et al.* (2002:5)).

7.2.2

Ordered dimensions

The GTVH (see [Chapter 6](#)) is presented by Attardo and Raskin (1991) as primarily describing joke similarity: that factor is the main source of evidence for their hierarchical ordering of knowledge resources (although the reasoning they give, from example jokes to resource orderings, is not watertight). Therefore, corroboration or refutation of the GTVH could come from empirical evidence about joke similarity. This is the motivation for the experiment described in Ruch *et al.* (1993), which explored two hypotheses:

First, the subjects will perceive some jokes as more similar and other jokes as less similar to one another; second, if the GTVH is correct, subjects will perceive a linear increase of similarity between pairs of jokes selected along the KR hierarchy.

(Ruch *et al.* 1993:127)

(Presumably the word ‘consistently’ should be inserted before ‘perceive’.) The latter hypothesis was explained more fully:

Thus, the GTVH predicts that given a joke A and a joke B differing, say, in a ‘low’ parameter like LA and a pair of jokes C/D differing in a ‘high’ parameter like SO, the first pair will be perceived as more similar and the second as less similar.

(Ruch *et al.* 1993:128)

The experimenters created three groups of jokes (blonde, chicken, light bulb). Each group consisted of a base joke (in question–answer form) and six variants, where each variant differed from the base joke in *one* of the GTVH parameters (LA, NS, TA, SI, LM, SO). For example, the chicken items were (52), (53), (54), (55), (56), (57), and (58).

(52) Why did the chicken cross the road?

It wanted to get to the other side.

(anchor item)

(53) Do you know the reason why the chicken decided to cross the road?

Because it wanted to get to the other side.

(different LA)

(54) The reason the chicken crossed the road is that it wanted to get to the other side.

(different NS)

(55) Why did the turtle cross the road?

It wanted to get to the other side.

(different TA)

(56) Why did the chicken eat an octagonal-headed worm?

Because it was hungry.

(different SI)

(57) Why did the chicken cross the road? Nothing ventured, nothing gained.

(different LM)

(58) Why did the chicken cross the road? He saw a blonde hen on the other side.

(different SO)

In view of the nature of the LA parameter – see [Section 6.4.2](#) – presumably *all* the non-base jokes differ from the base joke in their LA value, in

addition to the other, named, parameter. Also, (59) is stated as differing from (36) (in [Section 6.3](#)) *only* in its script opposition, but it must surely also differ in its situation, given the addition of such items as shoes and air deodorants.

(59) How many Poles does it take to screw in a light bulb? Five. One to take his shoes off, get on the table, and screw in the light bulb and four to wave the air deodorants to kill his foot odor.

(Ruch *et al.* 1993)

The subjects (over 500 undergraduate psychology students) were presented with pairs of jokes, where each pair consisted of a base joke and one of the variants from its group, and were asked to judge, for each pair, how similar the two jokes were, on a scale of ‘undecided’ plus 4 degrees of similarity. The results broadly show higher similarity where the parameter-difference is low in the GTVH hierarchy (and vice versa), as predicted (Ruch *et al.* 1993: Figure 3).

Examination of the graphs reveals some internal patterns, both across the three joke-groups and across the GTVH parameters. The clearest (inverse) proportionality of similarity is present for the lowest four parameters – LA, NS, TA, and SI – although one group (chicken jokes) place LA and TA very close, with TA significantly *higher* in similarity ratings (i.e. lower in the hierarchy) than NS. The average across all three groups shows a steady drop in similarity across these four parameters, as predicted. However, these parameters are exactly the parameters which we have argued ([Section 6.4.3](#)) to be general textual characteristics rather than humour-specific properties. The change from base joke to the LA item consists of a minor paraphrasing of the text, the change for the NS item is a paraphrase altering it from a question–answer pair to a declarative sentence, and the difference for the TA pair is the direct substitution of a noun phrase (e.g. *Irishmen* for *Poles*). The change of SI involves various alterations to the meaning. Hence we could account for the trend involving LA, NS, TA, and SI by these non-humour-related (and fairly plausible) hypotheses:

- A minor paraphrase retaining question–answer structure will result in less perceived difference than a minor paraphrase which changes the text into a single declarative sentence. (Hence, LA < NS.)
- A change which results in a near-exact paraphrase will make less difference than an alteration to the meaning. (Hence LA < TA, NS < TA, LA < SI, NS < SI.)
- A change which replaces exactly one noun phrase with another which is very similar semantically will make less difference than a set of changes which introduce several semantically different constituents. (Hence TA < SI.)

Although the experiment did not use any non-jokes in the data, it would be interesting to test whether comparable similarity judgements are made by subjects when presented with pairs of non-joke texts manipulated in similar ways.

For the more humour-related parameters, LM and SO, the position is less clear-cut. The three highest parameters, SI, LM, SO, all correspond to lower joke similarity than the items for the parameters LA, NS, TA, as predicted. For the parameters SI, LM, and SO, there are three predicted orderings of similarity scores ($SI > LM$, $SI > SO$, $LM > SO$), so for the three data groups (blonde, chicken, lightbulb) there are nine orderings to be verified. Of these, four are as predicted, five are not: the $SI > LM$ ordering fails for all three groups, and all three orderings fail for the chicken group. For the average across the groups, $LM > SO$ is corroborated, but SI and SO are so close that the difference is not statistically significant. This does not provide convincing evidence that the more humour-related parameters (LM, SO) are clearly ordered with respect to each other or with respect to the 'nearest' non-humour parameter, SI.

The GTVH-based similarity measure is based on the idea that jokes differ along several different dimensions, that some of these dimensions are more dominant or significant than others, and that the perceived similarity of jokes will reflect the ordering of these differences. That general formal arrangement makes sense, but (unless we subscribe to the GTVH) it is as yet unclear what the relevant dimensions of comparison are, or how they might be ordered. For example, one way in which jokes have traditionally been classified (e.g. Pepicello and Green (1984)) is by the type of ambiguity that they involve. It is not obvious how important this factor would be in rating the similarity of joke texts.

7.3

Extraneous information

For the reasons given above, it is not straightforward to try to define the 'essence' or 'identity' of a particular joke. However, some jokes do seem to contain information which is peripheral to the joke. We will adopt the term *extraneous* to refer to information which is not relevant to the logical or presentational structure of the joke, not even in the subsidiary manner of the colour-names in (41) and (42) (in [Section 7.1](#)). It might seem odd to have such a category, since it could be assumed that jokes do not contain redundant material (cf. the variant of the Gricean maxims in Raskin (1985:103) and Raskin and Attardo (1994:37)). However, two points have to be made here. First, our analysis framework is intended to be applicable to actual jokes, not carefully optimized jokes. Jokes which occur naturally do contain unnecessary material. Second, our category of extraneous material covers content which is not involved in the *logical or presentational structure* of the joke. As will be discussed in [Section 12.3](#),

there are other aspects of a joke, and some content might be both extraneous (logically) and yet play a supplementary role in enhancing the joke. For example, consider (60):

(60) A man walks into the front door of a bar. He is obviously drunk and staggers up to the bar, seats himself on a stool and, with a belch, asks the bartender for a drink. The bartender politely informs the man that it appears he has already had plenty to drink and that he could not be served additional liquor. The bartender offers to call a cab for him. The drunk is briefly surprised, then softly scoffs, grumbles, climbs down from the bar stool and staggers out the front door. A few minutes later, the same drunk stumbles in the side door of the bar. He wobbles up to the bar and hollers for a drink. The bartender comes over and – still politely if not more firmly – refuses service to the man and again offers to call a cab. The drunk looks at the bartender for a moment angrily, curses, and shows himself out the side door, all the while grumbling and shaking his head. A few minutes later, the same drunk bursts in through the back door of the bar. He plops himself up on a bar stool, gathers his wits, and belligerently orders a drink. The bartender comes over and emphatically reminds the man that he is drunk and will be served no drinks. He then tells him that he can either call a cab or the police immediately. The surprised drunk looks at the bartender and in hopeless anguish cries, ‘Man! How many bars do you work at?’

(from www.the-jokes.com)

This joke has a considerable amount of extraneous material, such as *with a belch, softly, curses, all the while grumbling and shaking his head*. We will return to this in [Section 12.3](#).

7.4 Summary

Despite the pre-theoretic intuition that some jokes are ‘the same’ as each other, a definition of joke identity is dependent upon having some formal definition of joke classes or structures, and no current proposals in this area offer a notion of ‘identity’. Joke similarity is perhaps more tractable, but this is dependent on having a formal theory of joke structure in which either there is some notion of a central core, or there are various dimensions of variation, perhaps ordered with respect to each other (although the latter is not essential). Moreover, the more general issue of non-humorous textual similarity is a complicating factor.

8

Manipulating interpretations

We examine the ways in which information is presented in certain forms of narrative joke, particularly those involving multiple interpretations.

8.1

Forced reinterpretation – more details

8.1.1

Ambiguity and interpretation

In our discussion in [Chapter 5](#) of the forced reinterpretation (FR) model of humour, we implied that an ‘interpretation’ is the same as the notion of a ‘meaning’ within linguistics, and that having more than one interpretation can be equated with being ambiguous (lexically, syntactically, etc). This limited view is adequate for certain jokes, such as (8) and (41).

(8) Why do birds fly south in winter?

It’s too far to walk.

(41) Why did the elephant wear red socks?

Because his green ones were being washed.

Both (8) and (41) work by letting the audience assume that the question has a particular focus (*fly south in winter* or *wear red socks*) and then providing a punchline which assumes a different focus (*fly* or *red*). Linguistic ambiguity (lexical-syntactic, in the sense of [Section 3.7](#)) is also directly involved in (25) and (26).

(25) A lady went into a clothing store and asked ‘May I try on that dress in the window?’

Well,’ replied the sales clerk doubtfully, ‘don’t you think it would be better to use the dressing room?’

(26) Postmaster: Here’s your five-cent stamp.

Shopper (with arms full of bundles): Do I have to stick it on myself?

Postmaster: Nope. On the envelope.

We can even include examples such as (61) in an FR analysis, if we allow an interpretation to include the pragmatic status of an portion of the text.

(61) What's the difference between an elephant and a watermelon?

(I don't know.)

You'd be a fine one to send to the store for a watermelon.

(Dienhart 1999)

This does not stretch the notion of 'linguistic meaning' very far, as (61) could be regarded as a case of semantic-pragmatic ambiguity.

Intuitively, the FR account also describes (18), in that the punchline triggers a reinterpretation of the set-up.

(18) 'Is the doctor at home?' the patient asked in his bronchial whisper.

No,' the doctor's young and pretty wife whispered in reply. 'Come right in.'

However, (18) does not involve linguistic ambiguity or two 'meanings' in the normal sense of the term; rather, it involves two ways that one of the characters in the joke (and the audience) could make sense of the other character's actions. It is just about possible to relabel this as semantic-pragmatic ambiguity, as in the case of (24), by claiming that there is ambiguity about the *perlocutionary effect* (Levinson 1983) of the query *Is the doctor at home?*, but this is a rather contrived argument, since it would class any action which could have more than one outcome as 'ambiguous'. The difference in the two 'meanings' of the query occurs at some higher level of the character's plan, and is not directly part of (say) the illocutionary force of the utterance (Levinson 1983). The non-linguistic aspect of the misinterpretation is even clearer in (62).

(62) Russian officers in an Eastern European country go to a tavern. They order beer. The waiter places coasters on the table and serves the beer. Later they order another round. The waiter returning with the beer finds no coasters. 'OK,' he tells himself, 'these are collectors,' and puts down another set of coasters. When the third round is ordered and brought out, there are again no coasters. Angry, the waiter puts the beer down on the table, but places no more coasters. One of the Russian officers protests: 'What's this? No more crackers?'

(Hetzron 1991:62)

This leads to a choice: we can still use the precise linguistic notion of ambiguity, and exclude examples like (18) and (62) from the generalizations of the FR account, or we can widen our notion of interpretation and capture a wider range of data. Here, we opt for the latter. One reason

for doing so is that it gives a better understanding of examples like (25) and (26). All of (18), (25), (26), and (62) can be seen as ‘funny stories’, in that they have a narrative structure absent from an example such as (8). Although (25) and (26) contain linguistic ambiguity, this ambiguity gives rise to an *event* (the act of uttering the ambiguous text) which can then be interpreted in more than one way. That is, all three examples involve alternative interpretations of *sequences of events*; the fact that in two cases the possibility for multiple interpretation arises from an ambiguous utterance is a relatively small detail (see [Section 3.8](#)).

It may be that this is the only way that ambiguity will contribute to a funny story.

Conjecture: in a narrative joke (*funny story*), if linguistic ambiguity is crucial to the joke, then it will occur not in the supporting narrative, but only in an utterance by a character in the story (reported directly or indirectly) or within some written material quoted within the story (e.g. a letter, a printed sign).

The ‘interpretation’ of any text, including a joke, usually extends beyond the literal meaning of the words and phrases involved. In relatively simple jokes, such as (6) and (25), the incongruity (ABSURD and/or TABOO content) is not inherent in the *literal* content of any of the sentences.

(6) Do you believe in clubs for young people?

Only when kindness fails.

In (25), the punchline *don't you think it would be better to use the dressing room?* implies (in some loose sense), but does not *state*, the idea that the lady had been suggesting using the window as a changing room. This is the overall interpretation when punchline and (the less OBVIOUS meaning of) the setup are taken together.

Sometimes the connection between the conveyed information and the propositions embodying the INAPPROPRIATENESS is very indirect. This can be illustrated by some examples in which the link is increasingly tenuous.

In (22), the proposition ‘a particular person is being knocked over every two minutes by cars’ could be said to be ABSURD, and is the entire alternative reading of the set-up. That is, this is a very direct connection, made oblique only by depending on the less OBVIOUS meaning of the text. Example (18) has a text whose ‘alternative’ (less OBVIOUS) interpretation implies ‘the wife-character is encouraging a surreptitious visit in her husband’s absence’, which then invites the plausible inference (not a strong logical consequence) that ‘the wife-character is inviting adulterous activities’. In example (9), the text conveys the basic information *The t is silent*

as in *Harlow*, which in turn implies the proposition expressible as ‘Jean’s name is spelt H-A-R-L-O-T’. There is then no logical implication, only some weaker relationship (‘hinting’?) that ‘Jean has something in common with a prostitute’, an idea with clear sexual associations. (8) is even more indirect. The potentially ABSURD image (flocks of birds trudging across continents) is not implied by *either* interpretation of the text: the ‘hidden’ meaning *denies* that this event happens. It could be argued that the statement that a particular event (the birds walking) is infeasible implicitly admits, perhaps even presupposes, that the event is conceivable, and moreover brings the concept of the putative event into the context of discussion.

There are examples where it is hard to attribute the humour to anything other than an indirectly summoned image, such as (63).

(63) A blonde woman gives birth to a child with bright red hair. ‘Oh,’ says the doctor, ‘Father a redhead, is he?’ ‘No idea,’ says the woman, ‘he didn’t take his hat off.’

(cited by Morreall (1996) from Palmer (1994))

Although Morreall says that the ‘humor here is based on the great degree of casualness with which the woman and her sexual partner conceived the child’, a text which simply said that two people conceived a child through very casual sex would not necessarily be funny, despite this being a suitably TABOO notion. Nor is there even a notion of violated expectations, since casual sex and conception are not uncommon. We can experiment with this text by replacing the final phrase *he didn’t take his hat off* with other sequences, some of which have essentially the same import with respect to the casualness of the conception event. Intuitions may vary, but lines such as the following seem to reduce the impact: ... *I’m not sure who the father is*, ... *we only met once*, ... *it was just a quick fling*, ... *we did it in the dark*.

It is even arguable that some of these variants turn the text into a non-joke. It is plausible to speculate that some of the humour comes from the notion of a man continuing to wear his hat during sexual intercourse.

It is uncontroversial to suggest that when a person hears (or reads) a meaningful text, some form of ‘reasoning’ takes place. This may not be rigorous logical deduction as defined in a traditional mathematical logic. Instead, it consists of some combination of processes such as plausible commonsense’ reasoning, backward reasoning of the sort sometimes known as *abduction* (Charniak and McDermott 1985: Ch. 8), some of the phenomena attributed to the manipulation of mental spaces (Fauconnier 1985, 1996), the invoking of familiar patterns of experience (Schank and Abelson 1977), or default assumptions which may be overturned later (Reiter 1978, Lascarides and Asher 1993). All this goes to ‘flesh out’ the bare content of the sentences in the text, resulting in a richer meaning.

It is unclear whether the extended scenarios listed above (in connection with examples such as (8), (9), and (62)) are properly part of the ‘interpretation’ of the text, or are better viewed as possible inferences from that interpretation. In the absence of clear evidence, we have taken the position in our formalization ([Appendix A.1](#)) that there is no principled notion of ‘interpretation’ in between the bare literal meaning and the entire set of inferences (or hints, or abductions, etc.) which may flow from that meaning.

Another reason for considering indirectly implied propositions (including presuppositions) in analysing a joke is that it facilitates connections between descriptions of joke mechanisms and broader theories of joke use or effect, including proposals for superiority or hostility theories of humour. Jokes which convey insults (or other contentious material) usually do so obliquely rather than in the literal meaning of the text.

We tentatively suggested (Ritchie 1999) that it may also be necessary to use a definition of ‘interpretation’ which is, in a sense, stricter than linguistic ambiguity. In particular, *the way in which a meaning is expressed* could be relevant to creating the joke, even though *the overall meaning* remains the same in the two possible expressions. That is, a text which unexpectedly switched between two different ways of expressing the same overall meaning could act as a joke. This conjecture is hard to substantiate.

(64) One prostitute said to another, ‘Can you lend me ten dollars until I get back on my back?’
(Suls 1972)

Although Ritchie (1999) suggests that (64) is an example of this, that joke is explicable more straightforwardly as a form of pun ([Chapter 9](#)) augmented with some notion of *thwarted expectation* ([Section 12.3](#)).

8.1.2

Relating set-up and punchline

There seems to be no evidence to suggest that the set-up in an FR joke is an abnormal text, grammatically, semantically, or in discourse terms. Hence we do not need to postulate any unusual linguistic mechanisms in order to characterize the interpretation of the set-up. The need for there to be a ‘most OBVIOUS’ reading of any ambiguous forms is also fairly conventional (if not fully understood theoretically).

It is less clear what defines the punchline, as there will be no explicit marker (except in special cases such as question–answer jokes). According to the FR model, the punchline can be recognized by the fact that it does not make sense (CONFLICTS) with the current (most OBVIOUS) interpretation of the preceding text. Nevertheless, it must be possible for at least the literal meaning of the punchline text to be interpreted semantically

(despite its lack of fit with the textual context), or this CONFLICT would not be detectable. Inspection of examples such as (6), (7), (8), (24), (25), (26), (28), or (41) (remember that all our examples are listed in [Appendix B](#)) suggests that the difficulty involves a lack of *discourse coherence*: a response does not seem to answer the question, or – in (62) – the utterance does not make complete sense with respect to the hitherto OBVIOUS interpretation. That is, the CONFLICT relation should *not* be identified with logical inconsistency, nor with the kind of semantic clash that results from mismatches such as applying a verb to a semantically inappropriate kind of object (violation of selectional restrictions, in the terms of Chomsky (1965)). In all these examples, the punchline *assumes* a particular, hitherto hidden, perspective, giving an effect very much like the failure of a *presupposition*. The audience, in these cases, has to make sense of the punchline by ‘working backwards’, a semantic (or discourse) phenomenon reminiscent of the use of abductive reasoning (Hobbs *et al.* 1993).

Once the punchline has been detected, there is then the question of how the meaning of the punchline relates to the possible interpretations of the set-up. (As pointed out in [Section 5.2.2](#), COMPATIBILITY is not necessarily the negation of CONFLICT.) The usual FR account is to say that the punchline evokes a less OBVIOUS interpretation of the set-up. Since we are not considering the finer details of processing, we can gloss over the question of how these interpretations are found, and how any search amongst them might happen: our aim here is to define the abstract relationships which may hold, not how they are computed.

In the examples above where (we suggest) the CONFLICT relation which marks the punchline is discourse incoherence ((6), (7), (8), (24), (25), (26), (28), or (41)), it also seems to be the case that the punchline is *discourse coherent with a less obvious interpretation of the set-up*. That is, COMPATIBILITY, for these examples, is some form of discourse coherence. For example, in (25), if the lady-character had intended to use the shop-window as a changing room, then the clerk-character’s response would be completely coherent; in (8), if the questioner intended to focus on *fly* as the means of movement, then the answer does not seem incoherent.

Notice that many of the above jokes in which the punchline is signalled by discourse incoherence also involve linguistic ambiguity (although, as noted earlier, (18) does not). It may be that linguistic ambiguity ensures that the joke-audience’s commitment to the more OBVIOUS interpretation is at a very basic level of discourse comprehension, so that anything incompatible with this reading causes incoherence at a similar level.

Conjecture: in a forced reinterpretation joke, if the differing setup interpretations are caused by linguistic ambiguity, then the punchline will cause discourse incoherence with the more obvious interpretation.

To sum up, we have a subclass of FR jokes in which CONFLICT consists of lack of discourse coherence, and COMPATIBILITY is discourse coherence. A more formal statement of these *discourse-coherence misunderstanding* jokes is given in [Appendix A.2.1](#).

8.1.3

Another set-up/punchline relation

There are some jokes that seem broadly to be FR jokes, but where the punchline is not detectable by the kind of discourse incoherence outlined above. Examples (65) and (66) follow this pattern.

- (65) John and his wife Mary were having a shower together in their upstairs bathroom when the doorbell rang. Mary heard the bell, got out of the shower, wrapped a towel around her, went downstairs, and opened the door. Their neighbor Charlie looked at her from the doorway, and said, 'Oh. I see that I got you out of the shower. Sorry about that.' 'That's all right,' Mary said, 'What do you want?'

Not too much ... my goodness you have beautiful skin. It's so pink from the shower. Mary, if I was to give you a hundred dollars, would you remove the towel from your upper body?' Mary thought about it for a minute, figured why not, for a hundred bucks, and removed the towel from her breasts. 'Wow,' Charlie exclaimed, 'they are truly beautiful. Listen, for another hundred bucks would you consider taking the towel all the way off?'

Why not,' Mary thought, 'that's a lot of money,' and she dropped the towel completely to the floor. Charlie had a good look, complimented her again on her fine looking body, reached into his pocket, took out two hundred dollars, gave it to her, and left. As she got back upstairs and was getting back into the shower, John asked her who was at the door. 'Just Charlie,' she said, as she started to rub his back. 'Charlie, eh,' said John, 'Did he give you the two hundred dollars he owed me?'

(from www.jokes2000.com)

- (66) Peter decided to go skiing with his buddy, Bob. They loaded up Peter's station wagon and headed north. After driving for a few hours, they got caught in a terrible blizzard. They pulled in to a nearby farm house and asked the attractive lady of the house if they could spend the night. 'I'm recently widowed,' she explained, 'and I'm afraid the neighbors will talk if I let you stay in my house.' 'Not to worry,' Peter said, 'we'll be happy to sleep in the barn.' Nine months later, Peter got a letter from the widow's attorney. He then went up to visit his friend Bob and said, 'Bob, do you remember that good-looking widow at the farm we stayed at?' 'Yes, I do.' 'Did you happen to get up in the middle of

the night, go up to the house and have sex with her?’ ‘Yes, I have to admit that I did.’ ‘Did you happen to use my name instead of telling her your name?’ Bob’s face turned red and he said, ‘Yeah, I’m afraid I did.’ ‘Well, thanks! She just died and left me everything!’

(from www.jokes2000.com)

In these examples it is less clear what, if anything, signals that a punchline has arrived: it does not appear to be discourse incoherence. What seems to be happening is that the new information, although completely fluent from a discourse point of view, renders one of the (hitherto) less OBVIOUS interpretations more OBVIOUS, making it now the most OBVIOUS. That is, if we assume that interpreting a text results in a set of possible interpretations, ordered by their OBVIOUSNESS, then the content of the punchline, taken together with the set-up, yields a different ordering.

One way of comparing these two arrangements is in terms of the logical relation between the punchline meaning and the set-up interpretation. In the examples where discourse coherence seems central (e.g. (25)), the punchline requires the less OBVIOUS interpretation in order to make sense, perhaps a form of presupposition. In examples such as (65), the punchline makes sense on its own, but actually *implies* the less OBVIOUS meaning. That is, the characteristic of this subclass of FR jokes is as follows: some of the information supplied by the punchline would, if supplied earlier in the text, have made the less OBVIOUS interpretation be the more OBVIOUS one. See [Appendix A.2.1](#) for a formal definition of the structure of these *punchline revision jokes*.

8.2

Implied incongruity

We have already pointed out, in [Section 8.1.1](#), that the location of the incongruity in an FR joke may not be the literal meaning, but some inference from there. There are also non-FR jokes where inferring an incongruity seems to be the central joke-creating device.

(33) Sitting on the side of the highway waiting to catch speeding drivers, a State Police Officer sees a car puttering along at 22mph. He thinks to himself, ‘This driver is just as dangerous as a speeder!’ So he turns on his lights and pulls the driver over. Approaching the car, he notices that there are five old ladies – two in the front seat and three in the back. The driver, obviously confused, says to him, ‘Officer, I don’t understand, I was doing exactly the speed limit! What seems to be the problem?’ ‘Ma’am,’ the officer replies, ‘You weren’t speeding, but you should know that driving slower than the speed limit can also be a danger to other drivers.’ ‘Slower than the speed limit? No sir, I was

doing the speed limit exactly ... twenty-two miles an hour!’ the old woman says a bit proudly. The State Police Officer, trying to contain a chuckle, explains to her that ‘22’ is the route number, not the speed limit. A bit embarrassed, the woman grins and thanks the officer for pointing out her error. ‘Oh, thank you,’ she says, ‘It’s a good thing you didn’t see us on Route 119.’

Example (33) makes use of a common device in narrative jokes, misinterpretation by a character. However, it is not misinterpretation of information presented in the set-up, and revelation of this misinterpretation does not provide the punchline, so it is not an FR joke. The punchline supplies further information which is not in itself humorous or incongruous, but which permits the inference of an amusing consequence of the already established misinterpretation. There is a sense in which this inference, or something very similar to it, could be made, once the misconception has been stated: driving at a speed numerically equal to the route number is bound to result in some very high speeds (particularly in countries such as the US where speeds are stated in miles per hour). However, the punchline both introduces a concrete instance of this possibility, and draws attention to that consequence. The fact that the driver-character’s behaviour is based on a misconception is not central to the working of the joke. It is the general rule adopted by this character, together with the punchline statement, which allows the inference.

Thus we have a further, non-FR, presentation of information, in which the punchline stimulates or draws attention to a particular inference from the set-up information. A formal definition of the presentational structure of these *punchline inference* jokes is given in [Appendix A.2.1](#).

8.3 Double entendres

There is a common class of joke in which a final passage of the joke (the punchline, essentially) allows more than one interpretation. If the joke is narrative, this multiply interpretable text is usually in the mouth of one of the characters. A typical example is (67).

(67) An English bishop received the following note from the vicar of a village in his diocese:

Milord, I regret to inform you of my wife’s death. Can you possibly send me a substitute for the weekend?’

(attributed by Raskin (1985) to Pocheptzov)

This is not an FR joke, as there is no punchline which imposes a hitherto unobvious reading on a prior portion of the story. The humorous effect results from a story character uttering the phrase with one intended interpretation, while the audience perceives another.

Similar non-narrative versions are possible, such as (68).

(68) It so happens that if there is any institution which is not susceptible to any improvement whatsoever, it is the House of Peers.

(W. S. Gilbert cited by Kelly (1971:6))

(Oddly, Kelly seems to be offering this as an example of a pun.)

In (67), the punchline does not serve to reveal hidden ambiguity in the preceding text, but it could be argued that such jokes work better if the audience momentarily perceives the reading which is *not* INAPPROPRIATE, and only later notices the other interpretation. It is interesting to consider why the audience should seek a further meaning, as (in contrast to FR jokes) there is no ill-fitting punchline (no CONFLICT) to trigger a search for a less OBVIOUS meaning. It may be that the joke-telling ritual, coupled with the fact that this is the end of the text, prompts the hearer to seek a hidden meaning (cf. [Section 5.3](#), in which the algorithm checks whether the text has ended). Attardo (1994: Ch. 3) proposes that there is reinterpretation, although he does not view this as a separate class of joke. His position (following the SETH – see [Chapter 6](#)) is that *all* jokes have this reinterpetative process underlying them, even where there is no closing text to stimulate the reanalysis, and he analyses (29) in these terms (see [Section 9.2](#)).

Some jokes consist of protracted double entendres, by having a narrative which sets up a radical misunderstanding by a character who then utters a sequence of remarks open to more than one interpretation, with one interpretation usually being TABOO. There is a joke, somewhat on the long side, in which a character writes a letter giving details of the WC in the village, under the impression that ‘WC’ means ‘Wesleyan Chapel’, not realizing that the letter-reader (and the joke-audience) interpret ‘WC’ to mean ‘toilet’. This leads to a succession of interlinked double entendres about matters such as seating arrangements and regularity of attendance. Notice that such a joke has some similarities to the class of joke outlined in [Section 8.6.2](#), in that the joke-audience is invited to be amused at the plight of a character who has misinterpreted something. Unlike the joke in [Section 8.6.2](#), this style of double entendre operates by having the character manifest this misunderstanding by producing text in which the joke-audience can perceive two meanings.

Even though a double entendre may not have the existence of a double meaning signalled by a punchline, the set-up/punchline division can still be applied to this class of joke, if we take the punchline to be the final, multiply interpretable, segment. In (67), the final clause *Can ... weekend?* constitutes the punchline. In (68), the whole text is ambiguous, but this could be viewed as having a null set-up (as in non-double-entendre ‘one-liners’, such

as (10) or (115)). The long joke based on the ambiguity of WC has a set-up consisting of the initial story, and a long, multi-sentence punchline consisting of the ambiguous section of text.

There is a closely related form in which there are two interpretations, but where the potentially amusing reading is the *more* OBVIOUS. Such a form typically occurs in maladroitness phrasings of mundane statements. Examples found in newspaper reports or advertisements are often seized upon and reported, or collected into anthologies (e.g. Parsons (1953)).

(69) Save time and cut fingers with a parsley mincer.

(*This Week* quoted by Parsons (1953:40))

(70) Imported Bedspreads by Glomar: Leaf design embroidered on fine polyester voile. Machine washable in champagne.

(*The Seattle Times* quoted by *The New Yorker*, re-quoted by Latta (1999))

Structurally, these are almost identical to a more conventional double entendre, in that there is another meaning, presumably intended by the writer. In (70), for example, the alternative reading has *in champagne* describing the colour of the items. In contrast to a typical double entendre, the prosaic interpretation, for this example, is less OBVIOUS than the ABSURD one. Such a difference is a very fine nuance, and might (for some examples) vary from audience to audience (as noted earlier, both the jokehood and the funniness of a text is dependent on the knowledge and attention of the audience). It may be that examples where the INAPPROPRIATE interpretation is more OBVIOUS have their full humorous effect only when the audience can also see the other interpretation: if the text seems simply to state baldly an absurd idea, that is not as funny, and may not even be seen as humorous. What can be concluded is that both forms can operate humorously: the INAPPROPRIATE element can be located in either the more OBVIOUS or the less OBVIOUS interpretation. (There is also the question of whether an accidentally humorous text counts as a joke, even when it is re-told for humorous purposes, but we will not digress into that here.)

It is straightforward to describe the structure of a double entendre in our formal framework ([Appendix A.2.1](#)): there is, either for the whole text or for its ending, an OBVIOUS interpretation and a less OBVIOUS interpretation. To make this more than simply an ambiguous text, the interpretations must have some further properties, such as INAPPROPRIATENESS, as discussed in [Section 8.6.4](#) below.

8.4

Prediction and contrast

In [Section 5.3.2](#), we discussed the issue of ‘prediction’ in the context of Suls’ two-stage model. The question is of broader applicability, and worth revisiting.

Prediction can happen at various linguistic levels – semantic, syntactic, lexical, etc. Such issues are familiar within computational linguistics and psycholinguistics (Crain and Steedman (1985), Fodor and Ferreira (1998)). Even in cases where a specific word seems to be predicted, it could be that the underlying prediction is at a semantic level, but is manifested lexically; alternatively, there might be an explanation in terms of statistics or semantic priming.

In the much-quoted (71), the final word overturns a very natural prediction of another lexical item (*table*).

(71) One more drink and I’ll be under the host.

(Dorothy Parker quoted by Katz (1996))

However, it is not obvious that this is the main factor in the joke status of the text, which is structurally a form of paradigmatic pun ([Section 9.3](#)). The additional impact created by the ruptured expectations may add to the humour, but as a *supplementary factor*; that is, an additional ingredient which enhances the humorous effect, while not being central to the text’s status as a joke (see [Section 12.3](#)).

A similar analysis could perhaps be presented for (72), or this joke could be seen as involving a predicted semantic content which enters into a humour-invoking relationship (some kind of CONTRAST) with the meaning of the punchline.

(72) Isn’t modern technology wonderful? I remember the excitement when we were the first family in our street to have cordless pyjamas.

(Arnold Brown, early 1990s)

The humour may involve the absurdity of cordless pyjamas being hailed as technologically exciting, but there is an implicit contribution from the notion of a *cordless telephone*. The question is: how is this phrase invoked? If we analyse the important relationship as mere similarity between the two *cordless* phrases, statically viewed, then this, like (71), has something in common with paradigmatic puns ([Chapter 9](#)). On the other hand, if we emphasize the dynamic aspect of text-understanding, and say that *telephones* is predicted, then we are classing this joke slightly differently. In either case, there can be CONTRAST between the two ‘cordless’ items, which may well contribute to the humour.

However we analyse (71) and (72), there is a strong intuition that the sudden deviation from expectation contributes to the impact of the joke, albeit as a supplementary factor (see above).

There are examples where it is more plausible that there is some kind of CONTRAST between a predicted meaning and the meaning of the punchline.

(73) A distinguished scientist was observing the heavens through the huge telescope at the Mt. Wilson Observatory. Suddenly he announced, 'It's going to rain.' 'What makes you think so?' asked his guide. 'Because,' said the astronomer, still peering through the telescope, 'my corns hurt.'

(Cerf 1964:102–3)

In (73), the set-up leads to the expectation of a learned scientific reason, rather than the banal folklore uttered. Neither the scientific nor the folksy account are either TABOO or ABSURD, and it seems probable that the humour comes from the difference between these two, i.e. some form of CONTRAST.

It is not hard to find examples like this where there seems to be a CONTRAST between a *predicted or expected* state of affairs and the actual outcome (sometimes an anticlimax), but there seem to be few (perhaps no) examples where the CONTRAST is incontrovertibly with a state of affairs *actually stated* in the set-up. There are some examples where the distinction is unclear, as in the (non-joke) example (74).

(74) A couple give each other, as simultaneous surprise presents, tickets for two particularly sought-after concerts. The husband smugly relates the devious sequence of contacts he used to obtain his tickets, and asks how his wife acquired hers. She replies, 'I phoned the box office.'

(from the US TV sitcom *Frasier*)

It is plausible to assert that the humour in (74) stems from the comparison of an elaborate, arcane method (explicitly described, *not* predicted) with a mundane, down-to-earth approach (also stated in the text). Alternatively, we could argue that there is an *expectation* that the wife's tale will be equally impressive, and the contrast is with that prediction. (There is a further angle, not relevant to our point about description versus prediction, that the relative popularity of the two concerts is reflected in the comparison.)

The rather long example (75) also seems to depend on the CONTRAST between two accounts of the mimed debate.

(75) Several centuries ago, the Pope decreed that all the Jews had to leave Italy. There was, of course, a huge outcry from the Jewish community, so the Pope offered a deal. He would have a religious debate with a leader of the Jewish community. If the Jewish leader won the debate, the Jews would be permitted to stay in Italy. If the Pope won, the Jews would have to leave. The Jewish community met and picked an aged Rabbi, Moishe, to represent them in the debate. Rabbi Moishe, however, could not speak Latin and the Pope could not speak Yiddish. So it was decided that this would be a ‘silent’ debate. On the day of the great debate, the Pope and Rabbi Moishe sat opposite each other for a full minute before the Pope raised his hand and showed three fingers. Rabbi Moishe looked back and raised one finger. Next, the Pope waved his finger around his head. Rabbi Moishe pointed to the ground where he sat. The Pope then brought out a communion wafer and chalice of wine. Rabbi Moishe pulled out an apple. With that, the Pope stood up and said, ‘I concede the debate. This man has bested me. The Jews can stay.’ Later, the Cardinals gathered around the Pope, asking him what had happened. The Pope said, ‘First I held up three fingers to represent the Trinity. He responded by holding up one finger to remind me that there was still one God common to both our religions. Then I waved my finger around me to show him that God was all around us. He responded by pointing to the ground to show that God was also right here with us. I pulled out the wine and the wafer to show that God absolves us of our sins. He pulled out an apple to remind me of original sin. He had an answer for everything. What could I do?’ Meanwhile, the Jewish community crowded around Rabbi Moishe, asking what happened.

Well,’ said Moishe, ‘first he said to me, “You Jews have three days to get out of here.” So I said to him, “Up yours”. Then he tells me the whole city would be cleared of Jews. So I said to him, “Listen here Mr. Pope, the Jews ... we stay right here!”’ ‘And then?’ asked a woman. ‘Who knows?’ said Rabbi Moishe. ‘We broke for lunch.’

(from www.jokesgallery.com)

Even here, it could be argued that there is a ‘prediction’ that the Rabbi’s exposition will be as theological as that of the Pope, and that there is a CONTRAST between predicted and actual. Notice in passing that the tone of the Rabbi’s contribution is established over several lines, which suggests that the humour is not created by a *sudden* reversal of predictions.

Another example where CONTRAST seems to be involved is (191) in [Section 11.5](#) below.

What this discussion suggests (not conclusively) is that we may need a notion of ‘prediction’ or ‘expectation’ in our text-processing model; a rough version of this is included in our formal definitions in [Appendix A.2.1](#).

The idea of prediction discussed here concerns expectation of specific content on the basis of the content of the set-up. This is quite different from any formal expectations that there will be a punchline, or that the text will follow some familiar presentation scheme (such as those discussed earlier in this chapter). Expectations about structure constitute some sort of meta-level prediction, whereby the ritual of joke-telling may arouse an anticipation that the text will accord with known joke patterns; see Deckers and Avery (1994).

8.5 Simple narration

There are many propositional jokes (in the sense of [Section 2.7](#)) in which the role of language is simply to convey the story of the joke; for example, (34), (35), (63). Such jokes do not even have the rather minimal requirement of requiring certain central facts to be withheld in the telling. It could possibly be argued that some examples rely on withholding information about the intricate thought processes of the central characters in these stories, but that is stretching a point. The characters in these stories do use language, and indeed the punchlines are often utterances by story characters, but this does not mean that the humour is dependent on the particular properties of the language forms used: any paraphrase would function as well (although some phrasings may work better than others, owing to subtle nuances of succinctness, rhythm, etc.).

It is trivial to state this in our formalization ([Appendix A.2.1](#)), although the definition describes any story, whether humorous or not; we will return to this point in [Section 8.6.4](#).

8.6 Other issues

8.6.1 *Misdirection and priming*

It has been observed (e.g. Dolitsky (1983, 1992)) that it may be necessary to withhold certain information from the set-up portion of a joke. Careful selection of information is crucial to forced reinterpretation jokes, where it can be seen as following these principles:

- nothing in the set-up must rule out the ‘hidden’ (less OBVIOUS) interpretation;
- nothing in the set-up should draw attention to the ‘hidden’ interpretation;
- it may be helpful to include details or choose phrasings which support or suggest the ‘non-hidden’ (more OBVIOUS) interpretation.

The first two of these principles are routinely adhered to in FR jokes (e.g. (18)). The third guideline is noticeably at work in (76).

(76) A pair of suburban couples who had known each other for quite some time talked it over and decided to do a little conjugal swapping. The trade was made the following evening, and the newly arranged couples retired to their respective houses. After about an hour of bedroom bliss, one of the wives propped herself up on her elbow, looked at her new partner and said ‘Well, I wonder how the boys are getting along.’

(quoted from Yamaguchi by Dolitsky (1992:41))

The main misdirection is in the choice of *conjugal swapping* rather than the more common (if less politically correct) *wife-swapping*, along with the vague phrases *the trade was made*, *newly arranged couples*. (This might seem to refute the conjecture in [Section 8.1.1](#) about the location of ambiguity within a story, but these misdirecting phrases are not, strictly speaking, linguistically ambiguous.)

A slightly subtler form of misdirection may occur by having one joke establish an expectation for another joke. Example (78), delivered phonetically, is typically preceded by (77).

(77) How do you get two elephants in a Mini?

One in the front, one in the back.

(78) How do you get /tʊ/ /weils/ in a Mini?

Across the Severn Bridge.

The first joke, (77), primes the audience to interpret the phonetic string /tʊ/ /weils/ as *two whales*, but the punchline treats it as meaning *to Wales* (this of course requires a dialect in which *wh* and *w* are indistinguishable). Notice that, in our terms, (77) is technically not part of the set-up, as that term applies to the part of the text which enters into some significant semantic/pragmatic relationship with the punchline. Instead, (77) is part of the context which influences the interpretation of the set-up (question) of (78). (Similarly, (1) is probably more effective if told after a few riddles of the form *What do you get if you cross ...?*) None of this has to be explicitly mentioned in our formal model. The basic text model that we are already assuming ([Section 3.9](#) and [Appendix A.1](#)) defines a text as having one or more interpretations, ordered by OBVIOUSNESS. This implicitly imposes the constraint that nothing must rule out the less OBVIOUS interpretation, and the second and third guidelines above more or less re-state the fact that the interpretations, for the chosen text, must be in a suitable order of OBVIOUSNESS. That is, the three principles above are, if anything, guidelines for constructing a text which meets our formal definitions, rather than being additional formal stipulations.

8.6.2

An outside view of misinterpretation

Although misinterpretation is a central mechanism in FR jokes, not all jokes where there is some misinterpretation count as FR jokes. For example, in (33), the misinterpretation (by the driver-character) is a contrivance within the set-up, not something revealed by the punchline.

Similarly, (79) uses linguistic ambiguity to create a misinterpretation, but not a reinterpretation of the set-up.

(79) There was a boy standing on a corner selling fish. He was saying, ‘Dam fish for sale, dam fish for sale.’ A preacher walked up and asked why he was calling them dam fish. The kid said, ‘I caught them at the dam, so they’re dam fish.’ The preacher bought some, took them home and asked his wife to cook the dam fish. His wife looked at him in bewilderment and said, ‘Preachers aren’t supposed to talk like that.’ The preacher explained why they were dam fish, and she agreed to cook them. When dinner was ready and everyone was sitting down, the preacher asked his son to pass him the dam fish. His son replied, ‘That’s the spirit, dad. Pass the fucking potatoes!’

(from www.jokes2000.com)

In (79), both interpretations of *dam(n) fish* are openly available (i.e. there is no hidden interpretation from the audience’s point of view). The misinterpretation in the punchline is by a character in the story. The humorous effect cannot derive from sudden revelation (to the audience) of a hidden meaning, but must be based on amusement at the actions of the misinterpreting character. The presentational structure in this joke is the simple narrative ([Section 8.5](#)) rather than any more oblique mechanism: the narrative describes a situation in which a character, having made a misinterpretation, makes a fool of himself. (See [Section 12.3](#) for some related discussion.)

In the traditional verbal/referential classification ([Section 2.7](#)), (79) would be a verbal joke.

8.6.3

What’s a punchline?

Although we have presented some textual structures in which the division between set-up and punchline seems fairly clear, this may not always be the case.

(80) A rabbit goes into a butcher’s shop and asks, ‘Have you got any lettuce?’ The butcher says, ‘We don’t sell lettuce here. You need the greengrocer’s across the road.’ The next day the rabbit comes into the shop and asks for some lettuce again. The butcher tells him,

Look, I told you yesterday, we don't sell lettuce. You need the green-grocer.' The rabbit comes in the next day and asks the butcher again, 'Have you got any lettuce?' The butcher goes mad. He says, 'Look, I'm sick of this. How many times do I have to tell you I don't sell lettuce. If you come in here asking for lettuce, I'm going to nail your ears to the floor.' The next day the rabbit comes in and asks the butcher, 'Have you got any nails?' 'Nails? No.' 'Right,' the rabbit says, 'Have you got any lettuce?'

(Ross 1998:43)

In (80), it seems plausible that the set-up continues at least as far as the phrase ... *and asks the butcher*. However, the remaining text (four utterances) then produces the amusing effect, and it is far from obvious where to draw a line part-way through this sequence. Perhaps all of this textual string constitutes the punchline. In contrast to the text structures outlined above (such as punchline conflict), there is no clear linguistic evidence for separating some particular end segment which has a special relationship with all that has gone before.

8.6.4

Further factors

The definitions given in this chapter (and formally in [Appendix A](#)) describe classes of *text* rather than classes of *joke*, in that a text could meet one of our definitions without being in any way humorous. A simple misunderstanding, or an accidentally ambiguous utterance, would be covered by these definitions even if completely unhumorous. In our earlier discussion of the FR model ([Chapter 5](#)), we suggested that an FR joke required some element of INAPPROPRIATENESS or CONTRAST to distinguish it from a mere misunderstanding. This would be statable more precisely using our formalization of FR jokes ([Appendix A.2.1](#)). The other types of joke discussed in this chapter would similarly need some further condition to transform them from describing types of text to defining classes of joke. We will explore this issue in more detail in [Chapter 12](#).

8.7

Summary

In [Chapter 5](#) we saw that there were, in the existing literature, (at least) two separate variants of the so-called incongruity-resolution model which made different claims about the structure and mechanisms of jokes. We have examined one of these variants further, and shown that there are at least two separate subvariants within it. We have studied some superficially similar, but slightly different, classes of joke. Formal definitions based on our discussion can be found in [Appendix A.2.1](#).

All of this analysis is directed at particular classes of joke, and does not purport to cover all jokes, not even all incongruity-resolution jokes, if such a class can be defined. Moreover, we have focussed solely on the way that *information is presented* in these (propositional) jokes: additional conditions on texts, to qualify them as jokes, will be discussed in [Chapter 12](#).

9

The structure of puns

We review a range of jokes involving phonetic similarity, and offer a structural definition which describes a substantial class of these jokes.

9.1

Puns

One of the simplest forms of joke is the *pun*, often known as a ‘play on words’. It is this very simplicity which makes puns of interest here. As argued in [Chapter 1](#), it is a useful first step (in the quest for an understanding of humorous phenomena) to be able to set down some ‘abstract syntax’ for a class of jokes, thereby clearing the ground for further investigation of how that class of jokes is used and why they are or are not funny. (See Attardo (1994: [Section 3.4](#)) for some further reasons for studying puns.) Here, we will attempt to find some principles which describe the basic structure of puns in English.

Puns are (at least within English-speaking culture) a very widespread and commonplace form of spontaneous humour. Whereas funny stories such as (18) or (62) are not normally created on the spur of the moment in ordinary social interaction, puns – albeit not very funny ones – are made quite freely. This may, of course, be another consequence of the structural simplicity of puns – anyone can build one:

It is the easy availability of puns which makes them a cheap and somewhat despicable type of humor for many individuals and social groups.

(Raskin 1985:141)

This means that an analysis of punning should take us (very slightly) beyond the limitations of studying only ‘canned’ jokes, i.e. jokes which are recounted in a relatively fixed textual form with little or no connection to the context.

Many puns, perhaps the majority, are not very funny. There is even a widely established habit of groaning in response to a pun rather than laughing. This is yet another empirical phenomenon which we will not

be able to explain unless we take the first step of defining what constitutes a pun.

Although we will present all the puns as written texts, this should be taken as conveying a basic phonetic form, as relations of phonetic similarity are crucial to the operation of puns (see Sections [3.2](#) and [9.7.4](#)).

We shall argue that ‘plays on words’ vary in their formal structures, and shall therefore divide this pre-theoretic group of jokes into several different subclasses. There are certain relatively systematic jokes ([Section 9.3](#)), for which we shall propose a fairly simple general structure. Then there are various subclasses where the analysis is less clear: those where the broad structure is clear but there is some essential factor which is hard to ascertain ([Section 9.4](#)), those where the regularities are clear but hard to formalize using conventional linguistic representations ([Section 9.5](#)), and a mixed collection where the mechanism is hard to formulate ([Section 9.6](#)).

A distinction is sometimes made (see various taxonomies reviewed in Attardo (1994: Ch. 3)) between *paradigmatic* puns and *syntagmatic* puns. In the paradigmatic variant, a particular substring appears in the text, and the joke depends on the similarity (or even identity) of that string to some other string not in the text. A syntagmatic pun has two (or more) substrings actually in the text, whose similarity (or identity) is the basis of the pun. (Cf. the use of these terms in traditional linguistics (Hudson 1971: Ch. 1), (Lyons 1977a: [Section 8.2](#))).

We shall analyse the paradigmatic form ([Section 9.3](#)), and then move on to the syntagmatic form ([Section 9.4](#)). First, however, we shall review some of the existing work in this area ([Section 9.2](#)).

9.2

Previous work

Puns have been written about copiously over recent centuries (see Sobkowiak (1991: Ch. 1) for a review), but most academic works are firmly rooted in disciplines such as literary studies or psychoanalysis, where a detailed structural analysis of the mechanisms of punning is not of interest:

The research literature on puns is vast, but largely inconclusive ... most of the studies on puns in literary criticism ... having little to say from a structural, or essentialist, perspective’ (Attardo 1994:111). Redfern (1984) discusses a wide range of data, but always very informally and never analytically (as Raskin (1987) caustically observes). Culler (1988) is similarly a collection of literary essays which take for granted that the reader knows what a pun is, and discuss how puns have tended to be used by various authors. Jespersen starts promisingly with this observation:

In the popular speech of all nations are found instances of a peculiar class of round-about expressions, in which the speaker avoids the

regular word, but hints at it in a covert way by using some other word, generally a proper name, which bears a resemblance to it or is derived from it, really or seemingly. The proper name used may be that of a place or of a person; it may be a name of real existence or one made only for the sake of the punning allusion.

(Jespersen 1933)

However, after this draft definition (which seems rather narrow, given its concentration on proper names), the rest of the article is nothing more than a dictionary-style list of examples of idioms and puns which involve real or invented proper names.

(81) Why doesn't a railway engine sit down?

Because it has a tender behind.

(Wright 1979:409)

Wright (1979) analyses the punning riddle given in (81) as containing a *neutral sensory presentation* (the phonetic form of *tender behind*), and two *context-indicators* (*railway engine* and *doesn't ... sit down*) which correspond to two *interpretations*. This very brief sketch (part of a much wider discussion of perceptual effects) has the same central intuition as our definition of a subclass of puns ([Section 9.3.6](#)), but is not developed in much detail and Wright does not consider a wide range of puns.

Sobkowiak (1991) gives an excellent discussion of the nature of the phonological (and metaphonological) knowledge which is involved in puns and word manipulation of various types, along with some statistical data on the occurrence of phonological patterns within certain corpora of puns. Although that work is thorough, methodical and thoughtful, its focus is on the nature of phonetic regularities, particularly phonological similarity, which is not what we are examining here.

Cornwell and Hobbs (1991) say that (despite the title of their article) they are *not* offering a 'definitive definition' of a pun. Instead, they discuss some aspects of punning, and make a number of observations that we would agree with: puns are ubiquitous, not all verbal ambiguities are puns, and if the punned item occurs more than once in the text it must be with different meanings. They also make the distinction between *initiator* jokes and *responder* jokes, which we shall return to in [Chapter 11](#). They claim that puns involve ambiguity, but say that this may include strings which are similar rather than identical phonetically. This is an interesting step: they are redefining a fairly well-established term – *ambiguity* – in order to preserve their generalization that puns always involve ambiguity. Here, we adopt a notion of ambiguity which seems to reflect traditional usage ([Section 3.7](#)), and build our definitions of puns from there.

Attardo (1994) provides an extensive review of the literature on puns, along with some proposals of his own. It is quite hard to discern what Attardo regards as the defining properties of puns, and hence his theoretical analyses are difficult to follow. He is certain (pp. 112, 133) that puns require ambiguity, with the ‘working definition’ that ‘any string that can refer to more than one object is ambiguous’. He appears to include ‘vague’ or ‘unspecified’ terms as ambiguous, though he gives no examples of puns using these subclasses of ambiguity. Also, he – quite reasonably – allows puns to include examples of paronymy (phonetic similarity but not identity). This means that, as with Cornwell and Hobbs, we cannot rely on traditional notions of ambiguity, nor even his own working definition. Attardo also cites, and adopts, a widespread consensus that puns always ‘involve two senses’. The meaning of ‘sense’ (or perhaps ‘involve’) would have to be stretched to cover examples such as those in [Sections 9.3.2](#) or 9.3.3 below, where the mere existence of another matching string seems to cause the pun regardless of its meaning. On the other hand, it is possible Attardo would not regard these as puns. He also says that one of the conditions for a joke to be present is that the two senses must be *local antonyms* (Raskin 1985); see [Section 6.4.1](#) for discussion of this term.

(29) Why did the cookie cry?

Because its mother had been a wafer so long.

Attardo fits (29) into Raskin’s SSTH ([Chapter 6](#)) by effectively having a null punchline: the end of the joke-telling stimulates the reader to look for another reading to the text (cf. discussion in [Section 8.3](#)). He also observes that the two senses of /əweɪfər/ are, as required, local antonyms, but does not explain why. His analysis is wholly *procedural*, in the sense that he outlines a process whereby a hearer might interpret (29); he does not offer a declarative or structural account of the precise properties that qualify this text as a joke (see [Section 2.10](#)). Although Attardo’s proposals are more elaborated than those of most other authors, they nevertheless are slightly lacking in detailed or formal definitions, and the logical connections between the data and the theory are not always clear.

Heller (1974) says that in all puns ‘a single manifesting mark signals more than one conceptual function’. He then sets out two important abstract properties of puns, each of which may have two possible values (thus leading to four major categories):

- (a) all the ‘conceptual functions’ may be evident when the ‘manifesting mark’ (word, phrase, etc.) is presented, or some may be concealed.
- (b) by the end of the joke, all the possible ‘conceptual functions’ may be retained as possibilities, or some may be removed (disambiguation).

This scheme is intended to cover not only simple puns where an ambiguous word (or word/phrase similar to another) occurs at the end of a joke, but also those such as (82), where an early (unobvious) ambiguity is resolved in an unexpected manner later in the text.

(82) Come forth, lazarus! And he came fifth and lost the job.

(Cited by Heller (1974) from Joyce's *Ulysses*)

As we will discuss in [Section 9.7.3](#), we have opted for a division into joke classes which analyses such jokes differently from more basic paradigmatic puns. Heller discusses at some length some of the linguistic (or orthographic) mechanisms which can contribute to the similarity of communicative signs in puns, a matter which we treat as a subsidiary notion (see [Section 3.8](#)).

Heller also extends his notion of pun to cover gestures and actions (an idea mentioned in passing by Attardo). He cites one amusing or embarrassing situation (typical of those used in situation comedies) based on a misinterpreted gesture, and cites (83) as a joke based on misinterpreted actions, classed in the same group as (82), but 'secondary or displaced (sociocultural)'.

(83) A sociology professor at the local college was fond of telling off-color jokes, a fact which embarrassed the young ladies in his class tremendously.... Finally the girls got together and decided that at the next hint of an off-color joke they would get up *en masse* and leave the classroom. They sat there primly awaiting his next lecture. Sure enough, he began with the words, 'There is a terrible shortage of prostitutes in Singapore.' The girls looked at each other, rose, and started to leave. He called after them, 'Don't go now, girls. The boat doesn't leave for Singapore until Friday.'

(Heller 1974:281)

This example would fit into the forced reinterpretation (FR) class (Chapters [5](#) and [8](#) above), being similar in many respects to (62).

It is clear that puns, loosely named, and the kind of jokes we have classed as FR are two major genres of joke, but there is no agreement on the defining characteristics of these classes, or whether they are even distinct. Attardo (via adoption of the SSTH) seems to subsume puns under FR, whereas Heller's typology of puns is wide enough to cover FR jokes. As we will discuss in [Section 9.7.3](#), we opt to regard these genres as different structural classes, at the cost of putting some examples which are informally puns (e.g. (2), (28), (82)) in the FR class, thereby perhaps forfeiting the right to claim that our definitions in the succeeding sections characterize 'puns' as they are informally labelled.

9.3 Paradigmatic puns

9.3.1 *Puns in context*

We shall start by returning to first principles, with some simple examples, and then gradually build up some generalizations, confining attention at this stage to paradigmatic puns. In most chapters in this volume, we generally dissect individual jokes in isolation, without much consideration of the context in which the joke is delivered. However, in the case of puns, we shall start by widening the scope of our analysis to take some account, albeit rather minimal, of the context (perhaps non-linguistic) of the joke-telling. The reason for this is twofold:

- as noted above, puns often occur spontaneously in real-life interactions, so it is desirable to have our description cover that form of joking;
- we will propose a generalization regarding puns which shows how spontaneous non-linguistic context and ‘canned’ linguistic context play a similar role.

To do this, we shall posit an abstract type of entity called a *context* (see [Section 3.9](#) and [Appendix A.1](#)). Informally, a context may include facts about the world, cultural information, salient objects in the surrounding environment, recently mentioned concepts, etc.

Puns which occur spontaneously in everyday life are not often very funny, and are rarely recorded for later use; hence, there is no obvious source of collected spontaneous puns. Some of the examples here (in particular, (86), (87), (89), (93)) were witnessed by the author in ordinary situations.

The first point is that simple use of a phonetic-lexical ambiguous word does not constitute a pun; it is not a *sufficient* condition. (84), uttered as a statement about a species of animal, is not a pun, despite the use of the ambiguous word *shell*.

(84) The tortoise has a protective shell.

(constructed item)

Even where the alternative meaning of a word could in principle make sense in the context of use (i.e. there is no semantic incompatibility with the surrounding linguistic material), ambiguity is not sufficient to form a pun, as shown in (85), where *shell* could equally refer to a discarded marine carapace or a stray artillery round.

(85) John found a shell on the beach.

(constructed item)

We cross the boundary to puns with (86), where *leak* is meaningful within the utterance and the homophonous *leek* is related to the surrounding context.

(86) A shopper is walking along, and a leek falls from his shopping bag to the ground, unnoticed. Another shopper calls out, 'Hey! Your bag's leaking!'

This illustrates the fact that if the compared items are phonetically identical, they must be lexically distinct (otherwise they would just be the same item, not two).

The crucial point in (86) is that both senses of the string *leek/leak* are relevant to the context of the utterance. The rather similar (87) does not, apparently, meet that condition.

(87) A shopper is standing with a bag containing groceries, including a prominently visible leek. A bystander remarks, 'You've got a leek in your bag.'

As it happens, (87) was observed in a situation where the speaker, by intonation, facial expression and posture, attempted to imply that there was some problem with the shopper's bag; that is, the speaker tried to contrive a context in which the *leak* sense *would* be appropriate, thereby manufacturing both a pun and a slight practical joke.

Phonetic-lexical ambiguity is not a *necessary* condition for a pun. In (88) and (89) there is a looser form of phonetic similarity between strings.

(88) Some South American stamps are un-bolivia-ble.

(Headline cited in *Have I Got News For You*,
BBC TV, May 2002)

(89) A man walks along a city street carrying what is obviously a fender (in the British sense): the ornate metal surrounding for a hearth. The object is a square U-shape, and he is carrying it so that it surrounds him on three sides. A passer-by observes 'Home is where the hearth is.'

In (88), there is similarity, but not identity, between *bolivia* and *believe*. The pun in (89) depends partly on the phonetic similarity between *heart* and *hearth*. The context is semantically linked to *hearth*, and arguably also to the whole phrase *home is where the hearth is*, as an observation implying that the man's hearth is in some way his portable residence. Equally crucial is the existence of the motto 'home is where the heart is', which allows the text *home is where the ... is* to evoke the word *heart*.

These examples illustrate certain central aspects of puns:

- (a) There is an implicit comparison between two textual strings which are to some extent phonetically similar. As will become clear from some examples, the degree of similarity may be quite weak (*paronymy* (Attardo 1994), *paraphony* (Dienhart 1999)).
- (b) One of these strings is part or all of the utterance.
- (c) At least one of the strings is semantically linked to the context in some way.

We shall refine these generalizations later.

An even finer (but necessary) distinction is that the portion of the utterance may be phonetically similar to (but lexically different from) only a subpart of the well-known phrase, as in (89). This means that in framing a detailed formal analysis, we have two options. In (89) we could consider the pun to be based on a comparison between *hearth* and *heart*, in which case the licensing motto consists of the substitution of the matching text *heart* in the utterance *home is where the hearth is*. Or we could analyse the pun as using the comparison between *home is where the hearth is* and *home is where the heart is*, in which case the established motto participates directly in the analysis, but only one part of it, *heart*, manifests lexical difference from the utterance. All these aspects are captured more formally in [Appendix A.2.1](#).

9.3.2 *Malapropisms*

A *malapropism* is a well-established humorous form, named after Mrs Malaprop, a character in Sheridan's 1775 comedy *The Rivals*. It involves the use of a word which sounds similar (but not identical) to an appropriate word for the context, but which is quite different in meaning. An example is (90).

- (90) Dorothy Stickney, invited to dinner with the President of the United States, was so nervous she quavered to him, 'What an imposing building the White House is! Who was the artichoke who designed it?'

(Cerf 1964:13)

Structurally these humorous items are extremely similar to puns, even if they are not usually so classed. (Some empirical work by Sobkowiak (1991) suggests that malapropisms typically allow a more liberal notion of 'phonetic closeness' than in a conventional pun.) In general, a malapropism is not delivered with humorous intent, but happens either through ignorance (the typical case) or a slip (as in (90)). However, a malapropism can

be used knowingly, as a deliberate joke. In the 1960s, a British boxer, to convey his minor role in the machinations of the sporting world, made the remark (91).

(91) I'm only a prawn in the game.

Although this was a classic malapropism, based on an error, it has been cherished and repeated by people when a suitable context arises. Also, this author witnessed (on a number of occasions), example (92), which was clearly said with humorous intent and with full knowledge of the incorrectness of the phrase (supposedly similar to *au revoir*).

(92) On departing for work, a person bids his family 'Olive oil.'

In our relatively limited investigation, we are not taking account of the speaker's intention, so the distinction between an unconscious and an intended malapropism will be ignored here.

Some of the distinctive structural features of a malapropism are:

- the uttered word/phrase is a genuine form in the language, not an artificially coined form just for this utterance;
- the uttered word/phrase has no connection whatsoever to the context – if it did, the utterance would be more likely to be classed as a pun;
- the utterance as spoken does not make sense in context, but it would be perfectly appropriate if the unuttered similar word/phrase were substituted into it.

In (90), the strings are *artichoke* and *architect*, where the latter is completely appropriate to the context and the former is not even loosely related. Whereas in (89) the utterance contains the *suitable* textual item, a malapropism contains the *unsuitable* item. This is essential, because the phonetic dissimilarity between the two strings means that uttering the appropriate form (*architect*) would not evoke the other item (*artichoke*): the mere presence in the language of a potentially similar (and contextually unrelated) phrase is a widespread phenomenon, and does not make an utterance amusing. There is a similar pattern in (88), where the utterance contains a valid word only if the other string, *believe*, is substituted.

A further observation is that the humorous effect of using a similar but incorrect word varies according to the degree of unsuitability of the word (or perhaps some form of absurdity/incongruity). For example, in everyday writing or speaking, people sometimes misuse certain words, such as *abdicate* for *abrogate*, or *disinterested* for *uninterested*. Structurally, these are malapropisms (and are usually perpetrated inadvertently), but they are not usually regarded as amusing. See [Section 9.7.7](#) for discussion.

9.3.3 *Textual contexts*

Although the above examples, such as (86) and (89), were cited here in textual form, the text served to describe a situation which actually occurred and where the context was not itself textual.

Utterances spoken shortly before the putative pun can also supply context, as in (93). (93) Some people are talking about their friend, Dick. Someone says 'Dick's ears are so filthy. You could plant potatoes in them.' Another remarks, 'Hm ... Dick 'taters.'

In (93), the two strings are *Dick 'taters* and *dictators*. The first is linked to the context, the latter is not linked to the context at all. A further point is that the lexical segmentation (i.e. the decomposition of the phonetic string into words) may differ across the two compared items. In (93), for example, the two words *Dick 'taters* are matched with the single word *dictators*. Moreover, if there were no such integrating connections (i.e. if two unconnected words had matched separately to *Dick* and *taters*) this would not be a pun. This appears to be crucial. As noted above with (89), the compared text which is *not* part of the utterance must have a sufficiently clear identity that the matching material in the utterance evokes it relatively easily. This identity can be, as in (93), as a word, or, as in (89), as a well-known phrase (a 'mimetic phrase' (Nash 1985:140)). If there were no such concatenation, nor any contextual linkage, and the spoken utterance contained a contextually appropriate word, the sentence would be a non-pun, as in (84).

Puns can be based on a non-linguistic context by being located in the name of a shop, restaurant or similar business, such as (94) or (95).

(94) A shop selling greetings cards: 'Cardiology'.

(95) A hairdressing salon: 'Curl Up and Dye'.

All of these conform to the pattern outlined above, where there are two phonetically similar strings, at least one being linked to the context, and the other being sufficiently well-known to be clear to the reader/hearer. The phonetically identical string may be just part of a word, illustrated, for example, in (94) above (*card*).

Another variation on the notion of context applies to newspaper headlines, which (at least in the UK) quite often contain or constitute puns (see Alexander (1997: Ch. 6) for a review of this phenomenon). In such a case, the context is the news story which follows the headline, so it could be said to be *succeeding* text which constitutes the context. There are also instances where a picture in a newspaper may form part or all of the context; that is, a pun located in a caption.

- (96) The UK's *Guardian* newspaper once published an article about an ongoing period of prosperity in Albania, along with a picture of the capital, Tirana, captioned 'Tirana booms today'.
- (97) A minor football team known informally as 'Caley Thistle' (where *Caley* rhymes with *alley*) soundly defeats Celtic (then the top team in the country) in a major competition. The next day, a newspaper headline reads: 'Super Caley Go Ballistic, Celtic Are Atrocious'.

(*The Sun*, 9 February 2000)

In (96), the entire phrase *Tirana booms today* is strongly linked to the context. It is similar to the word/phrase *tararaboomdeay* (from an old popular song) which is not linked in any way. Similarly, (97) has a strongly linked utterance, phonetically similar (though not very) to the song title and invented word *Supercalifragilisticexpialidocious*.

The context may contain both textual and broader non-textual information, as in (98).

- (98) In the 1960s, a UK Labour politician was having difficulty with the trade union movement, led by Frank Cousins, and also with his own party members, whom he was accustomed to addressing as 'brothers'. A political opponent remarked that his plans were being disrupted by 'his brothers and his Cousins'.

Here the two texts being compared are *cousins* (the kinship term) and *Cousins* (name of the trade union leader). The former is linked to the immediately preceding text item *brothers*, the latter is linked to the overall political context.

In examples such as (96), (97), and (98), it is not clear that there is a significant difference caused by the context being wholly or partly textual. The linkage between the punning items and the contextual utterances is not primarily textual in the sense of phonetic similarity or string matching; it occurs at the semantic or conceptual level. Hence, the contextual utterances are merely conveying particular propositions, or making salient certain concepts, which the pun items can then be linked to. If those propositions or concepts had been introduced or given prominence visually or by other non-textual means, they would still be available for punning, as in examples such as (89).

9.3.4

Puns with their own context

In the preceding section, we examined puns which were made relative to a context which may be wholly non-linguistic or which may be to some extent linguistic. To do this, we had to attempt, using language, to characterize

the context of joke-making, thus (reluctantly) converting the jokes into wholly textual items, at least as they appear within this chapter. The hope was that the reader would recognize these attempts at meta-level explanation of contexts as not being part of the joke itself, and accept that the original joke-telling was triggered by something other than the sequences of words that were used here to outline the context.

In contrast to this, there are puns which carry their own context around with them in textual form, making them self-contained and capable of being recounted in almost any actual context. We will now consider these ‘packaged’ puns, and argue that the same generalizations hold concerning the structure of these puns, providing we take the textual set-up to constitute the context.

Story puns

Perhaps the clearest example of a class of pun which establishes its own context verbally is what Binsted and Ritchie (1996, 2001) call the *story pun* and Norrick (1986) calls the *poetic story joke*. These consist of a narrative, sometimes quite long and rambling, which may not be at all humorous in its content, and which concludes with a phrase or sentence which acts either as a suitable final line in the story (e.g. a remark by a character) or as a summing up of the moral of the story. Crucially, this final punchline also sounds similar (sometimes with considerable distortion) to a well-known phrase or saying (e.g. a motto or famous quotation). Muir and Norden (1991) offer a large collection of these (all too long to be quoted here, unfortunately). (99) is an example where the pun is within the story, as a character’s utterance, and (100) is the more classic form where a moral is drawn.

(99) King Arthur had lots of knights who fared forth on coal-black charges to rescue beautiful maidens from dragons’ clutches, but did you ever know that one of them was mounted on a St Bernard dog? His name was Sir Marmaduke, and he and the St Bernard performed many a deed of derring-do. One evening, however, they were caught in a torrential thunderstorm, and sought shelter at a nearby tavern. ‘Reservation?’ asked the room clerk. ‘No,’ admitted Sir Marmaduke.

Sorry,’ said the clerk, ‘no room without a reservation.’ It was at this point that he discovered that Marmaduke was sitting astride his faithful St Bernard. ‘Hold on,’ said the clerk, ‘We’ll have to find *something* for you. I wouldn’t put out a knight on a dog like this.’

(Cerf 1964)

(100) Once upon a time, many years ago, there was a chieftain in a remote tropical village who owned an old and battered throne of which he

was very fond. One day, a visiting dignitary gave him a brand new and ornate throne, which the chieftain had to adopt immediately out of politeness.

However, he could not bear to part with the old throne which had served him so well, so he stowed it away in the roof area of his grass hut, in case it should be useful in the future.

Unfortunately the interior structure of his hut was too flimsy to support the weight of the large object, and it crashed through the grass ceiling, falling on the chieftain and killing him. The moral is that people who live in grass houses shouldn't stow thrones.

(Binsted and Ritchie (1996, 2001))

(For readers outside the relevant subcultures, it is necessary to know that there is a saying 'I wouldn't put a dog out on a night like this', used to comment on very inhospitable weather, and there is a proverb 'People who live in glass houses shouldn't throw stones'.)

The parallel between these and context-based puns should be fairly clear, particularly in the case of (89). The preliminary story establishes a context, albeit a fictional one, and the punchline is related to that context in just the way that was described in [Section 9.3.1](#).

Although story puns are in a sense 'funny stories', they do not, as Binsted and Ritchie point out, share many of the attributes of the conventional narrative joke. They are certainly not describable in the same terms as the various classes of narrative jokes described in [Chapter 8](#).

Example (101) is also a pun with its own context, and could perhaps be classified as a story pun, though its punchline achieves pun status through similarity to a single word, rather than to a well-known phrase or saying.

(101) A bear walks into a bar in Baraboo, Wisconsin, and sits down. He bangs on the bar with his paw and demands a beer. The bartender approaches and says, 'We don't serve beer to bears in bars in Baraboo.' The bear, becoming angry, demands again that he be served a beer. The bartender again tells him, 'We don't serve beer to bears in bars in Baraboo.' The bear, very angry now, says, 'If you don't serve me a beer, I'm going to eat that lady sitting at the end of the bar.' The bartender once again says 'Sorry, we don't serve beer to bears in bars in Baraboo.' The bear goes to the end of the bar, and as promised, eats the woman. He comes back to his seat and again demands a beer. The bartender says 'Sorry, we don't serve beer to bears in bars in Baraboo that are on drugs.' The bear says, 'I'm not on drugs.' The bartender says, 'Yes you are, that was a bar bitch you ate.'

(from www.jokesgallery.com)

This joke falls under the generalizations of [Section 9.3.1](#) if we regard the set-up story as forming the context.

Example (100) illustrates a common feature of puns: there are a few linguistic mechanisms acting on different parts of the textual form, causing an overall ‘similarity’ for some longer string of words. In (100) there is an instance of metathesis (*stow thrones/throw stones*) and an instance of phonemic similarity (*grass/glass*) combining to create a degree of similarity between the entire phrase *people who ... thrones* and the known proverb. Similarly, in (99), the pun is achieved not simply by the use of the ambiguous word *knight/night*. The linguistic effect uses a combination of two devices – homophony and metathesis – to create a similarity between the strings being compared – the phrases *a dog on a night like this* and *a knight on a dog like this*. So although ambiguity is an ingredient of (99), there is more to the issue of linguistic similarity than simple phonetic-lexical ambiguity (see Attardo (1994:115–8) for a review of some devices used within puns).

Punning riddles

There is no standard definition of a *riddle* (see [Section 11.4.2](#) for more discussion of this term), but many riddle-style jokes involve word-play. Typical examples are:

(102) What kind of murderer has fibre?

A cereal killer.

(Binsted 1996)

(103) What do you get when you cross a murderer with a breakfast food?

A cereal killer.

(Binsted 1996)

(104) What do you use to colour an animal?

Hare dye.

(Binsted 1996)

(105) What did the python say to its victim?

I’ve got a crush on you.

(Dienhart 1999)

These, and many like them, have similar patterns to the puns discussed earlier in this chapter, assuming the question part (set-up) constitutes the context. In (105), the punchline can be viewed as a way of saying *I’m crushing you*, similar to the idiomatic phrase *I’ve got a crush on you* (referring to infatuation).

Examples (102), (103), (104) exemplify joke-classes which Binsted's JAPE program (see [Section 10.5](#)) can construct. We assume that the communicative alphabet is phonetic, and that differences of spelling would not be detectable in the presentation of the joke. The wordings here, where one particular spelling is chosen, are those used by Binsted. As in example (14) earlier, one of the possible spellings may be more helpful to the audience in 'getting' the joke. All of these have a set-up question which introduces two concepts, and a punchline answer which contains words which are each linked to one of these concepts, and which together form a recognizable compound noun. Binsted (or more accurately, JAPE) also offers (106), which seems to be a slightly malformed pun.

(106)What do you call a naked bruin?
A grizzly bare.

(Binsted 1996)

As in some other JAPE examples, we appear to have to treat *bare* as a noun. Even then, there are two set-up concepts (nakedness, bear) and a punchline noun phrase whose head noun [*bare*]/[*bear*] has two possible lexical forms, each related to one of the set-up concepts. The adjective *grizzly* does not seem to contribute much to the pun. As with some earlier examples, the essential ideas of these jokes can be conveyed without using a question-answer form, for example (107), (108), and (109). (See [Chapter 11](#) for further discussion of the contribution of the question format.)

(107)To colour an animal, you could use hare dye.

(constructed item)

(108)When you cross a murderer with a breakfast food, you get a cereal killer.

(constructed item)

(109)As the python said to its victim, I've got a crush on you.

(constructed item)

(The last of these forms is quite common: 'As the X said (to the Y), *phrase punning on properties of X (and Y)*'.)

(110)What do you call a depressed train?
A low-comotive.

(Binsted 1996)

Example (110) is another example where just part of a word (*lo*) participates in the relationship of phonetic similarity. (Remember that we are assuming a phonetic communicative alphabet, so the difference in spelling is a mere artefact of the presentation here.) Both *low* and *locomotive* are related to the context.

Similarly, in (111), there is a resemblance between one part of the answer (*tub*) and the corresponding part of a known phrase (at least in the UK), *pub crawl*.

(111) What do you call a bath tour?

A tub crawl.

(Binsted 1996)

See [Section 9.3.5](#) below for further discussion of the formation of the punchlines for these riddles.

9.3.5

Linguistic acceptability

As argued in [Section 3.5](#), it is reasonable to adopt a general convention that humorous texts are grammatically well-formed, and, apart from their special humorous properties, are normally structured texts, broadly speaking. For many classes of joke, there is a further generalization that can be made, namely that the text is semantically and pragmatically well formed. That is, the text follows normal conventions of the language, such as the need for an answer to reply to a question, the tendency for consecutive sentences in a discourse to be coherently related, etc. When considering paradigmatic puns in context, this stronger constraint does seem to hold for most examples, although we shall consider some exceptions in [Section 9.5](#) below. When someone makes a pun, there is a well defined utterance which not only is grammatically and semantically well formed, but also fits pragmatically into the context. Usually this will be the actual string uttered (as in (86) or (97)), but in malapropisms (and examples like (88)) the contextually appropriate utterance requires switching the two similar strings. That is, the unspoken string must be substituted for the spoken one to form the semantically and pragmatically fluent utterance. The other string may be completely irrelevant to the context (as in malapropisms and (97)) or may be loosely associated (as in (98)). Although we shall talk of the string, or the text, as being appropriate to the context, what we mean is that if the text is assigned a conventional syntactic, semantic and pragmatic analysis, these will be felicitous in the context, without resorting to specialized devices.

It is straightforward to convey information with such a pun, or to integrate it into normal conversation. Many of the examples cited earlier in this chapter, while uttered for humorous effect, typically have a communicative role as well.

It might seem that the punning riddle examples of [Section 9.3.4](#) do not meet the criterion of linguistic appropriateness, as they contain hitherto unknown phrases such as *cereal killer* or *tub crawl*. However, it could be argued that these examples conform, with only a small amount of licence from the playful use of language. In English, noun-noun compounds (such as *hair dye*) are idiosyncratic in their semantic content. That is, the semantic relation between the modifier and the head is not predictable: a *dog collar* is a collar for use on a dog, but a *dog show* is a show featuring dogs, and a *flea collar* is a collar for eradicating fleas. Jokes such as (102) and (104) could be viewed as postulating noun-noun compounds with semantic relations which accord with the set-up question; in the case of *hare dye*, it is not even a particularly strange relation, given the existence of *hair dye*.

In the case where part of a word is substituted, as in (110), the linguistic justification for the formation is a rule that where a new notion is formed which in some way relates to or combines two existing ideas, a suitable word is some splicing of the nouns for these items. This is not a wholly fanciful rule, given formations such as *beefburger*, *workaholic*, *Irangate*.

Sometimes the spliced word may be a complete neologism, as in (112).

(112) What does a near-sighted ghost wear?
Spooktacles.

(Webb (1978) cited by Binsted (1996))

Nothing in the general definition we give below (and more formally in [Appendix A.2.2](#)) excludes non-words, as long as *at least one* of the strings involved meets the condition of being ‘linguistically appropriate’ to the context. Here, it is arguable that *spectacles* does meet this condition (although it does not address the question of a ghost’s needs perhaps being different from those of non-ghosts).

(Examples like (112) provide the motivation for allowing, in [Appendix A.1](#), a ‘lexical analysis’ of a string to assign an empty set of entries to certain portions of a string, since *tacles* does not correspond to an obvious lexical form.)

9.3.6 *A class of puns*

Let us draw together the preceding discussion into a summary of the components of (a wide variety of) paradigmatic puns (including malapropisms). In [Appendix A.2.2](#) we present a formal definition which covers what we call *linguistically normal paradigmatic puns*. Less formally, they have the following characteristics:

- (a) part of the utterance is phonetically similar (perhaps identical) to some other string not present in the utterance;

- (b) either the utterance, or the utterance with that other string substituted in, is contextually appropriate;
- (c) if the two substrings are identical, then they should be lexically analysable in different ways, and the lexical analysis of the one not in the utterance should either be linked semantically to the context, or should involve grouping together words which are separate within the utterance;
- (d) if the two substrings are merely similar, then the unspoken one should form a complete and recognizable linguistic unit (e.g. a complete word or an established phrase).

This class is not identical to the entire set of paradigmatic puns, as we shall see in later sections.

A few observations are in order here. We have abstracted away from the various devices that could render two strings ‘similar’ (see [Section 3.8](#) and the discussion in [Section 9.3.4](#)), treating the overall comparison of the relevant strings as primitive (or as a deferred problem – see [Section 2.14](#)).

Our definition ensures that where a single phonetic form is involved, there must be two meanings (otherwise any word would be a pun), and also that the unspoken phrase is not merely a phonetic imitation of a substring of the utterance, but either is itself contextually linked or groups together separate units; without these constraints, simple lexical ambiguity would meet the definition given here.

9.3.7 *Knock-knock’ jokes*

One of the most rudimentary forms of joke is the knock-knock joke, such as (113).

- (113) Teller: Knock, knock!
 Audience: Who’s there?
 Teller: Noah.
 Audience: Noah who?
 Teller: Noah good place to eat?

(Ross 1998:10)

Sobkowiak (1991) classes these as a form of pun, which is a plausible judgement. Their structure is extremely simple. There are two similar strings, one of which is a person’s name (or could be thus interpreted), and the other is the opening substring of an arbitrary sentence.

A formal definition of knock-knock jokes is fairly trivial; see [Appendix A.2.3](#). A text qualifies as a knock-knock joke even with a most banal ‘punchline’, as in (113). The meaning of that sentence is irrelevant (although suitable content may make the joke funnier; see [Section 12.3](#)).

It is just possible to see these jokes as a special case of the paradigmatic pun, in the following way. The initial utterances establish a context, so that when the initial part of the punchline is uttered, it is similar to a contextually linked item, the name. As in some forms of paradigmatic pun, many knock-knock jokes use the device of matching the name against a sequence of several words: *noah/know a*. However, in contrast to a linguistically normal paradigmatic pun (as we have defined this), the final sentence has no notion of contextual appropriateness.

9.4

Syntagmatic puns

Our pun examples so far have all been paradigmatic puns: of the two strings being compared, one has appeared in the utterance and the other has not. (Even in the case of phonetic identity, it could be argued that only one lexical string has been uttered, even though its phonetic form coincided with that of the other lexical string in the punning pair.) Let us now turn to syntagmatic puns: texts in which both of the similar strings occur, as in (114).

- (114) A husband and wife are at a wedding reception. The husband says he needs to absent himself to perform a certain natural function'. The wife scolds him, 'You don't talk about that sort of function at this sort of function.'

(TV comedy drama, *A Bit of A Do*, UK;
also Nobbs (1986:33))

Here the lexical items *function*₁ (mechanism) and *function*₂ (social occasion) are both strongly linked to the context, in that they do indeed make sense within this utterance. This appears to be necessary: two similar words where only one fits in with the context would not constitute a pun. Indeed, it is hard to see how one could construct such a sentence, given the overall principle that the entire utterance should make sense in context, since a constituent in the sentence which was completely unsuitable would render the sentence infelicitous.

- (115) It is better to be looked over than to be overlooked.

(Mae West)

This also contains both the compared strings, *looked over*, *overlooked*, related by morpheme-level metathesis. It is not obvious what acts here as the context. Although the line was probably first delivered in a particular situation, the fact that the quotation has survived as a self-contained witticism suggests that there is very little outside the text (apart from general society or culture, including perhaps the attribution to Mae West) that is needed

to support the pun. It could perhaps be that the surrounding sentence itself acts as the context, with the two strings fitting fluently into this linguistic matrix and hence being contextually connected. Binsted's JAPE constructs (116), (117), and (118), all syntagmatic puns.

(116)What do you call a strange market?
A bizarre bazaar.

(Binsted 1996)

(117)What's the difference between leaves and a car?
One you brush and rake, the other you rush and brake.

(Binsted 1996)

(118)What's the difference between a pretty glove and a silent cat?
One's a cute mitten, the other's a mute kitten.

(Ertner (1993) cited by Binsted (1996))

Although (117) and (118) are conventionally classed as *spoonerisms*, this can be seen as a particular subclass of pun. The string *brush and rake* is phonetically similar to *rush and brake*, and the concepts mentioned in these phrases are linked to, respectively, the concepts mentioned in the set-up, *leaves* and *car*. Both phonetically similar phrases are present, which means that neither need be a fixed or established phrase or saying: it appears that being strongly cohesive (forming a complete word or a well-known phrase) is necessary only where the form is not actually uttered and so must be evoked indirectly.

Similar remarks apply to (119).

(119)What's the difference between a hairy dog and a painter?
One sheds his coat, the other coats his shed.

(Ertner (1993) cited by Binsted (1996))

Syntagmatic puns generally meet the constraint (discussed in [Section 9.3.5](#) above) of being semantically and pragmatically well formed, as there is only one utterance available, and it is compatible with the context.

In some cases, one of the occurrences of the matching words may be omitted if it is strongly implied by the structure of the sentence.

(120) It was a lovely wedding. The cake was in tiers ... so was the bride.

(Chic Murray)

In (120), the text contains an elliptical construction apparently short for *the bride was in tears*.

It would be quite easy, using our formal linguistic definitions ([Appendix A.1](#)) to state formally that a syntagmatic pun must meet the following conditions:

- the utterance is linguistically appropriate to the context;
- there are two or more (non-overlapping) substrings of the utterance which are phonetically similar and which *either* are not identical *or* have different meanings (or both).

The problem is that this formulation constitutes a *necessary* set of conditions, but not a *sufficient* condition for a syntagmatic pun: all syntagmatic puns will (we claim) meet these conditions, but some texts which meet these conditions will not be syntagmatic puns. Informally speaking, the formula given above says little more than that the text must contain two similar strings, and that recurrence of the same word (or phrase) with the same meaning does not suffice. The empirical deficiency of this can be seen by considering a few texts which conform to the requirements listed above:

- (121) The car was towing a caravan. (constructed item)
- (122) The committee will meet to discuss the meat quota. (constructed item)
- (123) The city schools teach a variety of skills. (constructed item)
- (124) I am interested in linguistic schools, you are interested in linguistic skills. (constructed item)
- (125) My car needed a repair, so I took it to a repair-shop. (constructed item)
- (126) My car had a dent, so I took it to a dentist. (constructed item)
- (127) Fox-hunting is ‘the unspeakable in full pursuit of the uneatable’. (Wilde 1966:437)
- (128) That’s one small step for a man, one giant leap for mankind. (Neil Armstrong, 1969)

Intuitions may vary, but it would be surprising if anyone regarded (121) as a pun, despite the occurrence of the two occurrences of *car*. Similarly, (122) does not seem to be a pun, though it is possible to imagine it being uttered with a flippant intonation which emphasizes the two occurrences of *meat/meet*, thereby offering it as a pun. It is hard to see (123) as a pun, although puns are often based on less similarity than that between *schools* and *skills*. In contrast, (124) seems to be close to pun-hood, perhaps because of the implied parallel between the two phonetically similar strings. Although (125) does not seem to be a pun, (126), with very similar formal structure, could be seen as a pun; in (126), the absurdity of the content may be a contributory factor. What is interesting about (128) is that the double occurrence of *man* (together with some parallelism of linguistic structure) seems to be an important factor in making the remark epigrammatic: compare the hypothetical paraphrase ‘I’ve not moved very far, but the human race has made a significant breakthrough’, which is not nearly as effective or memorable. Nevertheless, (128) is not usually described as a pun or a humorous quip.

What can be drawn from this is that there are further factors needed to qualify an utterance as a syntagmatic pun, but it is not obvious what these are, even informally.

9.5

Linguistically abnormal paradigmatic puns

There are certain puns which, while superficially similar to the punning riddles in [Section 9.3.4](#) above, are not easy to formalize in precise detail. (In (129) we have used approximate phonetic representation, for the reasons discussed in [Section 3.2](#)).

(129) When is a door not a door?

When it is / ɔjær/.

(130) Why is coffee like the soil?

It is ground.

(Pepicello and Green 1984)

(131) Why is a goose like an icicle?

Both grow down.

(Pepicello and Green 1984)

In (130) and (131), there is a question which compares two entities, and a punchline text which is ambiguous, with each meaning capable of describing one of the entities in the question. Pepicello and Green classify these two examples separately, on the grounds that different linguistic devices are used to produce the ambiguity. Nevertheless, both examples

seem to be puns. This subclass also includes some of the JAPE examples, such as (132).

- (132) How's a nice girl like a sugary bird?
They're both sweet chicks.

(Binsted 1996)

As in many earlier examples, the essential content of these jokes can be restated without using a question, as in (133), (134).

- (133) Sometimes a door is not a door. Sometimes, it is /əjær/.

(constructed item)

- (134) Coffee is like the soil. It is ground.

(constructed item)

From the preceding discussion, it might seem that these puns fit neatly into the subclass of paradigmatic puns discussed earlier (and that they meet the formal definition in [Appendix A.2.2](#)). However, closer inspection reveals a slight problem with the condition that (at least) one version of the text provides a completely coherent utterance with respect to the context. The complication is that the ambiguous string has two meanings, and *neither* of these meanings in isolation makes complete sense in context. In (130), if we interpret the punchline using the sense of *ground* which means roughly *soil* or *earth*, the answer is an untrue statement about coffee; if we adopt the other meaning, the passive participle of *grind*, then the punchline does not answer the question. This would render rather murky the notion of the punchline being 'linguistically appropriate'. Although either lexical form would produce a grammatically well-formed sentence (so our basic assumption of grammaticality is not threatened), it is hard to give a formally precise account of how the text is semantically or pragmatically coherent.

There is a similar difficulty with (131), emphasized by the fact that (130) could have had a punchline *Both are ground*, and (131) could have had a punchline *It grows down*. The use of *both* also happens in (132), with similar complications.

It is not easy to give an elegant characterization of the way in which such punchlines deviate from complete linguistic appropriateness. The following is a possible account for (130), relying on the fact that ambiguity (with phonetic identity) is involved. If there were a lexical item whose semantic content was the *disjunction* of the senses of the two homophonous words, then the utterance would be felicitous (semantically and pragmatically) providing that lexical meaning was associated with the string. That is, if we had a word *ground*₃ which meant '*either soil/earth or*

having had grinding applied to it', then (130) would be a perfectly correct discourse.

This process might seem promising for this example, but becomes more problematic for (131), where the ambiguity is not a single case of phonetic-lexical ambiguity, but consists of two ambiguous words forming distinct syntactic combinations.

If this phenomenon were confined to quips such as (130) and (131), we could have tried to develop a solution based on the idea that such texts are produced by some radically different process from normal language generation, and obey different structuring rules from conventional discourses. One objection to this is that in (135), although the entire text is a deliberately constructed joke, the actual utterance (the punchline) functions as a normal remark within the story, and could be seen as conveying an opinion from one story-character to another.

(135) An American soldier, serving in World War II, had just returned from several weeks of intense action on the German front lines. He had finally been granted R&R and was on a train bound for London. The train was very crowded, so the soldier walked the length of the train, looking for an empty seat. The only unoccupied seat was directly adjacent to a well dressed middle aged lady and was being used by her little dog. The war weary soldier asked, 'Please, ma'am, may I sit in that seat?' The English woman looked down her nose at the soldier, sniffed and said, 'You Americans. You are such a rude class of people. Can't you see my little Fifi is using that seat?' The soldier walked away, determined to find a place to rest, but after another trip down to the end of the train, found himself again facing the woman with the dog. Again he asked, 'Please, lady. May I sit there? I'm very tired.' The English woman wrinkled her nose and snorted, 'You Americans! Not only are you rude, you are also arrogant. Imagine!' The soldier didn't say anything else. He leaned over, picked up the little dog, tossed it out the window of the train and sat down in the empty seat. The woman shrieked and demanded that someone defend her and chastise the soldier. An English gentleman sitting across the aisle spoke up, 'You know, sir, you Americans do seem to have a penchant for doing the wrong thing. You eat holding the fork in the wrong hand. You drive your autos on the wrong side of the road. And now, sir, you've thrown the wrong bitch out of the window.'

(from www.jokes2000.com)

On first glance, this may seem to be a straightforward pun, based on the two meanings of *bitch*. (If *bitch* is not deemed ambiguous, analysis as a pun becomes even more difficult.) Suppose we choose the *unpleasant woman* sense of the word. Then *the wrong bitch* would mean *the wrong unpleasant woman*,

which is not exactly what the character in the story was saying (i.e. it is not linguistically appropriate). Similarly, the utterance does not mean *the wrong female canine animal*. (This seems to be part of a general semantic uncertainty about how to formalize the semantics of the adjective *wrong* when it is applied to an ambiguous head noun: can *the wrong bank* mean ‘the river bank instead of the financial bank’, and if so, how is this best represented formally?) As with (130), if we posit an abstract lexical item with surface form *bitch* and the meaning ‘either an unpleasant woman or a female canine animal’, then the application of *wrong* may well be orderly, and the punchline would then fit in linguistically.

Syntagmatic example (120) also bends linguistic rules in a similar fashion: if the final clause had been *the bride was in tears*, then the use of *so* would not be strictly correct, as it should conjoin two identical predications; if the final clause had been *the bride was in tiers*, then it might not have been a plausible or accurate statement. Intuitively, this difficulty contributes to the humour.

In the case of syntagmatic puns ([Section 9.4](#)) we observed that the empirical facts were not clear, even informally. Here, the range of data which constitute these ambiguity-based puns is quite clear (and some of this class of pun have been computer-generated). What is difficult is formalizing a general and elegant statement of the phenomena using standard linguistic devices.

Suppose we were to drop the condition of contextual appropriateness from our definition. The drawback would be that this would class as a pun any text which caused a contextually-linked ambiguous word or phrase to appear, as in (136) and (137).

(136) Why is coffee like the soil? When camping, the tent should be pitched on level ground.
(constructed item)

(137) Why is a goose like an icicle? Fungus had started to grow down the front of the castle.
(constructed item)

This does not seem desirable, as these do not feel like puns, and our aim is to define as exactly as possible the set of puns (and hence the set of non-puns).

Alternatively, we can view these jokes as embodying a form of faulty reasoning. Let us compare them with a similar non-pun, (138).

(138) Why is ebony like coal?
It is black.
(constructed item)

We can see that the validity of the answer to (138) depends on the reasoning: ebony is black; coal is black; therefore ebony is like coal. A similar deductive sequence can be set up for (130): coffee is ground; soil is ground; therefore coffee is like soil. The difference is that *black* in the first reasoning sequence denotes the same concept in both statements, whereas *ground* has two different denotations. That is, in these riddles, the pseudo-correctness of the answer seems to rely on reasoning which uses the *surface English words* rather than some more abstract semantic representation of the text. This is an undesirably complicated notion of linguistic suitability.

9.6

Miscellaneous awkward examples

Example (139) may be a borderline example of an orthographic pun (see [Section 9.7.4](#)).

(139)How can you spell *enemy* in 3 letters? N-M-E.

(Raskin 1987)

In (139), there is play with two mappings from written symbols to phonetic strings. There is the *pronunciation* mapping, which assigns a sequence of letters its normal sound as it would be within a word (e.g. *enmesh*, *gunmen*, *unmentionable*). There is also the *spelling* mapping which gives the phonetic names of the letters; *nme* would be roughly ‘enn-emm-ee’. The joke seems to involve using the *spelling* mapping as if it were the *pronunciation* mapping. Alternatively, one could say that the string *nme* and the string *enemy* (regarded as sequences of letters) are being treated as similar because the *spelling* form of one is very similar, phonetically, to the *pronunciation* form of the other. (139) is, informally speaking, a play on words (and letters) but does not fit naturally under our definition.

A similar complication arises with (140).

(140)Which word is always spelt wrongly?
Wrongly.

(Alexander 1997:23)

Suppose this is delivered phonetically. If the *wrongly* in the set-up is pronounced in the normal way, then it can be interpreted only as an adverb, not a description of a spelling; the latter would be something like *double-yoo-arr-oh-enn-gee-ell-wye*. Hence, the joke is not internally coherent, a fact which might well be reflected in an aggrieved reaction by someone on whom this was inflicted. If the joke is delivered in written form, it is slightly less badly formed, if we allow the printed form *wrongly* to stand for both use of the adverb and the sequence of letters *w-r-o-n-g-l-y*. This could perhaps be seen as an orthographic pun ([Section 9.7.4](#)).

These examples fall outside our analyses.

9.7 Other issues

9.7.1 *Linking the other string*

Even if we once again confine our attention to linguistically normal paradigmatic puns ([Section 9.3.6](#)), we can see that many puns have additional relationships between their components.

In a basic pun, one of the two similar strings may not be semantically linked to the context at all. In (93) and (97), the non-present text string is irrelevant to the context; in malapropisms (and examples like (88)), it is the text within the actual utterance which has no semantic link. However, some puns contrive to have both strings linked to the context, as in (98). Despite their formal similarities, (99) and (100) differ in this respect: (99) has both strings linked to the context, (100) has only one. This is not to say that both strings are linguistically appropriate to the context (though (99) does achieve this). In (98) (in which the strings are phonetically identical) *Cousins* forms part of the actual utterance, and the utterance makes sense in context, whereas *cousins* is merely connected to *brothers* by a vaguer thesaurus-like association.

If we think of linguistic appropriateness as *strong linkage* and the more loose associative connection as *weak linkage*, then there are (normally) three possible varieties of linkage combination: one strong linkage only ((93), (97), (100), (90)), one strong and one weak linkage ((98)), two strong linkages ((99)). Most examples seem to have a strong linkage involving the actual utterance, though in malapropisms and some other examples the linkage requires the substitution of the non-spoken item.

Syntagmatic puns will normally have two strong linkages, as both of the similar strings have to fit into the utterance, thus achieving linguistic appropriateness (e.g. (114)).

It is interesting to consider, from the point of view of linkages, the pun in (141) (variants of which appear in many collections of jokes, with attribution to various wits, including Ben Jonson and Samuel Johnson).

(141) A well-known wit claimed that he could make a pun on any subject. He was challenged to make a pun on 'the King'. He replied, 'The King is not a subject'.

Intuitively, this is a pun. The word *subject* has two readings, which we can indicate as *subject*_{citizen} and *subject*_{topic}. To have this conform to our definition, we can take the utterance to be *The King is not a subject*_{citizen}. Although this is a true statement, it could be argued not to be linguistically appropriate to the context. However, with *subject*_{topic} substituted, the utterance is linguistically appropriate, as an objection (if a dubious one) to the challenge. Also *subject*_{citizen} is associated (weakly linked) to the context. There is

the additional link that by meeting the requirements for a pun, the utterance has a further link to the context (since puns are under discussion). Indeed, since the challenge is to produce a pun, once the utterance has qualified as a pun, it is arguably doubly linguistically appropriate. Hence, this joke has *three* linkages, with possibly two of them being strong links.

Lippman and Dunn (1995) offered the hypothesis that the more connections in a pun, the better it would be valued by human subjects. Their experiments, with the ‘Tom Swifty’ genre of joke (see [Section 10.3](#) below), seemed to support this, but the range of data, and the notion of ‘connection’ were both very limited. (See also Lippman and Dunn (2000).)

9.7.2

Intricacy of the match

It is tempting to suggest that the amount of contextual linkage (as outlined in [section 9.7.1](#) above) should be the main or only factor in determining the humorous quality of a pun, particularly as some of the less strongly linked examples ((92), (93)) seem so poor. However, there are examples which seem to suggest that the situation is not so simple. (97) for example, has only one strong linkage, the minimum for a pun. However, it is clear that many people were amused by this pun: it is a rare example of a pun which attracted news coverage in its own right, being quoted and discussed widely in the UK news media after it was published. Similarly, some of the story puns in Muir and Norden (1991) are very impressive, and produced copious laughter and applause when first broadcast, despite having (as with most story puns) just the one basic linkage. What may be happening in these examples is that the audience are impressed by the sheer ingenuity of the construction.

It would not be trivial to define a formal measure of the intricacy of a pun. A variety of factors might be relevant, such as the length of the strings compared, or the apparent irrelevance of the matched string (as in (97)).

9.7.3

Puns versus forced reinterpretation

Many people would class (142) as a pun, given that it is a play on the possible meanings of *spirit*.

(142) What do you call a person who puts you in touch with the spirit world?

A bartender.

(Dienhart 1999)

Example (28) earlier is similar, and so, arguably, is (6). These could be fitted within the scope of our definition, if we accept that some or all of the (textual) context comes *after* the ambiguous text; some of the set-up text could also count as context. That is, *bartender* acts as the context, and the (preceding) word *spirit*, or perhaps some longer phrase within *a person who puts you in touch with the spirit world*, is linked to that context. The other meaning of *a person who puts you in touch with the spirit world* is not linked to this context.

However, there are further differences between these jokes and those outlined in [Section 9.3.6](#):

- (a) These jokes depend on phonetic identity (usually phonetic-lexical ambiguity), with similarity not being sufficient.
- (b) The part of the context which occurs in the punchline (after the ambiguous word) is linked to only one of the interpretations; that is, there is a partitioning of the context text into ‘before’ and ‘after’ sections, with systematically different relationships to the two word-meanings.
- (c) It is central to the working of the joke that one of the meanings is less obvious than the other, based on the ‘before’ context.
- (d) The punning text falls in the set-up rather than the punchline.

Moreover, these examples can be analysed in FR terms, as in Chapters [5](#) and [8](#): there are two interpretations of the set-up, with the less obvious being revealed by the punchline. So how should such examples be categorized? This may appear to be a dilemma, but only if we adopt the theoretical aim of partitioning jokes into *disjoint* classes within a strictly hierarchical taxonomy. That is not our goal. Rather, we are looking for conditions which, when applied to linguistic forms which are otherwise well formed, indicate whether the form is a joke or not. If certain texts, such as (142), meet more than one set of conditions, that is not problematic. Nevertheless, here we will keep these two classes separate, at the risk of alienating those who might insist that any joke based on phonetic-lexical ambiguity (such as (6), (28), (142)) is a pun. We do not see it as crucial to frame a single definition that accounts for every example of the informal notion of ‘pun’. Rather, we intend to develop a set of class-definitions which, taken together, account for jokes, even if they do not correspond exactly to any pre-theoretic informal classification. We take the sketch of FR jokes given in Chapters [5](#) and [8](#) as defining one way of structuring jokes, and our definition of linguistically normal paradigmatic puns as defining another. It does not matter if jokes classed by other writers as ‘puns’ are spread across the two classes. The reason for maintaining this separation is a belief that our two definitions (of FR jokes and puns) characterize different possible joke structures, operating in distinct ways

and with different properties, and with different internal structures at the level of abstraction we have chosen.

9.7.4

Puns using orthography or other media

We have so far talked of puns in terms of strings of symbols. The assumption is that these would be phonetic units, as puns are generally based on the sounds of word and phrases. Given the discussion in [Chapter 3](#) about the possibility of using written representations of jokes, it is natural to consider whether there could be puns based not on the pronunciation of words but on their orthography. That is, is it feasible to have puns in which the symbol sequences under consideration are written, and the notion of ‘similar’ is defined in terms of similarity of written form (even though the spoken form might differ)?

(143) The medieval schoolmaster was determined to tame his unruly pupil. Therefore he threatened to flog the youth within an inch of his young life if he didn’t measure up and do his work properly. He added, ‘Thou hast better study thy lesson, Master Jonathan, or else methinks thou wouldst do well to get a weapon and defend thyself.’ He bowed stiffly and prepared to leave. Then the pupil made a *bow*.

(Heller 1974:274)

Heller offers (143) as a pun based on the two words whose written form is *bow* (saying that such puns ‘*cannot* be transmitted through a different medium’), but it is unclear whether this was a joke he found ‘naturally occurring’ (i.e. in some text or in an attested joke-telling) or whether it was constructed specially. Binsted also raises this issue, offering a constructed example, (144), to illustrate the point, which she says involves *pronunciation ambiguity*.

(144) What do you call a blizzard forecast on April 1st?
A wind-up.

(Binsted 1996)

A naturally occurring (syntagmatic) orthographic pun is (145), since *science* is not pronounced in the same way in its two occurrences.

(145) Processed cheese represents the triumph of science over conscience.

(Alexander 1997:36)

Such jokes seem to be extremely rare: the most widespread form of pun is firmly rooted in phonetic similarity.

Although Heller and Attardo (see [Section 9.2](#) above) both allude to puns in other media (such as gestures), they do not explore this in detail, and Heller has a wider notion of pun in any case.

Alexander (1997:21) suggests that ‘The Mouse’s Tale’ in Lewis Carroll’s *Alice’s Adventures in Wonderland*, where the text is typographically laid out in the shape of a mouse’s tail, is a ‘visual pun’. However, he concedes that it starts from the homophones *tale* and *tail*. In fact this homophonic relationship is crucial: this would not be a pun in a story told in a language where there is no such phonetic relationship. This seems to be a subtly different genre of joke, namely the manipulation of the *context* to set up a potential (perhaps unspoken) pun. That is, non-verbal circumstances are contrived which allow a verbal description which follows the pattern of a paradigmatic pun. Another example of a visually presented joke which is parasitic upon the verbal notion of the pun occurs in the 1963 film *It’s a Mad, Mad, Mad, Mad World*. A character is dying, rather theatrically, on a mountainside. As he expires, his leg jerks and involuntarily kicks a bucket. Here, the *context* has been manipulated to set up a situation in which a paradigmatic pun (*he’s kicked the bucket*) is *possible*, though it is not uttered. The juxtaposition of actions is presumably recognizable as a joke opportunity only to those who are familiar with the idiom *kick the bucket* (meaning *die*). Thus we have a stimulus which is entirely non-verbal, but where linguistic knowledge is critical to recognizing the humour, and the overall structure of the joke is essentially that of a linguistically normal paradigmatic pun (albeit unspoken).

The term *visual pun* should probably be reserved for a visual stimulus which has analogous properties to those which define a verbal stimulus as being a pun. The simplest case would be a visual item with more than one interpretation, each of which was associated with, or appropriate to, the context. An example (Geraint Wiggins, personal communication) is the logo used by the catering services on certain UK trains: it is interpretable either as a fork or as a section of rail tracks. Another subclass of pun which is related to written form (punctuation rather than orthography) is exemplified by (146).

(146) What makes men mean?

The letter a.

(Alexander 1997:22)

This could be viewed as a (double) paradigmatic pun, where the strings being compared are *men*/‘*men*’ and also *mean*/‘*mean*’. That is, the *use–mention* distinction is at the root of this pun. If the question were punctuated as *What makes ‘men’ ‘mean’?* the question would be a completely correct answer.

9.7.5 *Bilingual puns*

Sometimes a pun crosses the boundaries between two languages: while made partly or wholly in one language, it depends on phonetic similarity to a phrase in another language. The epitaph (147) is an example.

(147) Here lies Willie Longbottom, Aged 6. *Ars longa, vita brevis*

(Nash 1985:145)

This depends on knowledge of both English and Latin for its effect. Alexander (1997:35) cites some jokes which depend on cross-dialect similarity between phonetic strings, though in our taxonomy they would count as forced reinterpretation jokes (see [Section 9.7.3](#)). There used to be (perhaps still is) a fringe theatre group which called itself *Lumière and Son* (cf. *son et lumière*). This uses cross-language orthographic identity (English *son*/French *son*), cross-language semantic identity (English *and*/French *et*), and ordinary phonetic identity (*lumière*), as well as transposition.

It may be possible to generalize the definition in [Section 9.3.6](#) and [Appendix A.2.2](#) to cover cross-language or cross-dialect jokes, but we will not attempt this non-trivial task here.

9.7.6 *Multi-way puns*

(148) ‘I’m bored,’ yawned the piece of wood with a hole in it.

(Venour 1999:38)

Venour offers (148) as a (specially constructed) three-way pun, based on two meanings of *bored* and one of *board*.

(149) A wife is cleaning out a grate; she pulls the metal tray out and it makes a very loud and unpleasant sound; her husband exclaims ‘What a great noise!’

(Wright 1979:410)

Wright constructs (149), in which three ‘context indicators’ (fireplace, noisy friction, loudness) allow three readings of *great/grate*.

It might be possible to generalize our definition of linguistically normal paradigmatic puns ([Section 9.3.6](#) and [Appendix A.2.2](#)) to cover n-way puns, but there is not much relevant data to guide the wording of the extended definition.

9.7.7 *Beyond structure*

The patterns we have sketched here for some classes of pun are purely structural, in keeping with our aim of investigating the ‘abstract syntax’ of jokes. Although an arrangement such as that defined in [Section 9.3.6](#) indicates when a text is a pun (as opposed to some other form of joke), it may not cover all the factors relevant to making the text a joke, or making it funny. We remarked above that misuse of a word might be structurally a malapropism without being seen as funny, and that syntagmatic puns require more than just matching strings. It may well be that some, or all, puns require some further (semantic or pragmatic) component(s) to render them humorous, such as INAPPROPRIATENESS or CONTRAST, just as we have suggested that propositional jokes require these ingredients. Although we have not explored this issue here, it may be that successful puns have both the relevant structural properties and the right kind of semantic content; bad puns are perhaps those which merely meet the weakest structural requirements, but have little else to say (cf. Raskin (1985:141–2)).

9.8 Computational perspectives

As discussed in [Chapter 1](#), we regard computer testing as a useful way of checking the detail, precision, completeness, and correctness of symbolic rules. The definition framed in [Appendix A.2.2](#) may not be fully ready for computer implementation, but it is a step in that direction. It could be implemented as a procedure which tested whether a text contained a pun, if formal (and computationally tractable) definitions could be given of the following notions: *linguistically appropriate*, *context*, *concept*, *linked*, *similar*. As discussed in earlier chapters, one of the purposes of the outline definitions given here is to set out possible research subproblems, often in the area of linguistics or knowledge representation, whose solution (or partial solution) would contribute to a more articulated definition of humour, or more specifically of joke classes. So in a sense these undefined primitive notions should not be regarded solely as defects in the overall class definition; they can also be viewed as a decomposition of the problem.

The definition could also, with the same caveats about the primitive notions, be used in a computer program to generate puns, although there are various non-trivial decisions to be made about integrating the pun-guidance with normal natural language generation.

It is possible to construct a program which is devoted to producing self-contained puns (i.e. texts which contain their own context). However, the pun-generating programs of recent years (see [Chapter 10](#)) often have their pun-computations attached to, or embedded in, an *ad hoc* natural language generator, so that the relationship of the pun devices to general

text generation is obscured. More challenging, and more interesting, would be the construction of an opportunistic computational pun-maker, in which a general-purpose text-generator intended for some other task (e.g. as a friendly user interface for an instructional system) was able to use the general definition of a linguistically normal paradigmatic pun to guide its choice of phrasing and lexical items, so that it managed to produce contextually relevant utterances which also contained puns.

Takizawa *et al.* (1996) have implemented a pun detecting program for Japanese, which accepts a sequence of phonemic symbols and produces possible analyses of this in terms of sequences of Japanese words, rating each word-sequence with the likelihood that it is a pun, based on various heuristics. This handles syntagmatic puns, though it appears that in Japanese puns the two compared strings may overlap in the text.

9.9 Summary

We have considered a range of puns, both spontaneously occurring in everyday conversation, and passed on in 'canned' jokes. Generalizing from the regularities in these texts, we have given a definition of the linguistic structure of a subclass of paradigmatic puns. We do not have an equally detailed account of certain other paradigmatic puns, nor of syntagmatic puns, and we have chosen to class certain ambiguity-based examples as forced reinterpretation jokes rather than puns. With these exceptions, we believe that our definition captures fairly accurately the structural aspects of a traditional variety of 'pun'. As noted, these structural features may not be sufficient to form a humorous text, and further aspects may need to be present.

10

Some computational studies

We review recent work on computer modelling of verbal humour, and draw some conclusions about this field.

10.1

Introduction

In Sections [2.12](#) and [2.13](#), we discussed the possible role for, and ways of approaching, the computer generation of humorous texts. Since the early 1990s, there has been a growing interest in this form of research, and a number of small humour-generation programs have been implemented.

In this chapter, we shall review the following systems: the LIBJOG system (Raskin and Attardo 1994), the Tom Swifty generator (Lessard and Levison 1992, 1997), a simple riddle generator (Lessard and Levison 1993), the JAPE riddle generator (Binsted and Ritchie 1994, Binsted 1996, Binsted and Ritchie 1997, Binsted *et al.* 1997), the HCPP generator (Venour 1999, Lessard *et al.* 2002), the WisCraic pun generator (McKay 2000, 2002), and the HAHAAcronym acronym manipulator (Stock and Strapparava 2002, Attardo and Mele 2002).

After summarizing these systems, we will discuss some of the general issues raised by this style of work.

10.2

The LIBJOG joke generator

10.2.1

Behaviour

Raskin and Attardo (1994) discuss their LIBJOG series of programs (said to have started around 1988), which constructed light bulb jokes such as (36) and (150).

(150) How many Californians does it take to screw in a light bulb?

Two. One to screw it in and one to share in the experience.

(Raskin and Attardo 1994)

10.2.2 *Mechanisms*

The program uses a textual template with five slots: for a group name (e.g. *Poles*), for two numbers, and for two activities (e.g. *hold the light bulb*, *turn the table he is standing on*). The fillers for these slots are computed using specially constructed tables of information. In the first version of LIBJOG, this information consisted of a list of five-tuples, each tuple representing a valid set of text string fillers. Later versions used more indirection, involving tables of characteristic traits of named groups, but how the actual textual elements were computed from this is not specified.

Attardo and Raskin do not summarize their results or mention any evaluation.

10.3 The Tom Swifty generator

10.3.1 *Behaviour*

Lessard and Levison describe a program which, using an existing natural language generator, VINCI (Levison and Lessard 1992), produces a simple variety of pun, the *Tom Swifty*. This class of pun is exemplified by (151) and (152).

(151) 'Turn up the heat,' said Tom coldly.

(Lessard and Levison 1992)

(152) 'I'll have the lamb,' said Tom sheepishly.

(Lessard and Levison 1992)

That is, each example consists of a quoted utterance, and an attribution of this remark to Tom, with an adverb attached. There is some semantic link between the adverb and the meaning of the utterance.

10.3.2 *Mechanisms*

Lessard and Levison model these puns by identifying three parts of the text, and two relationships between these items. The *pivot* is the adverb (e.g. *coldly*). The *base* is some subpart of the adverb which is semantically related to the utterance (e.g. *cold*). The *target* is some subpart of the utterance which is linked semantically to the base (e.g. *heat*). The *formal bridge* is the relation between base and pivot (e.g. the creation of an *-ly* adverb from an adjective), and the *semantic bridge* is the relation between base and target (e.g. antonymy).

The VINCI program can generate normal English text, and so provides the syntactic structure for the pun. The Swifty generator operates by

imposing further constraints on the form of the text that VINCI is to generate, choosing from a finite list of possible values for the two bridge relations, and from fairly conventional vocabulary for the base, pivot and target.

There were various versions of the system. One variant restricted adjectives to those which could form *-ly* adverbs. This produced results in which the utterance was of the form *This NOUN is ADJECTIVE*, such as (153) and (154).

(153) 'This matter is opaque,' said Tom obscurely.

(Lessard and Levison 1992)

(154) 'This object is dull,' said Tom pointedly.

(Lessard and Levison 1992)

A more complex version of the system used lexical pointers between related words to implement the formal bridge, for example *boy* ¡ú *boyish* ¡ú *boyishly*.

Hence the semantic bridge could be from the target to a noun as base, with the required adverb (pivot) being formed by following links. This allows examples such as (155) and (156).

(155) 'I am not a nun,' said Tom monkishly.

(Lessard and Levison 1992)

(156) 'I am not a flirt,' said Tom coquettishly.

(Lessard and Levison 1992)

A further version operated by searching, from a noun base, for an adverb which is linked in some relatively indirect way, such as partial orthographic similarity, as in *pit/pithily*. This was implemented as a number of specially crafted lexicon mechanisms, and produced results such as (157) and (158).

(157) 'I stare at the flab,' said Tom fatuously.

(Lessard and Levison 1992)

(158) 'I notice the kitten,' said Tom categorically.

(Lessard and Levison 1992)

Although the syntactic forms of all the examples were generated in a relatively general manner, the lexicons were typically small and specially constructed. No evaluation of the output was carried out and Lessard and Levison give no indication of the overall quality of their results.

10.4 A simple riddle generator

10.4.1 *Behaviour*

Lessard and Levison (1993) sketch the workings of a program which, again using the VINCI generator, produces some basic forms of punning riddle, such as (159).

(159)What do you call a naked bruin?
A bare bear.

(Lessard and Levison 1993)

The program also produces conundrums such as (160).

(160)What has an eye and can't see?
A potato.

(Lessard and Levison 1993)

and riddles based on substring matching between words, as in (161).

(161)What kind of animal rides a catamaran?
A cat.

(Lessard and Levison 1993)

10.4.2 *Mechanisms*

All their mechanisms are based on special-purpose syntactic/semantic forms programmed in the VINCI notation, coupled with suitably prepared lexical data (e.g. obtained by a pattern-matching search of a machine-readable dictionary). A process which created (159) finds a noun (e.g. *bear*) having both an adjectival homonym (e.g. *bare*) and a synonym (e.g. *bruin*), then finds a synonym (e.g. *naked*) of the homonym. All these are then fitted into a syntactic frame to complete the text. In (160) the program is required to select a noun (e.g. *eye*) which names a body part, has a noun homonym, and has a role (e.g. *see*). Assuming the homonym is also a part of, or possessed by, some sort of object (e.g. *potato*), that object is computed. All these are then fitted into a syntactic frame.

Lessard and Levison give no summary of their results, and did not perform a full evaluation.

10.5 The JAPE riddle generator

10.5.1 *Behaviour*

The JAPE program (Binsted 1996) generated certain classes of punning riddles. Some of the better examples were the following.

- (132) How is a nice girl like a sugary bird?
Each is a sweet chick.
- (117) What's the difference between leaves and a car?
One you brush and rake, the other you rush and brake.
- (118) What is the difference between a pretty glove and a silent cat?
One is a cute mitten, the other is a mute kitten.
- (116) What do you call a strange market?
A bizarre bazaar.

10.5.2 *Mechanisms*

The relevant facts about a JAPE joke can be summarized semi-formally using various properties and relationships. This is most easily demonstrated with a simple example.

- (102) What kind of murderer has fibre?
A cereal killer.

Example (102) can be analysed thus (cf. Lessard and Levison's program, [Section 10.4](#) above):

There is a compound noun phrase (*serial killer*) such that its first word has a homophone (*cereal*). Substituting the homophone into the phrase (*cereal killer*) produces a phrase which can be used as the answer to a question which asks for something which shares the semantic properties of the original phrase (*serial killer*) and the homophone (*cereal*).

An abstraction of such a description was termed a *schema*, and it could be thought of informally as a network of constraints between words, parts of words, phrases, and lexical entries.

A schema also includes information about the message that the text must convey about these lexical items. For the schema outlined above, the specification could be loosely glossed as:

Create a descriptive phrase using the properties of the original noun phrase and of the substituted homophone, and arrange the text to say that this phrase means roughly the same as the phrase formed by substitution.

The JAPE procedure is to search the dictionary for words and phrases which satisfy a schema, thereby *instantiating* the schema. This also automatically supplies values for the textual specification portion of the schema, which can then be passed to the text-generation modules. These text modules have no humour-creating role, but are just a (slightly *ad hoc*) mechanism for carrying out this textual request. The text-generation rules contain a number of options, so there may in general be more than one way of constructing surface text which meets the specification supplied by the schema. Hence the same schema, and the same lexemes, can give rise to slightly different surface forms.

The unconstrained nature of the text generator means that rules could be written which produce text in direct response to arbitrarily chosen symbols provided from the schema; that is, a schema could make a very precise decision about the surface text, and dictate that decision to the generator. Although the intended division of labour is that the schema defines high-level abstract relations and the text generator ensures that these are faithfully realized in low-level concrete forms, it would be possible for the schema to pre-emptively choose some aspect of the surface text.

This method was used for jokes such as:

(110) What do you call a depressed train?

A low-comotive.

Here the schema specifies that there must exist lexemes (which in this example would be *low*, *locomotive*) with corresponding surface forms and subparts (here, *lo*, *comotive*, *low*). Then the answer part of the riddle must be exactly the concatenation of two of these surface texts (resulting in *low*, *comotive*); the answer text is (realistically) not defined in terms of abstract lexemes.

This emphasizes the unrestrained nature of the JAPE architecture. It is more like a library of routines which, used together, allow one to program up riddles. There is no suggestion that actual *formalisms* used within JAPE (particularly the text generation methods) form a substantive theory of how riddle texts should be produced.

(Ritchie (2003) gives full technical details of the working of JAPE.)

10.5.3 Evaluation

The output of JAPE was tested in a controlled fashion (Binsted *et al.* 1997). Texts produced by the program (using the online dictionary WordNet (Miller *et al.* 1990)) were filtered according to objective criteria (e.g. use of concrete vocabulary) based on the MRC psycholinguistic database (Coltheart 1981). A random selection of these were mixed with non-joke items of two classes: sensible question–answer pairs (e.g. *What do you get when you cross a horse with a donkey? A mule.*) and nonsense question–answer pairs (e.g. *What do you get when you cross a murderer and a ferry? A citrus fruit.*). Proportions of the different classes of item, and the order of presentation of items, were carefully controlled, being randomized where appropriate.

These data were shown to children aged eight to eleven, who were asked to state whether a text was a joke or not, how funny the text was on a five-point scale, and whether they had heard it before. The items were presented simultaneously in printed form and through headphones, read in a neutral intonation by an actor. There were 122 subjects, and each text item was rated by at least nine children, with most items getting twelve ratings.

Analysis of the results showed certain statistically significant patterns regarding the status of a text as a joke: human-generated jokes were reliably distinguished from non-jokes; JAPE-generated jokes were reliably distinguished from non-jokes; human-generated jokes were given a higher average rating than JAPE-generated jokes (as being jokes). Similar patterns emerged regarding funniness: human-generated jokes were funnier than non-jokes; JAPE-generated jokes were funnier than non-jokes; human-generated jokes were funnier than JAPE-generated jokes. There were other findings, based on a *readability* measure computed from the MRC psycholinguistic database: human-generated jokes were more readable than JAPE-generated jokes; human-generated jokes and JAPE-generated jokes put together show correlations between readability and both jokehood and funniness.

10.6 The HCPP generator

10.6.1 Behaviour

The Homonym Common Phrase Pun (HCPP) generator (Venour 1999) can generate simple texts consisting of a one-sentence set-up and a punchline consisting of a head noun preceded by an adjective or a noun modifier, sometimes with an indefinite article. Examples include:

(162) The surgeon digs a garden. A doc yard.

(Venour 1999)

(163) Joan hears wailing in the booth. Carrel singing.

(Venour 1999)

(164) Joan charms a man in the sack. A male bag.

(Venour 1999)

Venour lists 50 jokes produced by the HCPP generator.

10.6.2 *Mechanisms*

Venour used *schemata* (similar to those in JAPE) to specify the required lexical items and their interrelations. He defined ten schemata, but implemented only three. The examples quoted above represent one from each schema. The central idea is the *common phrase*, which is some well-established two-word phrase such as *raising cattle*, *down payment*, *dock yard*, *carol singing*, *mail bag*. These are treated as lexical items. There is a single node in a schema representing the phrase, but it has links to its two parts. There are then links from these parts specifying conditions on these items or on the linked items. In all the schemata, one of the phrase-parts is linked to a homonym or homophone.

Each schema also had a single surface *template* (again similar to the templates of JAPE) indicating where the lexical items are to be fitted into a skeleton text. Example (165) is generated from a different schema-template pair from that used for (162).

(165) Joan visits a surgeon in the garden. A doc yard.

(Venour 1999)

Each schema and its template are coded up together in a single piece of program written in the notation used by the VINCI text generator (Levison and Lessard 1992), so that the instantiation of these structural patterns is then carried out by the VINCI system. There are not separate data structures for schema and template, and the coding is very dependent on the way that the VINCI lexicon is represented.

10.6.3 *Evaluation*

With a lexicon of 36 common phrases and 240 words, the program produced, for each of the three schemata, 34, 7 and 9 jokes. Groups of 25 output items were shown to 16 subjects, who were asked to supply grades on a five-point scale (borrowed from an early study of JAPE (Binsted and Ritchie 1994, 1997)):

- 1 Not a joke. Does not make sense.
- 2 Recognizably a joke but a pathetic one.
- 3 OK. A joke you might tell a child.
- 4 Quite good.
- 5 Really good.

The average score was 2.81, with 21 jokes having average scores between 2.0 and 2.9 inclusive, but 22 jokes were awarded scores of 3, 4 or 5. The highest average score (4.4) went to (166).

(166) The sailor bears a stress. Pier pressure.

(Venour 1999)

No non-HCPP control items were included in the questionnaires, and only one question was asked. Venour does not explain how the groups of 25 items were selected or ordered.

10.7

The WisCraic pun generator

10.7.1

Behaviour

The WisCraic program (McKay 2000, 2002) could produce simple puns either in question–answer form (e.g. (167)), as a single sentence (e.g. (168)), or as a two-sentence sequence (e.g. (169)).

(167) Who broke the woman's hart?
The cruel deer-keeper.

(McKay 2002)

(168) The performing lumberjack took a bough.

(McKay 2002)

(169) Your mate Johnny is a hard-up deer-keeper. He really needs doe!

(McKay 2002)

The system could also give a brief explanation of the workings of a pun which it had created; McKay suggests that this could form part of software to support second-language learning. An example output is as follows:

The friendly gardener had thyme for the woman!
The word time, which is part of the idiom
[have, time, for, someone]
is a homonym of the word thyme.
A HOMONYM is a word that sounds like another word.

[LINK] between thyme and gardener:
 |thyme is a type of plant
 |a gardener works with plants
 friendly', which is associated with the idiom
 [have, time, for, someone]
 was selected from other adjectives as it has the highest
 imagability score:
 439

10.7.2 *Mechanisms*

The WisCraic program systematically constructs a structure which ensures that the necessary constraints are met, roughly as follows:

- (a) Select an idiom or well-established verb phrase from an available list.
- (b) Find a homonym for one of the words in the phrase.
- (c) Substitute that homonym into the phrase.
- (d) Find a noun which is a compatible subject for the verb phrase and which is semantically associated with the homonym which has been inserted in the phrase.
- (e) Make this noun the head of the subject.
- (f) Find an adjective which is semantically associated with the original idiom/common phrase.
- (g) Insert this word in the subject noun phrase.

To allow this computation, the program had the following knowledge bases:

- Idioms: a set of common verb phrases, each with a set of associated adjectives.
- Professions: a set of occupations, each with semantically related vocabulary.
- Dictionary: general set of nouns, verbs, adjectives, with phonetic information and various psycholinguistic parameters.
- Lexicon: links between words, for use in generating explanations (not jokes).
- Grammar: various syntactic information needed for sentence construction and substitution.

10.7.3 Evaluation

The WisCraic output (21 types, each appearing a number of times to produce 76 tokens) was tested by 15 subjects. Each subject was presented with ten items to judge: six WisCraic-generated and four control items. The controls consisted of one text from each of the following types:

- A sentence using the idioms/common phrases normally, e.g. *The mysterious man disappeared into thin air!*
- A WisCraic joke but with the adjective omitted from the subject phrase, e.g. *The lumberjack took a bough!*
- Uses of the idiom with homonym substituted but no other links, e.g. *The strong policeman met the woman half whey!*
- A sentence which meets the WisCraic formula except for the homonym substitution (which would be possible), e.g. *The strong fisherman showed the woman his muscles!* (These were termed *suggested jokes*.)

Subjects were asked to judge texts on three measures: being a joke (yes/no), funniness rating (0 to 4), and cleverness rating (0 to 4). Where an item was marked as not being a joke, the subject had to allocate it to *one* of four classes: uses unknown words; an ordinary sentence; does not make sense; a failed attempt at a joke.

McKay lists three hypotheses as having been confirmed: WisCraic output are jokes (84 per cent of WisCraic items were classed as jokes, and 10 per cent as failed attempts at jokes); cleverness correlates with funniness (of the (five) jokes rated five times or more, 80 per cent had the same funniness and cleverness ratings); 'suggested' jokes are weak jokes (over 40 per cent of the fourth class of control item were classed as jokes, and had below average funniness ratings). Two hypotheses were not confirmed: WisCraic items with adjectives removed are less funny, and normal use of idioms are judged as ordinary sentences. McKay does not give overall funniness or cleverness ratings for the complete set of WisCraic items, but provides bar charts which indicate that most of the output texts were rated around 2 on both the scales.

Although the evaluation findings (including comments by the subjects, reported in McKay (2000)) are interesting, the study is on a very small scale and has some methodological weaknesses. In particular, the average number of times each joke was rated was 3.6, with one joke rated only once and another rated ten times. McKay does not provide measures of significance for his results, and the suspicion must be that some of the effects are not statistically strong. However, the bald fact that 84 per cent of the WisCraic items were rated as jokes is interesting.

10.8 HAHAcronym

10.8.1 *Behaviour*

The HAHAcronym program (Stock and Strapparava 2002) is capable of two separate computations:

- (a) Given an abbreviation based on the initials of words (such as *MIT*, for *Massachusetts Institute of Technology*), it can generate a phrase with the same initials and with an amusing meaning (such as *Mythical Institute of Theology*).
- (b) Given two English words (or abbreviated phrases) (e.g. *fast*, *processor*), it will generate an acronym (i.e. a phrase in which the initial letters of the words spell out another word) loosely describing the two concepts expressed by the two words (e.g. *INERT: Inconclusive Non_parallel Electronic equipment for Rapid Toggle*).

10.8.2 *Mechanisms*

In order to rephrase existing abbreviations, the system carries out the following steps:

- Parse the phrase. A conventional syntax tree is constructed for the full form of the given abbreviation, using a simple formal grammar.
- Select terminal nodes for alteration. Usually the syntactic head of the phrase is retained unaltered, and other words such as adjectives are available for mutation.
- Find and perform possible substitutions. The revised phrase should retain the same initial letters, and its prosody (rhythm) should be similar to the original; the CMU pronouncing dictionary (www.speech.cs.cmu.edu/cgi-bin/cmudict) was used for this.

The substitution of words is guided by the following knowledge sources:

- Semantic field oppositions. Entries in the WordNet online dictionary (Miller *et al.* 1990) were annotated with *domain labels*, which in turn were organized into a hierarchy. A separate table listed pairs of domains which were ‘opposed’, such as *religion* and *sex*. In general, replacing a word with one from an opposing domain was preferred.
- Antonym sets. Information from WordNet about which adjectives were antonymous (e.g. *fast/slow*) was helpful.
- A-semantic dictionary. A list was put together of adjectives and adverbs which are generally useful for the acronym task, such as *abstrusely*, *exorbitantly*, *weirdly*. These were used if nothing better could be found.

Generation of a novel acronym from a pair of key words involves:

- Select a candidate word. A word which is in a funny or 'ironic' relationship to the initial meaning (s) is chosen.
- Generate a phrase. A phrase whose initial letters match the chosen word has to be generated, trying to satisfy various constraints such as staying close to the domain (semantic field) of the input meanings, and having systematic semantic relations (e.g. antonymy) to the inputs.

10.8.3 Evaluation

Both the acronym-reanalysis and the acronym-generation functionalities were tested.

Acronym-reanalysis was tested by giving 20 subjects a set of items consisting of a genuine acronym, its correct expansion, and a small number of possible distorted expansions. The subjects had to rate each expansion for funniness on a five-point scale, allowing a maximum possible funniness score, across all subjects, of 100, i.e. 5 points from each of 20 subjects. The mean score was 59.61 (*mildly funny*), a standard deviation of 8.91, minimum 40, maximum 80.

A set of 87 generated acronyms were presented, along with the two words which were input to the generation, to 20 subjects (undergraduate students), who were asked to rate the acronym for being 'funny/ironical' (1–5). Subjects were told these were inventions of an advertising agency as names for programs. The mean score was 56.37 (above *not very funny* and just below *mildly funny*), with a standard deviation of 8.01, a minimum of 40 and a maximum of 75.

As a control, the reanalysis program was run without the semantic filters and heuristics, thereby creating syntactically valid but otherwise random acronyms. These were rated for funniness by a different set of 20 subjects, resulting in a mean of 42.05 (*not very funny*), with a standard deviation of 9.76, minimum 27, maximum 67. This indicated that the semantically controlled reanalyser performed statistically significantly better than chance.

10.9 Discussion

Some generalizations can be made about all the systems listed above:

- (a) The coverage is extremely narrow, with no pretensions to general joke-creation.
- (b) The mechanisms make use of lexical information, relying on the fact that the chosen genres do not require vast amounts of world knowledge or reasoning.

- (c) All the programs are generators of humorous text rather than interpreters which analyse jokes. Although the written accounts strive to present the humour knowledge in a relatively non-procedural manner, it is not clear how the procedural parts could be ‘run backwards’ to ‘parse’ jokes. In some ways, this is no worse than some linguistic theories, past and present, which, while not proposed as theories of language *processing* nevertheless are stated in a unidirectional way. As noted earlier ([Section 2.10](#)) many informal accounts of the workings of humorous texts are also phrased unidirectionally, usually from the viewpoint of the hearer of the joke rather than the producer.
- (d) None of the systems is intended to be an ‘intelligent’ or ‘creative’ agent (see [Section 2.12](#)). The projects explore the (linguistic) structure of specific classes of jokes, and the software works through its procedures blindly, exhaustively churning out results.

Specific observations about the pun-generators are:

- There is no explicit use of a theory of humour in the design of the systems; the various relationships used are based on structural regularities which the designers have discerned in example texts (although Venour outlines how his structures can be seen as embodying ideas from the GTVH). If humour results from any ‘incongruity’ in the output texts, this is purely fortuitous.
- It is not too difficult to fit most of the texts generated by these programs into our definition ([Section 9.3.6](#) and [Appendix A.2.2](#)) of linguistically normal paradigmatic puns: the initial part of each text establishes a context with two concepts mentioned in it, and the final part of the text mentions both concepts using some form of phonetic similarity.
- There is no notion of supplying a particular input parameter to the system: the program always runs with all its knowledge bases.

The HAHAAcronym project differs from the pun projects in the following ways:

- In both its functionalities, it starts from some input item (either an existing abbreviation, or two words), whereas the pun-generators start only from their basic knowledge bases.
- It could be argued that the use of domain opposition appeals to some variant of the incongruity approach to humour, possibly a version of ‘script opposition’ ([Chapter 6](#)).

In a sense, these programs exemplify a small methodological paradigm for building and evaluating joke generators, which we could set out as follows (although no system adheres totally to this doctrine):

- *Separation of knowledge.* All data sources used by the program should have a clear status as contributing to the humour, or being merely supporting information (e.g. linguistic knowledge). This makes explicit what the implemented model has to say about humour, as opposed to accidental details of the working program. For example, the schemas in JAPE and HCPP or the domain label oppositions in HAHACronym are humour knowledge, while the text generation facilities in these systems are not.
- *Generality of supporting data.* Any data source (e.g. the lexicon) which is not classed as humour-creating should be constructed so as to be neutral, i.e. not accidentally smuggling humorous effects into the process.
- *Encapsulation of support.* If a supporting mechanism (e.g. text generation) is required, then the manner in which it achieves this is not important to humour research, providing it supplies the appropriate functionality and obeys the *generality of supporting data* convention. The VINCI-based systems use a relatively principled text-generator, whereas JAPE and WisCraic have *ad hoc* mechanisms, but this does not affect the claims made about humorous structure.
- *Control of input data.* If a system takes input parameters (as HAHACronym does) then any experimental testing (as opposed to debugging during development) should have a well-defined procedure for selecting these values (e.g. by independent agents, randomly).
- *Evaluation of results.* Where a program claims to generate humorous output, this should be tested by presenting the output in a controlled manner to suitable human judges. All the normal requirements of psychological experimental design should be followed. The evaluation of JAPE is the most sophisticated of those listed above, and the HAHACronym testing is also quite extensive. The WisCraic evaluation suffers from having very small numbers and uneven coverage of test items.
- *Selection of output data.* If the output of a system is to be evaluated in a controlled way, then the choice of output items (from the system's total output) should be made in a systematic and unbiased way. Of the systems summarized above, only the JAPE evaluation records its procedure.
- *Use of scales.* One non-trivial issue to be decided is what questions should be asked of the human judges. The use of a five-point funniness scale is popular, but care is needed in labelling the scale. The Venour/JAPE scale had a lowest point which was 'Not a joke', thereby conflating jokehood and funniness. Binsted *et al.* (1997)'s test of JAPE separated jokehood and funniness into separate questions, but presented funniness as a symmetrical scale about a central point, which implies that there is a neutral notion of 'not funny' and then some

negative' forms of (un)funniness. They also omitted to offer a 'not sure' option for the jokehood question. The scale used for HAHAcronym – where the points are *not funny*, *not very funny*, *mildly funny*, *funny*, *very funny* – is more convincing. However, this means that comparisons of absolute numerical results across the different projects are of little interest. Informative results can nevertheless be presented by focussing on statistical relations (e.g. correlation, difference) between separate measures *within* an evaluation, a course followed by both Binsted *et al.* (1997) and McKay (2002). None of these projects seemed to take account of work on psychometric studies of sense of humour, where scale-based testing is more established (e.g. the seven-point scale from 'not at all funny' to 'very funny' in the 3WD test (Ruch 1992)).

The whole issue of establishing a sound methodology for evaluating computer systems which create subjectively assessed artefacts is not trivial; see Pearce and Wiggins (2001), Ritchie (2001), Pease *et al.* (2001).

10.10 Summary

This work is relevant to the linguistic methodology outlined in earlier chapters. These projects involved detailed description of the structural characteristics of a small class of jokes; their rules were computer tested, and the rule-generated texts were evaluated by human subjects. Although the computations defined by the pun-generators can be seen as a special case of the general definition of puns in [Chapter 9](#), it is not immediately clear how to generalize their rules or architectures to a wider class of jokes.

11

Pragmatic and discourse issues

We discuss, informally, some factors involving the linguistic organization of the joke beyond the single sentence.

11.1

Questions and interaction

Superficially, it might seem that jokes such as light bulb jokes ([Chapter 6](#)), riddles, and knock-knock jokes ([Chapter 9](#)) require some active participation by the audience, so that our pragmatic model should include not only a joke-audience but a *joke-teller*. Although jokes of these classes are often told interactively, with a brief interjection by the audience, the crucial point here is that this interjection is completely formulaic and contributes no substance to the actual joke. In the case of a riddle (or a light bulb joke), a successful telling of the joke requires the audience merely to indicate ignorance, by saying *No* or *I don't know*, or something similar. In a knock-knock joke such as (113), the first audience line is fixed, and the second is defined as the second teller's line followed by *who*? Such classes of question-answer jokes are often transmitted in joke books, which emphasizes that the humour (or at least the text's status as a joke) does not depend on actual verbal interaction between two people. Indeed, when riddles and light bulb jokes are included in joke books, the text is often presented simply as question and answer, without the (predictable) intervening response.

We are focussing on conditions for a text to be a joke, but not (as yet) the factors which enhance its funniness (see [Section 12.3](#)). This suggests we can ignore for the moment the assignment of joke-telling roles.

Conjecture: the structural factors which qualify a question-answer text as a joke are independent of which interlocutors tell which parts of the joke.

As noted earlier, the waiter-character in (24) can be seen as misinterpreting the diner-character's utterance as a boast or cry of pleasure rather than as a complaint, forcing the audience to see this alternative interpretation.

(24) Diner: Waiter! There's a fly in my soup!

Waiter: Please don't shout so loudly – everyone will want one.

That is, (24) can also be classed as a forced reinterpretation joke. It is not obvious how best to characterize the misinterpretation, but it could be said to be a misassignment of *perlocutionary effect* the intended effect of the utterance (Levison 1983). A similar structure occurs in (61).

(61) What's the difference between an elephant and a watermelon?

(I don't know.)

You'd be a fine one to send to the store for a watermelon.

In (61), the questioner, using semantic-pragmatic ambiguity, manoeuvres the other participant into expecting a riddle (for which an unthinking, formulaic confession of ignorance is appropriate), then reacts as if the question had been a genuine request for information. The misinterpretation here is describable as a misassignment of *illocutionary force* (Levison 1983): loosely speaking, the type of speech act. In a sense, these questions act as practical jokes, making the interlocutor appear foolish. This is an example where using the question–answer format is crucial.

If such jokes are conveyed in their entirety, for example in a joke book, then the text is an example of forced reinterpretation. In such presentations, the intervening *No* or *I don't know* has to be included, as the punchline does not make sense otherwise. However, these jokes are typically expressed interactively, with the respondent (second speaker) acting as the joke-audience and the first speaker as the joke-teller. In order to model or describe the workings of such joke-tellings, we need to be able to represent or classify the pragmatic status of an utterance (e.g. in terms of illocutionary force), so as to state formally (or semi-formally) the joke-interpreting experience of the audience (or victim) of one of these jokes (see also [Section 12.3](#)).

11.2

The initiator/responder distinction

Cornwell and Hobbs discuss the importance of social interaction in the making of puns, and draw a distinction between an *initiator* pun and a *responder* pun:

In an initiator pun, the speaker uses a sound, or closely related sounds, and contrives to shift the listener from one meaning to another.

(Cornwell and Hobbs 1991:205)

In a responder pun, the listener gives a sound a meaning different from that being used by the initiator and does so in such a way as to cause the initiator to note the second meaning too.

(Cornwell and Hobbs 1991:206)

Examples of initiator puns abound, as they are the conventional form where the joke-teller is in charge and the joke-audience merely listens. Responder puns are exemplified by the dialogues (170) and (171).

(170)A: I feel like a cup of tea.

B: You don't look like one!

(Cornwell and Hobbs 1991:206)

(171)A: My ice-cream is dripping.

B: Oh that's funny I thought it was ice-cream!

(Cornwell and Hobbs 1991:208)

In these jokes, the second participant deliberately misinterprets what the first speaker has said for (supposedly) humorous purposes. The overall text, with both utterances, is structurally a forced reinterpretation joke, in our terms; indeed, that seems to be a necessary consequence of the Cornwell-Hobbs definition of a responder pun. The resulting joke differs from a conventional (initiator) joke in terms of who has control over this joke structure, and whether the first speaker intends a joke. What it does share with an initiator joke is that the person delivering the punchline is the one who creates the joke; the difference is that the other speaker both supplies the set-up and acts as the joke-audience.

Leaving aside the finer issues of what constitutes a pun – (170) and (171) do not fall into the structural class of puns that we have outlined in [Section 9.3.6](#) – it seems clear that this distinction is applicable to jokes other than plays on word meanings, since contrived misunderstandings with humorous intent occur in everyday life, and they are not all based on phonetic-lexical ambiguity. Examples (25) and (26), although recounted as narrative jokes, could be seen as descriptions of scenes in which a responder joke occurs.

Although the initiator/responder distinction would be important in a full theory of joke-telling and human interaction, it is beyond the scope of our enquiry. We abstract away from the details of which participants deliver the set-up and the punchline, and analyse the resulting text and its relation to the joke-audience.

11.3

The contribution of questions

Even if we choose to analyse question–answer jokes at a level which ignores the potential for interactive delivery, we should still consider the way that a question may enhance the humour of a text.

In cases where a joke has two possible formulations, one in a question format and one not (e.g. (8) and (172), (41) and (173)), the question version often seems intuitively more amusing.

(172) My father explained to me why birds fly south in winter.

It's because it's too far to walk.

(constructed item)

(173) I went to the zoo yesterday. I was surprised to see that the elephant was wearing red socks, but it turned out that his green ones were being washed.

(constructed item)

This is not a firm empirical fact, and some careful experiments would be needed to establish whether the question format is generally more effective. If this were shown to be the case, we would then have to seek an explanation.

Notice that these variants, although not using direct questions, rely on sentence constructions (an indirect question in (172)) which in some way allow an ambiguity of 'focus', loosely speaking.

One possible reason for using a question–answer structure might be the options for arranging the information within the text. The question–answer structure provides a very natural way for a particular piece of information (the answer) to be presented last.

(142) What do you call a person who puts you in touch with the spirit world? A bartender.

If we were to rephrase (142) without using a question, plausible wordings might be (174) or (175).

(174) A bartender is a person who puts you in touch with the spirit world.

(constructed item)

(175) A person who puts you in touch with the spirit world is called a bartender.

(constructed item)

Of these, (174) is the more ‘unmarked’ linguistic structure, but it undermines the joke by re-ordering the critical information. (It is arguable that this sentence is still humorous, but in a different way.) Version (175) does not destroy the joke, but it resorts to a slightly less neutral or natural sentence structure than (174). The original question–answer form achieves both fluency and the correct ordering.

The second ‘explanation’ for the putative effectiveness of the question format could be that it allows the audience some time to ponder (particularly if the joke is being told interactively, rather than read in a joke-book), and this may in some way increase the impact of the answer. Timing has long been acknowledged to be an important aspect of oral joke-telling. It would be interesting to carry out an experiment in which question–answer jokes, and non-question paraphrases such as (175), are presented to subjects with varying delays between set-up and punchline, to see how this affects ratings of the funniness of jokes. Wilson (1979) describes some experiments on timing, but the issue is still not fully explored.

More broadly, even non-question–answer jokes may raise ‘questions’ in the audience’s mind during the set-up, and this may contribute to the joke’s impact. For example, (62) and (176) (the original, slightly longer, variant of (33)) both have mild puzzles in their initial narratives, with the punchline resolving the puzzle.

(176) Sitting on the side of the highway waiting to catch speeding drivers, a State Police Officer sees a car pattering along at 22mph. He thinks to himself, ‘This driver is just as dangerous as a speeder!’ So he turns on his lights and pulls the driver over. Approaching the car, he notices that there are five old ladies – two in the front seat and three in the back – eyes wide and white as ghosts. The driver, obviously confused, says to him, ‘Officer, I don’t understand, I was doing exactly the speed limit! What seems to be the problem?’ ‘Ma’am,’ the officer replies, ‘You weren’t speeding, but you should know that driving slower than the speed limit can also be a danger to other drivers.’ ‘Slower than the speed limit? No sir, I was doing the speed limit exactly ... twenty-two miles an hour!’ the old woman says a bit proudly. The State Police Officer, trying to contain a chuckle, explains to her that ‘22’ is the route number, not the speed limit. A bit embarrassed, the woman grins and thanks the officer for pointing out her error. ‘But before I let you go, Ma’am, I have to ask ... Is everyone in this car OK ? These women seem awfully shaken and they haven’t muttered a single peep this whole time,’ the officer asks. ‘Oh, they’ll be alright in a minute officer. We just got off Route 119.’

(edited from www.jokes2000.com)

If this speculation is correct, then we have a humour-related factor which is separate from the more structural properties of the joke: the set-up introduces a question which the punchline (directly or indirectly) answers. This is hard to formalize, not only because it cuts across the various joke classes, but because it is not clear, in terms of knowledge representation, what counts as a ‘question’ in this sense.

On the evidence of the examples considered so far (all the question–answer jokes, and a few narratives such as (176)) it does not appear that the question factor is essential to the status of the text as a joke – it simply increases the ‘funniness’. As has been demonstrated, question–answer jokes can usually be rearranged into non-question form (except for certain inherently interactive examples, as in [Section 11.1](#) above). Also, narrative jokes with implicit questions may have very similar variants with no question: compare (176) and (33).

This leads us to the following:

Conjecture: for any narrative joke, and most question–answer jokes, where a question (direct or implicit) contributes to the humorous effect (other than via ambiguity about illocutionary force), there is a similar joke without this question which relies upon the same central mechanism and content.

Following our earlier discussion, we also have the following speculation:

Conjecture: if two jokes have identical semantic content but one is stated declaratively and the other as a question–answer pair, the latter will generally be slightly funnier.

A possible experiment to explore this issue would be to have subjects rate, for jokehood and funniness, jokes which contained a question and near-equivalent jokes with the question element eliminated (such as (142) and (175), or (176) and (33)).

The raising of questions in narrative jokes such as (176) is comparable to the way in which non-humorous stories often raise questions for the audience in order to hold the attention, and then provide a satisfactory outcome by resolving these questions (Palmer 1991, Laurel 1993). It may well be that the ‘funniness factor’ proposed above (i.e. that a joke will be more effective if it involves a question which is later answered) is not specific to humour, but is part of some more general satisfaction that an audience experiences when a puzzle is solved or a question answered.

11.4 Questions, descriptions, and riddles

11.4.1 *Reinterpreting questions*

Examples (4) and (28) may not at first appear to be FR jokes, but it is arguable that they conform to this pattern.

(4) What's black and dangerous and sits in a tree?
A crow with a machine gun.

(28) What is grey, has four legs, and a trunk?
A mouse on vacation.

In these, the set-up (question) seems to have a natural answer (an elephant for (28), a panther or similar animal for (4)), but instead a different reply is supplied. In (28), the alternative answer corresponds naturally to a reinterpretation of the set-up, particularly the phonetic-lexical ambiguity of *trunk*. In (4), the nature of the reinterpretation is less clear, but it could be viewed as a higher-level ambiguity concerning the possible referent(s) of a phrase, what might be termed *semantic-referential* ambiguity.

This class of joke seems to be structurally similar to a style of joke often used by comedians and satirists, in which a verbal description is given which seems to be talking about one person, group, situation or event, but is suddenly revealed to be about something else. This form has often been used on topical comedy shows on BBC Radio such as *Weekending* and *The News Huddlines*; (177) is a typical example, delivered in a context where bribery and corruption of international cricket teams was being discussed.

(177) The police have been investigating a group of people hell-bent on preventing the England cricket team from winning. But there is not much they can do about the selection committee.
(*The News Quiz*, BBC Radio 4, 17/18 September 1999)

It is not at once obvious where these jokes manifest what we have called INAPPROPRIATENESS. In (4) and (28), the answer (punchline) seems to contribute some ABSURDITY, although it could be argued that the integrated interpretation of the (reinterpreted) question and the answer is ABSURD, as in the general FR situation. In (177), there is no ABSURDITY in the punchline itself, but the humour may come from the (sudden) comparison of the two targets of the description, the initially implied one and the eventually revealed one. That is, this class of joke may involve what we called, in [Chapter 5](#), CONTRAST. Or there could be some INAPPROPRIATENESS in the integrated interpretation of the whole text.

11.4.2 *Riddles*

There are other classes of joke (or puzzles) to which the notion of a ‘description’ is central. These texts have a clear concrete division into set-up and punchline, often in the form of question and answer, or posed problem and solution. As with (4) or (28), the ‘description’ (a set of attributes or characteristics of some entity) is supplied by the set-up (typically a question) and the punchline (typically an answer) supplies some entity or situation which matches, or is alleged to match, that description.

The clearest and simplest example of this is the *riddle*. Riddles, particularly their linguistic aspects, have been widely discussed (Pepicello and Green (1984) give a typical review), but there is no accepted definition. Dienhart offers this ‘working characterization’:

A riddle can be viewed as a discourse class involving a two-part structure: an initial text’ (the ‘precedent’) in which the riddler supplies a series of clues (generally insufficient or misleading) from which a second text (the sequent) is to be surmised by the riddlee. The precedent is often, but not necessarily, expressed in the form of a question. Together, the precedent and the sequent make up the riddling text as a whole. This text establishes a link or (‘bisociation’) between two ‘scripts’ (or ‘frames of reference’). The riddler, through the precedent, deliberately plays up one script and then, through the sequent, awakens the riddlee’s awareness of a second script.

(Dienhart 1999:104)

This is both a delimitation of the set of data and a suggestion of how the forms work: the first three sentences of this passage outline the data, but the final two sentences are more in the nature of a theoretical analysis (which seems heavily influenced by Raskin (1985), and adopts essentially what we have called the forced reinterpretation process). It is not clear that all of the ‘riddles’ which Dienhart subsequently discusses accord with the analysis he sketches, although it is a very plausible account of some question–answer jokes, such as (8) or (142).

Discussions of riddles generally identify them as question–answer forms, but Dienhart’s characterization (above) points out that an initial question is not necessary. Although the initial part may not be a question *syntactically*, a riddle poses a problem, and so is in a sense a question, even if indirectly.

To some extent, there is a continuum of how humorous these question–answer discourses are intended to be. Some are genuine problems of a ‘brain-teaser’ nature, such as the traditional item (178).

(178) In spring I am gay,
 In handsome array;
 In summer more clothing I wear;
 When colder it grows,
 I fling off my clothes;
 And in winter quite naked appear.
 (solution: a tree)

(Dienhart 1999)

or the simpler (179).

(179) What has a mouth but cannot eat?
 A river.

(De Palma and Weiner 1992)

This style of riddle appears to have been a common form in Western culture. Maranda (1971) discusses the structure of examples from Finnish.

Because of the tendency to treat the question–answer format as a defining characteristic of riddles, some humorous texts which are not truly problems to be solved are often classed as riddles. Raskin (1987:446) refers to ‘the conundrum – the punning riddle, the one the hearer can never solve’, but (surprisingly) he offers (6) as the typical example of this. Pepicello and Weisberg (1983) say (31) is merely a *parody* of a riddle, but Dienhart offers (180) (which can be seen as an FR joke) as a riddle:

(180) What is a *cloak*?
 The mating call of the Chinese frog.

(Dienhart 1999)

In such examples, a particular form of verbal interaction which has traditionally been associated with problem-posing has been borrowed (some might say subverted) in order to achieve humorous effects. It could be argued that many of these examples are not truly riddles, a term which should perhaps be reserved for the genuine problem-posing texts (even if they are not humorous).

11.4.3 *Describing absurdity*

Sometimes a question–answer joke requires an answer which depicts some particularly strange situation, as in (181).

(181) What goes cluck-cluck bang?
 A chicken in a minefield.

(Dienhart 1999)

The situation may even be completely bizarre, as in (182).

(182) What goes ninety-nine clop?
A centipede with a wooden leg.

These jokes are arranged as follows. There is a situation, preferably rather odd. There is a description of that situation, which may itself be rather peculiar. There is no ambiguity or punning. The order of presentation is also important, as the variants (183) and (184) barely work as jokes.

(183) A chicken in a minefield goes cluck-cluck bang. (constructed item)

(184) A centipede with a wooden leg goes ninety-nine clop. (constructed item)

11.4.4 *Summing up*

This leaves us with various questions:

- (a) Are examples such as (178) to be classed as humorous texts, and therefore to be accounted for by any complete theory of jokes?
- (b) If not, how do they differ from more obviously humorous examples such as (28), (142), and (180)?
- (c) Do extremely absurd examples, such as (181) and (182), operate in the same way, or are they yet another subclass?
- (d) Is the difference just a matter of degree, with all these examples sharing a basic form, varying in the extent to which some humour-inducing factor(s) (e.g. 'incongruity') are involved?
- (e) What does the question-answer structure contribute, and how vital is it?

As noted above, examples where there seems to be an obvious or natural answer ((4), (28), (180)) could be viewed as FR jokes; that is, part of the humour may come from the discrepancy between the predictable answer and the one supplied in the punchline. On the other hand (178), (179), (181) and (182) would normally *not* bring to mind an obvious answer, resulting in them acting like a 'puzzle' (albeit a near-impossible one in some cases). (179) could also be analysed in FR terms, since the punchline forces a reinterpretation of *mouth* (and hence the whole question set-up), even though (unlike (28)) it does not have an apparently obvious answer. Similarly, (178) could be seen as an FR example, if we include the reinterpretation of words such as *clothing* and *naked* in *metaphorical* senses. This would mean that we were classing as FR jokes both question-answer

pairs where the ‘reinterpretation’ involves the meaning of the question and those where it is the inferred ‘natural’ answer which is found to be wrong.

The remaining questions without obvious answers – (181), (182) – do not fit so neatly into the FR category, nor into any of the other presentational classes in [Chapter 8](#), nor any of the classes of puns in [Chapter 9](#). We shall propose a further subclass of jokes to contain them, albeit a rather narrow class. A *description-based* joke consists of these ingredients:

- (a) a *description*, *D*
- (b) a scenario or set of entities, *S*
- (c) a statement, explicit or implicit, that *D* describes or is true of *S*
- (d) the relationship between *D* and *S* is sufficiently unobvious that *D* would not normally evoke *S*
- (e) *S* is in some way strange or absurd.

A possible further factor is that *D* may be strange or absurd, though the notion of absurdity for verbal descriptions is rather unclear. Puzzle-style riddles such as (178) or (179) meet the first four of these conditions, but not the requirement that *S* is absurd.

Notice that (4), which we have already analysed in FR terms, would fit neatly into this class.

- (4) What’s black and dangerous and sits in a tree?

A crow with a machine gun.

Indeed, in a particular telling of the joke where the audience did not notice the ‘obvious’ answer to its question (i.e. they perceived it as being a completely open question), it could not, on that occasion, function in FR terms. However, it could still be comprehended as a joke by virtue of its membership of the class of description-based jokes. Similar remarks would apply to (28). This might seem to imply that we have an example which meets the structural requirements for more than one of our subclasses of joke. However, closer consideration shows that this is not the case. The membership of the text in a joke class is dependent on properties which are *relative to the knowledge and perceptions of the joke-audience*. Hence, if the audience’s knowledge, beliefs, and attention are such that the ‘obvious’ answer is *not* noticed, then the classification of the text is, for that audience at that moment, not an FR joke. If the audience does notice the obvious meaning, or has it pointed out afterwards, then the joke will be, relative to that state of knowledge, an FR joke.

To speculate further, examples like (4) or (177) may suggest that an audience finds it pleasurable to discover that a descriptive text can refer to two very disparate entities in an unforeseen way.

11.5 Repetition and sequencing

Many narrative jokes involve a sequence of events which repeat, to some extent, a pattern of activity. On closer inspection, it can be seen that sequences or repetition can be used to achieve different purposes in different jokes, and that the repetition may not be as central to the joke as it first appears.

A very common use of repetition is that it establishes a pattern, which is then modified in some way on the final (often the third) cycle. Relying on the pattern, the punchline can succinctly state some ‘incongruous’ action or idea. (In a dissertation on the automatic detection of key signature in music (Steedman 1973), the following ‘Principle of Congruence’ is offered: ‘No non-congruent event will occur until such a point in the piece that the event can be perceived unambiguously to be non-congruent, by the prior establishment of sufficient framework.’)

Consider example (185).

(185) An Englishman, a Scotsman and an Irishman were trying to get in to see the Olympics without tickets. The Englishman walked around the stadium and saw a pole lying on the ground and picked it up. He walked to the entrance and said, ‘Peter. England. Pole vault.’ The guards let him in without hesitation. The Scotsman found a manhole cover, carried it under his arm to the entrance, saying ‘McGregor. Scotland. Discus throwing.’ The guards let him in also. The Irishman was very frantic, since both his friends were now inside. He walked around the stadium and found a roll of barbed wire. He picked it up, walked to the entrance and said, ‘Murphy. Ireland. Fencing.’

(edited from rangsk.hypermart.net/humor.html and other WWW sites)

Here, the set-up sequence of ‘English’ and ‘Scottish’ is used to establish a pattern or template which the punchline can be set against. (The use of England and Scotland in the set-up is presumably in compliance with joke conventions, rather than to increase verisimilitude, as these nations do not compete separately at the Olympics.) Consider a shortened form, (186), of this joke.

(186) An Irishman was trying to get in to see the Olympics without tickets. He walked around the stadium and found a roll of barbed wire. He picked it up, walked to the entrance and said, Murphy. Ireland. Fencing.’

(constructed item)

This version contains the same salient features and seems to work as a joke, but may be harder to understand without the establishment of the

pattern *name-country-sport*. That is, the set-up acts to *facilitate* understanding of the punchline, the latter being a further instance of a template introduced in the earlier repetitions. Although (185) also uses a stereotype of the stupid Irishman, the central idea could be conveyed without that; consider (186) amended to feature a Canadian, or even a nationality-free version such as (187).

(187) A man was trying to get in to see the Olympics without tickets. He noticed that competitors were allowed in if they gave their name and their chosen event, and were carrying their equipment. He walked around the stadium and found a roll of barbed wire. He picked it up, walked to the entrance and said, 'John Smith. Fencing.'

(constructed item)

The version in (187) also shows how the set-up need not use a repetitive narration to establish the pattern: the text could *summarize* a sequence of events (number unspecified), without the text itself having a cyclic structure.

Sometimes a joke which appears to be of the same class as (185) is even less dependent on repetition and sequencing in the set-up, as for example (188) (which also exercises the Irish stereotype).

(188) An Englishman, a Scotsman and an Irishman were looking for jobs with a construction company. The foreman decided to test them, so he set each of them to put in telegraph poles for a day. At the end of the day, he asked the Englishman how many he had put in. 'Twenty,' was the reply. Then he asked the Scotsman, who also replied 'Twenty'. Then it was the Irishman's turn and the foreman asked him how many telegraph poles *he* had been able to put up in that day.

Three', said the Irishman. The foreman blinked, 'Three?' he repeated, puzzled. 'Only three? That's not very many. The Englishman and the Scotsman have both put in twenty each – how do you account for that?' 'Ah,' replied the Irishman, 'they may indeed have put in twenty each, but did you see how much they left sticking up out of the ground?'

(shortened from Raju (1991))

Here, there is no real need to establish a pattern. The central idea is a misconception about what it means to put in a telegraph pole. In order for the punchline to work, there is a need to contrast this mistaken view with (preferably a concrete instance of) the conventional one, but this can be done more succinctly. Consider the variant in (189) (which also has the national stereotyping removed).

(189) A company engaged in erecting telegraph poles recruited a new labourer, and the foreman set him to work on his own for the first day. At the end of the day, the foreman asked the new worker how many poles he had put in. 'Three,' was the reply. The foreman blinked, 'Three?' he repeated, puzzled. 'Only three? That's not very many. All the rest of the workers have put in about twenty each – how do you account for that?' 'Ah,' replied the recruit, 'they may indeed have put in twenty each, but did you see how much they left sticking up out of the ground?'

(constructed item)

Why, then, does (188) adopt the three-nationality sequence? It may be that this is done partly to incorporate the Irish stereotype (i.e. for offensive purposes), although that could be done in a shorter variant comparable to (189). Another reason may be that this form is a conventional pattern in joke-telling. It may also be that repetition has some role in building up tension or creating momentum (see [Section 12.3](#)).

A further example where repetition does not appear to be essential to the logic of the joke is (190).

(190) Ethel Barrymore was told that a hated rival of many years' standing had taken unto herself a new husband – and had made a full confession of past indiscretions to him. 'What honesty! What courage!' marvelled the reporter. 'What a memory!' added Miss Barrymore.

(Cerf 1964)

In this joke, the contribution from the reporter-character is not essential to the joke, but somehow enhances the telling.

(191) Two Irishmen were digging a ditch directly across from a brothel. Suddenly, they saw a rabbi walk up to the front door, glance around and duck inside. 'Ah, will you look at that?' one ditch digger said, 'What's our world comin' to when men of th' cloth are visitin' such places?' A short time later, a Protestant minister walked up to the door and quietly slipped inside. 'Do you believe that?' the workman exclaimed, 'Why, 'tis no wonder th' young people today are so confused, what with the example clergymen set for them.' After an hour went by, the men watched as the local parish priest quickly entered the whore house. 'Ah, what a pity,' the digger said, leaning on his shovel, 'One of th' poor lasses must be ill.'

(edited from www.geocities.com/Tokyo/7928/joke)

In (191), repetition establishes a pattern which is then treated differently on the final cycle, not necessarily in an inherently incongruous way.

A CONTRAST is thus presented between the two ways of treatment or perspective. (As with the examples in [Section 8.4](#), it is moot whether the CONTRAST is with the views expressed in the set-up, or with some predicted position.)

Yet another form of repetition is where the sequence progresses in some way (e.g. characters boasting of greater prowess in some field, or outdoing each other in tales of hardship). The final cycle continues this progression, using extreme exaggeration or implausibility to create an excessive instance of the pattern; for example, (192) (which also contains, in its opening phrase, a good example of what we have called – in [Section 7.3](#) – *extraneous* information).

(192) A wealthy grain merchant bought a new car and was enjoying his first ride in it no end until a motorcycle cop stopped him and suggested a visit to the local magistrate. ‘He was doing sixty,’ reported the cop. ‘Nonsense,’ declared the merchant. ‘I never got her up above forty.’ The merchant’s wife put in her two cents’ worth at this point. ‘He wasn’t going faster than twenty-five,’ she averred. A friend who had been riding in the back seat added, ‘I’d say we were virtually at a standstill when this officer came along.’ The magistrate threw up his hands and cried, ‘Stop right now – before you folks back into something.’

(Cerf 1964)

Another example of this genre is (193), in which the set-up consists of a ‘sequence’ of just one cycle, thus emphasizing that if the pattern is conveyed clearly enough, repetition may not be needed.

(193) George S. Kaufman, after a self-made millionaire boasted, ‘I was born into this world without a single penny.’: When *I* was born, I owed twelve dollars.’

(Cerf 1964)

The extent to which a repetitive *narrative structure* is needed will vary across these jokes. Although in some cases, it is straightforward to use a brief summary of a sequence of events (see (187), (189)), it might be more difficult in others (e.g. examples of exaggerated outdoing, though see (193) above).

To summarize:

- although the function of repetition in a joke set-up is generally to establish a pattern, there is more than one way that such a pattern can be used in the punchline;

- in many cases, it would be possible to establish the pattern either by summarizing a sequence of events or in some more explicit manner, so that repetitive narrative structure is not strictly necessary;
- repetition may not be central to the joke, and its role may be relatively minor, peripheral or even redundant;
- even where it is not logically essential, it may be that repetitive narration acts as a supplementary factor, adding impact to the joke.

11.6 Summary

Question–answer structure can contribute to a joke in various ways, and is sometimes present implicitly even when there is, syntactically, no question. Some jokes are conventionally told interactively, but need not be, while others are more like practical jokes and require participation. Not all question–answer jokes are conundrum-style riddles.

Repetition can also play a number of roles within jokes, and it is important to distinguish repetition of content from repetition within the narrative structure.

12

Speculations on joke structure

We bring together some of the strands from earlier chapters, and offer some proposals about what could be the principal factors within a joke.

12.1

Towards a notion of ‘structural description’

Given the analogy between linguistic description and joke description put forward in [Section 2.3](#), it is interesting to consider what would count as a *structural description* of a joke. A grammarian working within a particular linguistic framework can annotate a sentence to indicate what structural constituents or abstract properties that sentence has, thereby making explicit all the linguistically interesting aspects of the sentence. It would be interesting to devise a comparable tabulation of descriptive statements available to a joke-theorist for stating the humour-relevant components of a joke. The GTVH ([Chapter 6](#)) offers exactly such a scheme – a joke is a six-tuple of parameters – but we will consider possible alternatives.

Such a classificatory arrangement would constitute a theory of joke structure, if it were to be given the status of a substantive claim about what does/does not constitute a possible joke. Our work is not yet at a stage where we can make such predictive claims. That would need more investigation of varying joke classes, and further refinement of the necessary abstractions. Nevertheless, it is hard to resist the temptation to try to draw together some of the patterns we have proposed for classes of jokes and to see what further generalizations are possible. In this chapter, we shall be rather speculative, suggesting some broad but tentative ideas about what might constitute the semantic and pragmatic components of jokes. Although we believe that theories should be formally stated and based on clear evidence, there is a role for the informal expression of intuitions, and the discursive presentation of preliminary suggestions; that is what this chapter inclines towards. Our concerns will still be relatively structural, and we will not even consider the question of what causes humour in general, or why jokes are funny.

In this chapter, therefore, we will consider, as far as possible, the various *types of information* that are relevant to the workings of jokes.

It may be that there are not very many joke classes that are definable in terms of their use of linguistic devices. The two broad classes of joke that we have concentrated on – puns and forced reinterpretation – have attracted a great deal of attention in linguistic treatments of humour, and may even be the most documented and discussed joke classes. It could be that as one explores outwards from these language-related joke classes, one finds very few patterns in the language of jokes.

To explicate the humorous effects of such jokes we would need to go beyond language and attempt to characterize the logical (or pseudo-logical) relationships between objects and events described by the joke. A complete study of that issue would merge into the much broader issue of what events and situations are funny and why, and that is beyond the scope of this work. What we will do here is try to generalize some of our earlier analyses, considering the various ingredients of a joke.

Let us concentrate for the moment on propositional, rather than linguistic, jokes ([Section 2.7](#)).

12.2

The content of propositional jokes

12.2.1

Delivery mechanisms

All the definitions which we have given (principally in [Chapter 8](#)) relate to the *delivery* of the humorous content rather than the actual detection, comprehension or appreciation of humour. We could therefore label these classifications as *delivery mechanisms*.

What, then, is the functionality of a delivery mechanism? That is, what does it map a joke text to, abstractly? As presented so far, it is simply a property that can be true or false of a text: the joke either has punchline-revision or it does not, for example. As we have noted already (e.g. [Section 8.6.4](#)), we need to connect these presentational definitions to some notion of INAPPROPRIATENESS or CONTRAST occurring within a specific interpretation. It is noticeable that several of the delivery mechanisms (discourse-coherence misunderstanding, punchline revision, punchline inference, double entendre) naturally give rise to a definition of two possible interpretations, loosely speaking before and after the processing of the punchline. Hence, we could consider defining a joke in terms of some condition on these two interpretations. At first glance, it might seem that one of the two interpretations has special status, as being ‘hidden’ or ‘indirect’, and that this is where the INAPPROPRIATENESS should be located. This would be the less OBVIOUS reading in discourse-coherence misunderstanding, in punchline revision and in double entendre, and the inferred interpretation in punchline inference. However, this suggestion

does not hold up well. As observed in [Section 8.3](#), the INAPPROPRIATENESS may be located either in the more or in the less OBVIOUS interpretation. Even for FR jokes, matters are more complex. (66) has the same inferential pattern as (65). However, it is hard to argue that the (initially less obvious) interpretation established by the punchline of (66) is somehow ABSURD or TABOO. Intuitively, the audience is led to expect an scenario which is mildly improper (accidental pregnancy and deception over paternity) but then receives a *less* TABOO interpretation. That is, if there is any INAPPROPRIATENESS it is in the reading that the audience should, if the joke works, find more OBVIOUS.

Hence, even for delivery mechanisms where a pair of interpretations seems to be involved, it does not seem adequate to single out one of these as the only available position for INAPPROPRIATENESS. The generalization may be that the INAPPROPRIATENESS occurs in *either one* of the two interpretations.

For simple narratives ([Section 8.5](#)), it would appear that the final interpretation is the only available location for any further factor. Nevertheless, it might be possible to view simple narration as a delivery mechanism which also provides two interpretations: immediately pre-punchline and post-punchline (as with the other delivery mechanisms discussed above). In contrast to the other delivery mechanisms discussed above, there is no choice about where the INAPPROPRIATENESS should be located: it must be post-punchline.

In the light of this, a possible formal generalization would be to have a delivery mechanism provide an *ordered pair* of interpretations (a primary and a secondary one).

Of the examples we have analysed in earlier chapters, very few of them seem to need a comparison of two interpretations (what we have dubbed CONTRAST), but some such relation does seem apt in certain cases ([Section 8.4](#)). (Also, the widely cited idea of script opposition – [Chapter 6](#) – is a form of CONTRAST). If CONTRAST is involved, it can be between the two interpretations supplied by the delivery mechanism.

Of course, this still defers as subsidiary research problems the relevant definitions of CONTRAST and of INAPPROPRIATENESS, or what we have proposed as the latter's two subtypes, ABSURD and TABOO.

It is still an open question whether this provisional notion of delivery mechanism will still be a useful generalization when further joke classes have been analysed.

As argued in [Section 8.1.1](#), the crucial INAPPROPRIATENESS may be quite indirectly conveyed by the text.

This leaves us the following speculative and rather vague proposal:

Conjecture: a text constitutes a propositional joke if its text embodies one of the delivery mechanisms defined earlier, and

(one of) the interpretation(s) indicated by that delivery mechanism implies or somehow alludes to information which is INAPPROPRIATE in some way, or if there is some form of CONTRAST between the interpretations.

This criterion is offered only as a *sufficient* condition for jokehood, not a *necessary* one. It would require both some refinement (i.e. further conditions) and some proper supporting evidence before it could be regarded as strongly corroborated.

Anecdotal evidence does suggest that impromptu remarks which hint at or conjure up TABOO images are commonly perceived as amusing. That is, the provisional conjecture above might be generalizable to cover innuendos or ‘wisecracks’ made in everyday contexts, as well as crafted jokes.

12.2.2

Is anything essential?

Another (rather flimsy) observation is that a joke which uses nothing but simple narrative (i.e. has no reinterpretations, double meanings, or any of the other presentational devices discussed in [Chapter 8](#)) seems to require more than a mere trace of ABSURD or TABOO material. Instead, it needs – as in (34) and (35) – some inherently bizarre logic (perhaps a special subclass of ABSURDITY) in order to be funny. The conjecture here is that a mild form of INAPPROPRIATENESS may produce a joke if conveyed sufficiently obliquely; e.g. (6), (26), (33). On the other hand, a simple narrative presents the final situation fairly baldly, and may need a more extreme form of ABSURDITY to achieve humorous effect; e.g. (188).

It is also interesting to consider whether the INAPPROPRIATENESS ingredient is absolutely necessary (for jokehood), or merely extremely helpful. Relevant evidence would be a joke involving one of the delivery mechanisms discussed in [Chapter 8](#), but no obvious INAPPROPRIATENESS or CONTRAST. These are hard to find, but (194) or (195) might be examples (both using discourse-coherence misunderstanding).

(194) ‘Excuse me, how long will the next bus be?’

About 12 metres.’

(195) Guess what was on the TV last night?

A vase of flowers.

It could be argued that (194) offers the notion that the questioner is interested in the physical length of a bus, which is ABSURD. This is not wholly convincing, and in (195) it is even harder to find something INAPPROPRIATE. Example (1) may be similar.

In contrast, jokes such as (30) and (31) seem to contain little more than a question–answer structure and copious amounts of ABSURDITY. As noted in [Section 2.4](#), it may be that there is no single core structure for all jokes, but simply a cluster of factors which can be put together in various ways to create humour.

12.3 Supplementary information

In [Section 7.3](#), we drew attention to the fact that a joke may contain content which is not strictly necessary to the central idea of the joke – *extraneous* information – and observed that such material was not necessarily completely useless. That is, this peripheral content or structure might enhance the funniness of a text, even if it was not essential to its joke-hood. We refer to such supporting devices as *supplementary information*, or *supplements*. Some are classes of content (that is, they refer to sets of propositions conveyed by the joke), while others are structural (being properties of the way the joke is presented). Many of them can be seen as properties of general narrative, rather than being peculiar to jokes.

We will summarize, informally, some of the ways in which supplements may contribute to a joke.

Inappropriateness

Although this factor seems to play a crucial role in ensuring that certain texts are jokes (at least if the speculations of [Section 12.2](#) above are broadly correct), it also functions in a supplementary role. Some jokes have ABSURD content over and above that which is conveyed by the punchline, and this may add to the funniness of the text. Examples such as (30) and (31) are cases in point, as are jokes which involve human-like behaviour by animals or vegetables or even inanimate objects. In Hofstadter and Gabor (1989), the notion of a suitable evoked image (over and above the logical structure of a joke) is suggested as a contributory factor to the funniness of two example (sexual) jokes. In (62), the notion of the soldiers chewing up the beermats could perhaps be deduced from the text's interpretation, and this image may enhance the funniness, while not being central to the logic of the joke. This factor may also be applicable to non-humorous texts, rendering them more interesting.

Question-solving

This factor was discussed in [Section 11.3](#). If the set-up of a joke raises a question or puzzle (explicitly or implicitly) and the punchline provides a solution (explicitly or implicitly), then the joke may be more satisfying. As noted previously, this sort of improvement may also occur in non-humorous texts.

Thwarted expectations

Some jokes seem to lead the audience to expect a particular *textual* ending, but the punchline deviates from this prediction (cf. [Section 5.3.2](#)). Here we refer to the narrow case where there is a very specific prediction of some *wording*, not just the wider case where (as in all FR jokes) the audience's interpretation is confounded. Such punchlines can produce a very marked surprise, particularly where the unexpected material is the last word or phrase of the joke, as in (71) or (72).

Disparagement

The text conveys an insulting message which accords with the audience's beliefs. This may be a generalized insult (applied to a whole set of targets) which is merely inferrable from information stated/IMPLIED regarding specific targets (i.e. characters in the joke); for example, (62). A direct insult (i.e. one aimed at the joke-audience) does not involve this supplementary factor *from the perspective of that audience*, but may well do so *from the perspective of another party*. That is, a text which directly insults the recipient is unlikely to have increased funniness for that audience/target. In such cases, a particular remark may, when first uttered, not be perceived as being supplemented (enhanced) by disparagement, but may do so when recounted later (or when overheard on the first occasion). The effect of a joke is always relative to some set of beliefs. Although it is widely recognized that many jokes convey offensive attitudes, often about particular nationalities (cf. (62), (185)), this does not mean that such insults are sufficient in themselves to create humour. Examples of jokes which insult particular targets or butts usually contain other mechanisms which qualify them as jokes.

Superiority over joke character(s)

Sometimes, a joke describes a situation in which some character acts foolishly, incompetently, etc. (e.g. (79)). It may be that this allows the audience to feel superior to the character. This differs from disparagement in that there is no generalization to some wider set outside the joke. Disparagement may be aimed at a specific person (e.g. a famous politician) or may be directed at a whole class of people (e.g. politicians, Englishmen) but in either case these targets *exist outside the joke*; any joke-characters who are shown as flawed are in the joke to represent the real target, in society. What we are calling superiority is directed only at an imaginary joke-character, with no scope for generalization beyond the world of the joke. (As is often commented – for example, Raskin (1985:36), Attardo (1994:49) – Hobbes (1840:46) proposed this as the *essential* factor in stimulating laughter.)

Embellishment

Very often, extraneous content is added by a joke-teller to enrich the audience's experience, by making the story more vivid or dramatic. The examples given earlier of extraneous information in example (60) are generally embellishments, and are quite a normal part of narration, even in non-humorous stories.

Dramatic tension

This is closely related to (and blends into) embellishment; the only distinction between the two is that embellishment is additional detail, whereas dramatic tension is structural. Examples of this would be repetition (see [Section 11.5](#)), or other narrative devices for prolonging the set-up. This could interact with question-solving. It is also a general technique, used in non-humorous narratives.

Facilitation

Any device (of content or structure) which allows the logical presentation to work more smoothly counts as a contribution to facilitation. For example, in [Section 8.6.1](#) we noted how certain choices of wording helped to ensure that a particular interpretation appeared more obvious. Also, if we compare (186) and (187), the set-up of (186) explicitly establishes the pattern for the punchline (name, country, event), which makes the punchline more immediately intelligible. Once again, this factor also appears in non-humorous narratives.

Parallelism

This factor – which is of broader application than just jokes – could perhaps be called *intertextuality* (Norrick 1989), (Attardo 1998:237). If the joke text, or some aspects of it, echo or imitate some familiar form, perhaps even a familiar joke form, this can enhance the effect. Parody is a particular case of this, as are catch-phrases. Certain puns (see [Chapter 9](#)) consist of little more than intertextuality at a detailed level.

Prank

A joke may be used to make the audience feel foolish, by tricking them in some particularly unfair way. In such cases, the joke acts as what is normally known as a 'practical joke'. Examples of this include pseudo-riddles based on semantic-pragmatic ambiguity ([Section 11.1](#)). There is a similarity between this and what we have termed superiority over

joke-character(s), the difference being that the individual being looked down upon in this case is the audience. The resulting ridicule may contribute pleasure to the joke-teller, or even to a third party (neither joke-teller nor joke-audience); the extent to which it creates enjoyment for the targeted audience may vary.

Ingenuity

A joke which impresses the audience with its cleverness may have extra impact. Although many puns are often regarded as poor examples of wit, occasionally a pun is deemed worthy of admiration, such as (97) or (141); this may be an acknowledgement of the creativity that produced them.

Funny words

It is a much-repeated part of comedic folklore that certain words are funnier than others, and that certain letters (e.g. in English, *k*) sound more amusing. If this is true, then the funniness of a joke will be enhanced by the inclusion of such material (cf. comments in [Section 7.1](#) about the choice of colour words for (42)).

The classes of supplementary information given above have some resemblance (at least in their names) to some of the *techniques* of Berger (1998: Ch. 1). Berger lists 45 techniques (including exaggeration, insults, ridicule, absurdity, repetition, embarrassment) which he says account for all humour. Berger's proposals are by no means a formal theory of humour, and it is not completely clear what some of his techniques are, even informally. An important distinction between Berger's account and ours is that he (apparently) claims that a text which manifests one of his techniques will be humorous; that is, each technique is a sufficient condition for humour. We do not suggest this: devices such as repetition or exaggeration are not sufficient for jokehood.

12.4

Extending to linguistic jokes

Both propositional and linguistic jokes allow the presence of extraneous material, and the types of supplement listed above are equally applicable to either of these classes. However, it is unclear that the facets discussed above in [Sections 12.2](#) apply to puns. In the classes of pun examined in [Chapter 9](#), the conveyed information is often totally lacking in humour, or in any trace of INAPPROPRIATENESS; the humour appears to emanate from the form of the text rather than its overall content (but see the discussion in [Section 9.7.7](#)). Although we could decide to call our definition of 'linguistically normal paradigmatic puns' ([Section 9.3.6](#) and

[Appendix A.2.2](#)) a ‘delivery mechanism’, it would be a very different entity, formally, from the information-presentation methods discussed in [Chapter 8](#) and [Section 12.2](#). In particular, there is no natural way of defining how a pun involves the expression of two interpretations (in the sense we are using, namely a consistent set of propositions).

This seems to underline the distinction which we have drawn between propositional and linguistic jokes. The former are defined in terms of propositional interpretations (*situations*, in the terminology of [Section 3.9](#) and [Appendix A.1](#)), Hence we can place further conditions on these interpretations: how they are conveyed by the delivery mechanism, and how they present content which is INAPPROPRIATE (a property of situations), or how they enter into CONTRAST (a relation between situations). Linguistic jokes (or at least those puns that we have considered) do not have their primary humour-producing components located in the (single) interpretation expressed by the text, and delivery mechanisms (in our formal sense) play no part.

12.5 What’s in a joke?

We can sum up the rather tentative discussion in this chapter as follows. Jokes fall into two broad classes: *propositional* and *linguistic*. A propositional joke has a *delivery mechanism* which indicates how the linguistic processing of the text can give rise to a *pair of interpretations* (loosely, pre- and post-punchline). It seems that such a joke must involve *either* CONTRAST between these two interpretations, *or* some degree of INAPPROPRIATENESS in – or indirectly suggested by – exactly one of these interpretations. A linguistic joke, on the other hand, is defined by a configuration of linguistic elements, involving notions such as phonetic similarity, segmentation into words, etc. It has no comparable component of delivery mechanism.

Both linguistic and propositional jokes rely on background information, may include extraneous information and may contain supplementary content/structure.

Hence, a full account of the components of a particular joke would consist of a statement of whether it was a linguistic or propositional joke, coupled with further details, as follows.

Linguistic jokes:

- a statement of which joke-definition is involved;
- an indication of how the parts of that definition are instantiated by particular structures (substrings, etc.).

Propositional jokes:

- the delivery mechanism used, and any relevant details (in particular, which text segment forms the punchline);

- *either*: an indication of which of the two interpretations provided by the delivery mechanism is linked to the INAPPROPRIATE content, and what that content is,
- *or*: a note that CONTRAST is the central factor, and an indication of what parts of the two interpretations constitute the CONTRAST.

Both classes:

- a note of any background information needed, indicating for each part whether it is accommodatable or not;
- a note of any supplementary aspects of the text;
- a note of any extraneous content.

It is important to note that the above summary of joke structure is an attempt to generalize over just the material so far studied. It is not yet a substantive theory of joke structure, and makes no predictive claims. It is an outline of the regularities observed in some jokes, and could be seen as a provisional draft of a theory of joke structure (as discussed in [Section 2.14](#)).

This checklist can be illustrated with a simple example, an informal analysis of (28).

(28) What is grey, has four legs, and a trunk?

A mouse on vacation.

This is a propositional joke, using the discourse-coherence misunderstanding delivery mechanism with punchline *A mouse on vacation*. The post-punchline interpretation implies the INAPPROPRIATE notion that a mouse takes a trunk on holiday. The only background required is basic knowledge of the English language (including word meanings). The image of a mouse carrying a trunk may be indirect supplementary content. The question–answer format contributes supplementary structure. There is no extraneous content.

Similarly, we can summarize a linguistic joke such as (29).

(29) Why did the cookie cry?

Because its mother had been a wafer so long.

The second sentence (the answer) forms a linguistically normal paradigmatic pun relative to a context made up of the first sentence (question). The phonetically similar text strings are (orthographically) *a wafer* and *away for*. The personification of edible pastries may constitute supplementary INAPPROPRIATE content. The question–answer format contributes supplementary structure. There is no extraneous content.

These analyses may seem rather thin and uninformative, but there are two points to be made. First, we can write such terse summaries because we have already given detailed definitions of the central structuring devices (here, discourse-coherence misunderstanding and linguistically normal paradigmatic pun). These then carry most of the import of the analysis. Second, we have set out to address only the *abstract syntax* of jokes, and hence our analyses are simply a structural characterization of the components of the joke. A structural analysis is a basis for further investigation, not an exhaustive statement of every aspect of a joke.

13

Future directions

We suggest some ways in which the work here could be extended.

13.1

Summary so far

The discussion in previous chapters are a few steps in a long-term strategy for explicating the mechanisms of humour.

We have:

- offered relatively precise definitions of a few classes of jokes, although for propositional jokes the formality extends only to what we have referred to as the *delivery mechanism*, leaving other aspects relatively unformalized;
- given sketches, of varying degrees of detail, of certain other classes, including certain sorts of puns;
- discussed some of the pragmatic or discourse devices that are used in jokes;
- considered some other joke classes, including linguistically abnormal paradigmatic puns, orthographic puns, and riddles;
- surveyed recent work on computational modelling of verbally expressed humour, suggesting that a fruitful methodological paradigm is developing in that subfield;
- offered some informal speculations, based on the findings of earlier chapters, on what the main components of a simple joke might be. In the next few sections, we will assess how our proposed programme could be pursued further.

13.2

Further analysis

As noted above, we have given descriptions of only a few joke classes, and these descriptions are mostly rather sketchy. If we are to develop empirically based generalizations, which can then become the basis of theoretical

hypotheses about joke structure, then more data has to be analysed: more instances of the joke classes we have considered, further classes of joke, and – as needed – further aspects of joke structure. Also, the class definitions we have drafted here need to be refined and developed. As noted in our presentation, we have glibly delegated to ‘further research’ the exact definitions of some central constructs, both linguistic (some of the primitives of [Appendix A.1](#)) and more humour-related (principally INAPPROPRIATENESS). These definitions will not materialize of their own accord.

13.3 Greater formality

Although one of our aims was to achieve greater formality than past work on verbally expressed humour, we have made only slight progress in that direction. Many of our discussions (e.g. Chapters [11](#), [12](#)) have been very informal. Even where these discussions may have shown a genuine pattern in the data, the ideas need to be formalized more thoroughly. It would be particularly interesting to study the issue of ‘distorted logic’ as used in jokes such as (34) or (35), to discover if there are any formalizable statements regarding humorous logic.

13.4 More rigorous methodology

In discussing example jokes, we have generally relied on our own intuitions about which texts are or are not jokes, and have assumed that these judgements are not problematic (cf. Attardo (2001: Section 1.6.3)). Although this is exactly comparable to the working practice, for many years, of generative linguists manipulating grammaticality judgements, it would be desirable to develop some slightly more objective and genuinely scientific way of corroborating/falsifying hypotheses.

13.5 Psychological modelling

One route to more thorough testing of hypotheses would be through the framework of experimental psychology. This would necessitate the formulation of hypotheses which had psychological consequences, in particular testable consequences. At the moment, the ideas set out in the earlier chapters have little or no psychological content. Although it might be possible, and extremely interesting, to develop these ideas into some form of cognitive model, such extensions are still in the future.

13.6 Computer modelling

As noted in [Chapter 10](#), there are the beginnings of a standard, and suitably thorough, methodology to guide the building and evaluation of

joke-generating software. It would seem, therefore, that computer testing would be a suitable way to test the internal coherence and the empirical accuracy of joke-class descriptions. However, in the field of humour there is as yet little overlap between theory and computer implementation. Even descriptive work, as illustrated in our earlier chapters or in Attardo (2001), is quite separate from the small-scale software projects of recent years. More sophisticated computer modelling is held up by two factors: first, theories of humour are rarely stated with sufficient detail or formality to allow implementation; second, processing of humorous forms beyond simple puns requires the program to have various other capabilities or resources, such as being able to produce/interpret natural language, being capable of subtle and flexible inferences, and having a vast store of knowledge about the real world.

13.7 From 'what' to 'why'

The programme of work advocated in this volume is aimed at pinning down just *what* makes a joke, and (eventually) *what* is funny. That leaves as a conceptually separate, but extremely difficult, question: *why* are these forms funny? That is, in the longer term, once the descriptive work has mapped out the territory more clearly, a theory is needed which will provide a reason for the funniness, or jokehood, of the structures uncovered by the empirical explorations. It would be particularly desirable if this theory were actually explanatory with respect to humour, in the sense that it showed that the patterns of funniness arose from some general principles regarding human behaviour (e.g. perceptual, cognitive, neural), or resulted from some evolutionary factors. This brings us back to our remarks in [Section 1.2](#). There is certainly no shortage of difficult questions.

Appendix A

Formal definitions

A.1

Linguistic constructs

In this section we will set out our assumptions about the linguistic entities available as building blocks for describing the mechanisms of jokes. These are not properly part of a theory of humour, but are external to such a theory. As stated in [Section 3.1](#), we are not putting forward a linguistic theory or a model of language processing. Rather, we are specifying what constructs we need from a linguistic theory in order to state our generalizations about joke structure as elegantly as possible. Also, many of these constructs are *primitive* as far as the humour analyses are concerned, in the sense that no definition is offered for them. They are simply unanalysed atomic notions, whose import is left to the reader's intuition, or to a suitable linguistic theory. The informal intended meaning of the terms should be relatively clear from the names chosen and earlier discussion. Any fuzziness is either irrelevant to the exploration of humour, or else has to be resolved empirically by further study of the data (i.e. jokes).

The constructs listed here are assumed (explicitly or implicitly) in various chapters.

Alphabets

There is a phonetic alphabet PHON (a set of atomic symbols), and any string over this alphabet (i.e. an element of PHON*) is known as a *phonetic string*. There is an orthographic alphabet LETTER (a set of atomic symbols), and any string over this alphabet (i.e. an element of LETTER*) is known as a *written string*. The term *text string* will be used to cover both phonetic strings and written strings. There is a similarity measure for phonetic strings, in the form of a relation SIMILAR, such that $(t, t') \in \text{SIMILAR}$ (also written $t \sim t'$) means that they are 'sufficiently similar' in some sense. Although we shall often invoke the notion of phonetic similarity, particularly when discussing *puns* ([Chapter 9](#)), it is as yet unclear what counts as similarity between strings for humorous purposes, and this abstract relation is a provisional measure pending further empirical work. It might be useful

to posit a mapping ALIKE from $\text{PHON}^* \times \text{PHON}^*$ to $[0,1]$, along with some threshold value ϵ ($0 < \epsilon < 1$), such that $t \sim t' \text{ iff ALIKE}(t, t') > \epsilon$. However, such further detail does not seem to be relevant to any humour matters.

There is a partial mapping PRONOUNCE from LETTER^* to the powerset of PHON^* (i.e. written strings to sets of phonetic strings). This indicates how a written word should be pronounced, and is partial because not every sequence of letters can be spoken, and involves sets of phonetic strings because there may be more than one pronunciation. There is a partial mapping SPELL from PHON^* to the powerset of LETTER^* (i.e. phonetic strings to sets of written strings). This indicates how a sound may be spelt and involves sets of written strings because there may be more than one spelling.

The following notation will allow us to write down some expressions more tersely: given strings U, T_1, T_2 where T_1 is a substring of U , write $U[T_1|T_2]$ for U with T_1 replaced by T_2 .

Syntax

There is a set SYN of *syntactic categories*; these are, informally, labels such as ‘noun’, ‘verb’, etc. All other syntactic information is bundled up into the *text-meaning mapping*. A text-meaning mapping M takes a string T and returns an ordered set $(LM, <)$ (possibly empty) of pairs, where each pair (sem, T_i) consists of a *literal meaning* (see below, under Semantics) and a final substring T_i of T . Informally, M consumes an initial substring of T (leaving a remainder T_i) and converts this initial substring into a representation of its meaning. There is a set of these pairs to allow for ambiguity (at any level), and the ordering ‘ $<$ ’ represents OBVIOUSNESS: the various possible segmentations are ordered according to their relative OBVIOUSNESS. Also, $(LM, <)$ always has a maximal element, being the most OBVIOUS reading. A string T is *grammatically well formed* if M applied to T either contains an element (sem, ϵ) where ϵ denotes the empty string, or contains an element (sem, T_i) where T_i is grammatically well formed. That is, a well-formed string is segmentable into a sequence of substrings each of which is suitable for translation into semantic form.

As will be defined below, a text-meaning mapping takes one further argument, representing the current semantic context.

Semantics

The unit of meaning will be the *proposition*, which can be thought of as a fact or hypothesis or statement, and which could be formalized as a well-formed formula in a suitable logic. For the moment, this is just a primitive notion. *Inconsistency* between propositions is represented as a property which

can be true of a set of propositions. This does not play a central role in our definitions (in particular, it is not identified with humour-creating ‘incongruity’), but is merely a condition imposed on various proposition-sets.

A text, or even an initial substring of a text, will – if grammatically well formed in our sense – be regarded as conveying an item called a *situation*, which can be thought of informally as a collection of propositions, some of which are more prominent than others. More formally, a situation is defined in terms of the operations that can be applied to it and the relations into which it can enter. There is a function ALL-PROPOSITIONS which, applied to a situation, returns a consistent set of propositions (intuitively, all the statements which are ‘believed’ in that situation). Another function SALIENT-PROPOSITIONS returns a subset of those returned by ALL-PROPOSITIONS (intuitively, the currently prominent propositions in that situation).

There is a set of *concepts* (also primitive), and a binary relation MENTIONS between propositions and concepts, indicating which concepts are explicitly part of the content of a given proposition.

There is also a relation ASSOC between concepts, which is symmetric and reflexive. This is intended to capture the kind of associative linkage that has been discussed in the literature on semantic nets, whereby (for example) the semantic form of a word like *nurse* would be linked to the meanings of words like *doctor*, *hospital*, *medicine*, etc. (that is, items in a similar *semantic field* (Lyons 1977b: 583)).

This allows a concept *C* to be *salient* in a situation *S*: this is the case if there is a proposition $P \in \text{SALIENT-PROPOSITIONS}(S)$ such that $P \text{ MENTIONS } C$. Also, a concept *C* can be *linked* to a situation, if there is a concept C_j such that C_j is salient in *S* and $(C, C_j) \in \text{ASSOC}$.

Situations can also be operated on by inference, accommodation and information update (see below).

A literal meaning, as produced by the text-meaning mapping, consists of a *pair* of sets of propositions. The first one (which may be empty) represents *prerequisites* for that constituent (i.e. presupposed facts which must be true in the situation for that item to be coherent in the discourse), and the second set represents the *core meaning* of the sentence (i.e. what is being asserted, queried or ordered by the sentence). The prerequisites are intended to include the notion of *presupposition*, interpreted very broadly (see [Section 3.6](#)).

We can now define more precisely the operation of a text-meaning mapping, introduced in the Syntax passage above. A text-meaning mapping will take two arguments: as well as the text substring, there will be a *situation* (see above), whose presence models the way in which current knowledge or recently mentioned concepts may influence the translation into meaning.

Lexicon

A lexical entry is a quadruple $\langle Phon, Ortho, Syn, Sem \rangle$ where $Phon \in \text{PRONOUNCE}(Ortho)$ is a phonetic string, $Ortho \in \text{SPELL}(Phon)$ is a written string, and Syn is a syntactic category. Sem is a *lexical semantic item*; we shall not specify what such an entity is, beyond stating that there is an operation which can act upon it: LEXCONCEPTS applied to a lexical semantic item will return a set of concepts. A *lexicon* is a set of lexical entries. A lexicon L induces a (partial) mapping LEX_L^P from phonetic strings to sets of lexical entries, where $LEX_L^P(w) = \{E \mid E = \langle w, Ortho, Syn, Sem \rangle\}$, and similarly LEX_L^W from written strings, defined by $LEX_L^W(w) = \{E \mid E = \langle Phon, w, Syn, Sem \rangle\}$. The mapping is partial because not every phonetic fragment corresponds to a word with a syntactic label or a meaning. Where the lexicon maps a string to a non-singleton set, that string is deemed to have several lexical entries (i.e. to be ambiguous). Given a lexicon L , a *lexical analysis* of a text string T is a finite sequence $((W_1, E_1), \dots, (W_n, E_n))$ such that $W_1 \dots W_n = T$, and for $1 \leq i \leq n$, either $E_i \in LEX_L(W_i)$ (if $LEX_L(W_i)$ is defined) or $E_i = \emptyset$, the empty set (otherwise). (Here $LEX_L = LEX_L^P$ or LEX_L^W , depending on whether T is phonetic or written.) Where $((W_1, E_1), \dots, (W_n, E_n))$ is a lexical analysis of a string T , we say that T *realizes* $E_1, \dots, E_n$, and that $(T, ((W_1, E_1), \dots, (W_n, E_n)))$ is a *lexically analysed string* (relative to the lexicon). A lexical analysis $((W_1, E_1), \dots, (W_n, E_n))$ of a string T is said to be *shorter* than a lexical analysis $((W'_1, E'_1), \dots, (W'_m, E'_m))$ of T , iff $n < m$. (Intuitively, the shorter analysis must group together the parts of T into larger portions, somewhere in the string.) A lexical analysis $((W_1, E_1), \dots, (W_n, E_n))$ is said to be *linked* to a situation S if any of the concepts appearing in the lexical semantic items (via LEXCONCEPTS) in any of the lexical entries $E_1, \dots, E_n$ are linked to S .

Given an alphabet (with a similarity relation), and a lexicon for strings over that alphabet, there are various possible conditions that can be phrased which can be true or false of a lexeme, a string or some combination of these. For example, if the syntactic label NOUN occurs in the lexicon, then there is an associated predicate *noun* which can be applied to a lexeme; the similarity of two text strings t_1, t_2 can be stated in predicated form as *similar*(t_1, t_2); if we define a relation SYNONYM (such that two lexical entries are synonyms if their semantic items are identical), then there is an associated predicate *synonym* which can be applied to two lexemes; if a string T realizes a lexical sequence $E_1, \dots, E_m$, then there is a predicate, which we can call *surface-form*, such that *surface-form* $(\langle E_1, \dots, E_m \rangle, T)$. (Strictly, there are two variants, *surface-form*^P for phonetic strings and *surface-form*^W for written strings.) In this way, choice of an alphabet and lexicon indirectly defines a collection of predicates (unary and binary) which can be used to state conditions on text strings and lexemes. We shall call these the *lexically definable predicates*; they are a rough approximation to the type

of information that is typically used by current computational pun-generators ([Chapter 10](#)).

Our version of the lexicon is in some ways less sophisticated than that typically discussed within linguistics. It represents an association between information and surface texts, but the manner of computing this association is of no interest. Therefore, it is not essential to include any form of morphological processing to ensure that various forms of a verb are all related to the same lexical entry in some way. In our abstract version, *walk* and *walks* are separate entries. Explicating the relation between these forms is delegated to linguistics.

Inference

We are not concerned with the exact formal nature of the reasoning process, so we shall simply assume the existence of an *inference operator* *I* which encapsulates all the ‘filling in of gaps’ which goes on during text understanding (see [Section 8.1.1](#)). This operator takes a situation and produces a set of possible results, each result being a situation (a different possible set of conclusions from the given set of initial facts). This set of results is ordered, with a maximal element, to represent how OBVIOUS each result is (relative to the others), the maximal element being the most OBVIOUS interpretation.

Relativity

To capture the fact that the status of a text as a joke and its degree of funniness are relative to some body of knowledge – cultural knowledge, social context, personal beliefs of the joke-audience, etc. – we will include in our formalization a structured entity called a *text-interpreter*. This is not intended in any way as a realistic model of a human listener; rather, it is a bundling together of the miscellaneous factors relative to which a text (particularly, a joke) is to be interpreted. A *text-interpreter* will be taken as consisting of the following components:

- Linguistic knowledge. This component will contain all the sets and relations listed above under Syntax, Semantics, and Lexicon, in particular the text-meaning mapping *M*.
- Beliefs. This is a situation *PB*, such that ALL-PROPOSITIONS (*PB*) represents knowledge/beliefs about the world, including cultural assumptions, social prejudices, etc.

It is arguable that the inference mapping, *I*, should form part of the text-interpreter, to make the inference relation relative to the audience and context. However, for the moment we have left it out, thereby assuming

that formal inference is general; idiosyncratic logical connections would therefore have to be represented as explicit statements, just as material implications in conventional propositional logic can give the effect of specific logical steps.

Accommodation and updating

We shall approximate discourse-understanding by assuming that each most OBVIOUS literal meaning produced by M is in some way assimilated into the existing context. We shall treat the two parts of the literal meaning – prerequisites and core meaning – separately. An abstract operator, ‘ $\hat{\cdot}$ ’ will check the prerequisites against the ongoing situation, and (if possible) assimilate the conditions into a revised situation. If ‘ $\hat{\cdot}$ ’ succeeds, we will say that the situation *can accommodate* the prerequisites. Similarly, we shall represent information update or belief revision using an *update* operator, ‘ $\cdot$ ’. That is, we delegate to others (logicians, discourse theorists, etc.) the task of stating how new information is normally assimilated during text understanding. The update operator will act on the core meaning of the text segment and the current situation, to produce a revised situation.

A similar, but subtly different, primitive is a relation SUITABLE between situations (as defined above) and text strings; this embodies the contextual appropriateness of [Section 3.9](#). Although it might seem that there should be some axioms relating ‘ $\hat{\cdot}$ ’ SUITABLE and ‘prediction’ (see below), these are distinct notions. Accommodation and prediction refer more to transient effects in the assimilation of the content of an utterance; contextual appropriateness concerns the ‘fit’ between utterance and surroundings, regardless of any hurdles that had to be crossed in the interpretation of the utterance. The only connection may be that an utterance which cannot be accommodated cannot be interpreted (see below) and so cannot be deemed to be SUITABLE. However, we do not make use of any such inter-relations in our joke analyses.

Text interpretation

The various building-blocks established above allow us to sketch a basic model of how a text string is transformed into a set of propositions. We shall assume that the text T is grammatical, and hence can be partitioned (by the text-meaning mapping) into n substrings (essentially, constituents), each convertible into a set of literal meanings.

For each constituent in turn there are (logically) several stages: mapping (via M) to a set of possible literal meanings, selection of the most OBVIOUS meaning, checking of prerequisites, absorption of core meaning, inference, and selection of the most OBVIOUS result of the inference. The process

is repeated from this new state for the remaining portion of text. Initially, M acts on T (which we can regard as T_1) starting from PB (the situation from the text-interpreter), computing $M(T_1, PB)$, which will be an ordered set $(LM_1, <)$ of candidate literal meanings each paired with a remainder substring.

Each element of LM_1 contains a literal meaning of the form (FC_1^j, CM_1^j) , a pair consisting of prerequisites and core meaning. The next step is to select $((FC_1^{max}, CM_1^{max}), T_2^{max})$, the maximum element of $(LM_1, <)$, which contains the most OBVIOUS literal meaning (and the remainder substring corresponding to that meaning).

If PB can, via $\tilde{\text{accommodate}}(FC_1^{max})$, then inference can take place, starting from $((PB \odot FC_1^{max}) \oplus CM_1^{max})$. The result of this will be of the form $(R_1, <)$ where R_1 is a set of situations and ' $<$ ' indicates the OBVIOUSNESS-ordering on situations. The maximum element from this set, S_1 , is then the interpretation so far of the text, and can be used as an input to the interpretation of T_2 (i.e. in place of PB). In this way, interpretation proceeds segment-by-segment, with OBVIOUSNESS, prerequisites and inference involved at each step. (It might be helpful to assume that M is sufficiently subtle that it can take account of the relationship between FC_i and S_{i-1} , to the extent that FC_i^{max} is unaccommodatable in S_{i-1} only if there is no alternative meaning of the textstring which can be accommodated in S_{i-1} ; that is, FC_i^{max} will normally be accommodatable unless *all* the available literal meanings are unaccommodatable).

We can define this more precisely and less procedurally as follows.

DEFINITION: INTERPRETATION SEQUENCE

Let T be a text. Relative to a text-interpreter (M, PB) an *interpretation sequence* for T consists of a sequence $T_1, \dots, T_n$ of final substrings of T and a sequence of situations $S_0, \dots, S_n$ such that:

- (a) $S_0 = PB$;
- (b) $T_1 = T$;
- (c) for $1 \leq i \leq n$, there is an element $((FC_i^{j(i)}, CM_i^{j(i)}), T_i^{j(i)})$ of $M(T_i, S_{i-1})$ such that:
 - (i) $T_i^{j(i)} = T_{i+1}$ for $1 \leq i < n$;
 - (ii) $T_i^{j(n+1)} = \Lambda$;
 - (iii) S_{i-1} can accommodate $FC_i^{j(i)}$;
 - (iv) S_i is an element of $I((S_{i-1} \tilde{\text{accommodate}} FC_i^{j(i)}) \cup CM_i^{j(i)})$;

If, in addition to these conditions, for every $1 \leq i \leq n$, $((FC_i^{j(i)}, CM_i^{j(i)}), T_i^{j(i)})$ is the maximum element of $M(T_i, S_{i-1})$ and S_i is the maximum element of $I((S_{i-1} \tilde{\text{accommodate}} FC_i^{j(i)}) \cup CM_i^{j(i)})$, then $(\langle T_1, \dots, T_n \rangle, \langle S_0, \dots, S_n \rangle)$ is said to be the *most obvious interpretation sequence* for T .

DEFINITION: INTERPRETATION

Let T be a text, and $(\langle T_1, \dots, T_n \rangle, \langle S_0, \dots, S_n \rangle)$ be an interpretation sequence for T (relative to some text-interpreter (M, PB)). Then the situation S_n is said to be an *interpretation* for T . If $(\langle T_1, \dots, T_n \rangle, \langle S_0, \dots, S_n \rangle)$ is the most obvious interpretation sequence for T , then S_n is *the most obvious interpretation* for T .

Any interpretation which is not the most obvious interpretation is said to be a *less obvious* interpretation.

Notice that there is no requirement that the content of the situations increase gradually as the text is interpreted, so there could be changes of facts at a step in the sequence, if the accretion of information with “is” is non-monotonic.

Prediction

We build semantic prediction into our text-interpretation model as a primitive term: the *prediction relation* is a relation between interpretation sequences and sets of proposition.

DEFINITION: SATISFYING A PREDICTION

Let T be a text of the form $T_1 \bar{T}_1 \bar{q}$ and $(\langle T_1, \dots, T_{n-1} \rangle, \langle S_0, \dots, S_n \rangle)$ be an interpretation sequence for T (relative to some text-interpreter (M, PB)). Let P be a set of propositions such that $((\langle T_1, \dots, T_{n-1} \rangle, \langle S_0, \dots, S_{n-1} \rangle), P)$ are in the prediction relation. Then S_n *satisfies the prediction* P if $P \models \text{SALIENT-PROPOSITIONS}(S_n)$.

A.2**Joke structures****A.2.1****Information presentation in propositional jokes****DEFINITION: PUNCHLINE CONFLICT**

A text T has *punchline conflict* with punchline $T_1 \bar{q}$ iff:

- (a) $T = T_1 \bar{T}_1 \bar{q}$ where $T_1 \bar{q}$ is non-empty;
- (b) $S_1 \bar{q}$ is the most obvious interpretation of $T_1 \bar{q}$
- (c) $M(T_1 \bar{q} S_1 \bar{q})$ has a maximum element $((FC^{max}, CM^{max}), \bar{q})$;
- (d) $S_1 \bar{q}$ cannot accommodate FC^{max} .

DEFINITION: PUNCHLINE RESOLUTION

A text T has *punchline resolution* with punchline T_i iff:

- (a) $T = T_i \hat{\cup} S$ where T_i is non-empty;
- (b) S is the most obvious interpretation of T_i ;
- (c) $M(T_i \hat{\cup} S)$ has a maximum element $((FC^{max}, CM^{max}), \hat{\Delta})$;
- (d) there is a less obvious interpretation S_j of T_i which can accommodate FC^{max} .

DEFINITION: DISCOURSE-COHERENCE MISUNDERSTANDING

A text T constitutes a *discourse-coherence misunderstanding*, relative to a text-interpreter (M, PB) , iff T has punchline conflict and punchline resolution both with the same punchline T_i

DEFINITION: PUNCHLINE REVISION

A text T has *punchline revision* iff the following hold:

- (a) $T = T_i \hat{\cup} T_j$ where T_i is non-empty;
- (b) the most obvious interpretation sequence for T_i is $(\langle T_1, \dots, T_{n-1} \rangle, \langle S_0, \dots, S_{n-1} \rangle)$;
- (c) SALIENT-PROPOSITIONS applied to the core meaning part of the maximum element of $M(T_i \hat{\cup} S_{n-1})$ contains a set of propositions P such that there is a sequence of situations $S'_1, \dots, S'_{n-1}$ such that $S'_{n-1} \neq S_{n-1}$ and, for some $k, 1 \leq k \leq (n-1)$:
 - (i) for $i < k, S_i = S'_i$;
 - (ii) S'_k is the maximum element of $((S'_{k-1} \odot FC_k^{max}) \oplus P) \oplus CM_k^{max}$ where $((FC_k^{max}, CM_k^{max}), T_{k+1})$ is the maximum element of $M(T_k, S'_{k-1})$;
 - (iii) for $k < i \leq n-1, S'_i$ is the maximum element of $((S'_{i-1} \odot FC_i^{max}) \odot CM_i^{max})$ where $((FC_i^{max}, CM_i^{max}), T_{i+1})$ is the maximum element of $M(T_i, S'_{i-1})$.

DEFINITION: PUNCHLINE INFERENCE

A text T has a *punchline inference* iff:

- (a) $T = T_i \hat{\cup} T_j$ where T_i is non-empty;
- (b) T_i has most obvious interpretation S_j ;
- (c) $M(T_i \hat{\cup} S_j)$ has a maximum element $((FC^{max}, CM^{max}), \hat{\Delta})$;
- (d) when SALIENT-PROPOSITIONS is applied to the maximum element of $I((S_j \hat{\cup} FC^{max}), CM^{max})$ this yields a set of propositions which has a subset A which:

- (i) is not a subset of CM^{max} (i.e. they must be inferred);
- (ii) is not a subset of SALIENT-PROPOSITIONS(S_j) (i.e. prior to the punchline, they were not salient).

DEFINITION: DOUBLE ENTENDRE PUNCHLINE

A text T has a *double entendre punchline* iff:

- (a) $T = T_1 \hat{\cup} S_j$ where T_1 is non-empty;
- (b) S_j is the most obvious interpretation of T_1 ;
- (c) $I((S_j \hat{\cup} FC^{max}) \cap CM^{max})$ has more than one member, where (FC^{max}, CM^{max}) is the maximum element of $M(T_1 \hat{\cup} S_j)$.

DEFINITION: EXPECTATION CLASH

A text T of the form $T_1 \hat{\cup} S_n$ has an *expectation-clash punchline* iff the most obvious interpretation sequence of T_1 is $\langle T_1, \dots, T_{n-1} \rangle$, $\langle S_0, \dots, S_{n-1} \rangle$, P is a set of propositions such that $(\langle T_1, \dots, T_{n-1} \rangle, \langle S_0, \dots, S_{n-1} \rangle, P)$ are in the prediction relation, the most obvious interpretation sequence for T is $\langle T_1, \dots, T_n \rangle, \langle S_0, \dots, S_n \rangle$, and S_n does not satisfy the prediction P .

DEFINITION: NARRATING A SITUATION

A text T is a *narrative resulting in S* iff S is the most obvious interpretation of T .

To formalize the suggestions of [Section 12.2.1](#), the following definition could be used.

Definition: a *delivery mechanism* is a mapping from a text-interpreter and a text T to a pair of situations (S, S_j) such that S is an interpretation of an initial substring of T (not necessarily proper) and S_j is an interpretation of T .

Then we could specify the situation-pairs as follows. For discourse-misunderstanding, they would be S_j from the definition of punchline conflict and the maximum element of $I((S_j \hat{\cup} FC^{max}) \cap CM^{max})$ from the definition of punchline resolution. For punchline revision, they would be S_{n-1} and S'_n in that definition. For punchline inference they would be S_j and the maximum element of $I((S_j \hat{\cup} FC^{max}) \cap CM^{max})$ from that definition. For double entendre they would be two members of $I((S_j \hat{\cup} FC^{max}) \cap CM^{max})$, one being the maximal. For a simple narrative, the resulting interpretation S must be based on an interpretation sequence $(\langle T_1, \dots, T_n \rangle, \langle S_0, \dots, S_n \rangle)$ where $S_n = S$; the two situations would then be S_{n-1} and S_n .

A.2.2

Linguistically normal paradigmatic puns

For the purposes of this definition, we will assume that what we called a ‘context’ in [Chapter 9](#) can be modelled by our formal object *situation*.

DEFINITION: LINGUISTICALLY NORMAL PARADIGMATIC PUNS

Given a lexicon L , a text string U constitutes a *linguistically normal paradigmatic pun* relative to a situation C iff there are lexically analysed strings (T, P) , (T_i, P_i) relative to L , such that:

- (a) T is a substring of U ;
- (b) $T_i \neq T$;
- (c) either (C, U) or $(C, U[T_i|T])$ is *SUITABLE* (i.e. one of these strings is linguistically appropriate to C);
- (d) if $T = T_i$ then:
 - (i) $P_i \neq P$;
 - (ii) either P_i is shorter than P , or P_i is linked to C ;
- (e) if $T_i \neq T$, then T_i forms a complete and recognizable linguistic unit (e.g. a complete word in L or an established phrase).

A.2.3

Knock-knock jokes

DEFINITION: KNOCK-KNOCK JOKES

Given a lexicon L , a text string U constitutes a *knock-knock joke* iff U is of the form $T_1 T_2 T_3 T_4 T_5 T_6$ and:

- (a) $T_1 = \text{knock knock}$;
- (b) $T_2 = \text{who's there}$;
- (c) T_3 is a personal name (in L);
- (d) $T_4 = \text{who}$;
- (e) $T_5 T_6$ constitutes a well-formed sentence;
- (f) $T_3 \neq T_5$.

We have segmented the text in this way (separating T_1 and T_2 , for example), to allow the further distinction between *interactive* delivery – $T_1, T_3, T_5 T_6$ delivered by the joke teller, and $T_2, T_3 T_4$ spoken by the joke-audience – and the *passive* reception of the joke – in which the audience reads or hears U in its entirety ([Section 11.1](#)).

Appendix B

Examples

- (1) What do you get if you cross the Atlantic with the *Titanic*?
About half-way.
- (2) 'My brother's a naval surgeon.'
Wow, they do specialise nowadays, don't they?'
- (3) What do you get if you cross a sheep with a kangaroo?
A woolly jumper.
- (4) What's black and dangerous and sits in a tree?
A crow with a machine gun.
- (5) Fashion is a form of ugliness so intolerable that we have to alter it every six months.
- (6) Do you believe in clubs for young people?
Only when kindness fails.
- (7) What is black and white and /r/ all over?
A newspaper.
- (8) Why do birds fly south in winter?
It's too far to walk.
- (9) Jean Harlow kept calling Margot Asquith by her first name, or kept trying to: she pronounced it *Margot*. Finally Margot set her right. 'No, no, Jean. The *t* is silent as in Harlow.'
- (10) There are 10 kinds of people in the world – those who understand binary numbers and those who don't.

- (11) TOO MUCH SEX makes you shortsighted.
- (12) 'Too much sex [*voice drops to a whisper*] makes you hard of hearing.'
- (13) Airline passenger: Where does this door go to-o-o-o ...
- (14) What did one sheep say to the other?
I love ewe.
- (15) He may not have been actually disgruntled, but he was certainly far from grunted.
- (16) Do you believe in clubs for young people?
Only when kindness fails, my friend.
- (17) 'Do you believe in clubs for young people?' someone asked W. C. Fields.
Only when kindness fails,' replied Fields.
- (18) 'Is the doctor at home?' the patient asked in his bronchial whisper. 'No,' the doctor's young and pretty wife whispered in reply. 'Come right in.'
- (19) Why did the antelope?
Nobody gnu.
- (20) Why did the antelope?
Nobody gnu.
- (21) What goes 'Pieces of nine! Pieces of nine!'
A parrotty error.
- (22) In New York, someone is knocked down by a car every two minutes. He's getting pretty fed up about this.
- (23) This was a mythical beast with the head of a lion and the body of a lion, but the body of *another* lion.
- (24) Diner: Waiter! There's a fly in my soup!
Waiter: Please don't shout so loudly – everyone will want one.
- (25) A lady went into a clothing store and asked 'May I try on that dress in the window?'
Well,' replied the sales clerk doubtfully, 'don't you think it would be better to use the dressing room?'

- (26) Postmaster: Here's your five-cent stamp.
Shopper (with arms full of bundles): Do I have to stick it on myself?
Postmaster: Nope. On the envelope.
- (27) An ancient, wizened man accosts a lady of the night and inquires as to her rates. She replies, \$5 on the floor, \$10 on the couch and \$15 in bed.' As he hands the hooker \$15, she remarks, Okay, once in bed,' to which he objects, 'No, three times on the floor.'
- (28) What is grey, has four legs, and a trunk?
A mouse on vacation.
- (29) Why did the cookie cry?
Because its mother had been a wafer so long.
- (30) Why did the elephant sit on the marshmallow?
Because he didn't want to fall into the hot chocolate.
- (31) Why do elephants paint their toenails red?
So they can hide in cherry trees.
- (32) If your son flunks out of school and is illiterate and anti-social, what can he grow up to be?
An Italian policeman.
- (33) Sitting on the side of the highway waiting to catch speeding drivers, a State Police Officer sees a car puttering along at 22mph. He thinks to himself, 'This driver is just as dangerous as a speeder!' So he turns on his lights and pulls the driver over. Approaching the car, he notices that there are five old ladies – two in the front seat and three in the back. The driver, obviously confused, says to him, 'Officer, I don't understand, I was doing exactly the speed limit! What seems to be the problem?' 'Ma'am,' the officer replies, 'You weren't speeding, but you should know that driving slower than the speed limit can also be a danger to other drivers.' 'Slower than the speed limit? No sir, I was doing the speed limit exactly ... twenty-two miles an hour!' the old woman says a bit proudly. The State Police Officer, trying to contain a chuckle, explains to her that '22' is the route number, not the speed limit. A bit embarrassed, the woman grins and thanks the officer for pointing out her error. 'Oh, thank you,' she says, 'It's a good thing you didn't see us on Route 119.'
- (34) O'Riley was on trial for armed robbery. The jury came out and announced, 'Not guilty.'
Wonderful,' said O'Riley, 'does that mean I can keep the money?'

- (35) Fat Ethel sat down at the lunch counter and ordered a whole fruit cake. 'Shall I cut it into four or eight pieces?' asked the waitress. 'Four,' said Ethel, 'I'm on a diet.'
- (36) How many Poles does it take to screw in a light bulb?
Five. One to hold the light bulb and four to turn the table he's standing on.
- (37) How many Poles does it take to screw in a light bulb?
Five. One to hold the light bulb and four to look for the right screw-driver.
- (38) Who supports Gorbachev? Oh, nobody. He is still able to walk on his own.
- (39) Why does a donkey eat thistles?
Because he's an ass.
- (40) What do you get when you cross a horse with a donkey?
A mule.
- (41) Why did the elephant wear red socks?
Because his green ones were being washed.
- (42) Why did the elephant wear blue socks?
Because his yellow ones were being washed.
- (43) 'Why is your dog wearing brown boots?'
Because his black ones are being mended.'
- (44) Why do firemen wear red braces?
To keep their trousers up.
- (45) Why do firemen wear strong helmets?
To protect their heads.
- (46) Why do birds fly south in winter?
It's too cold in the north.
- (47) Why did the boy take a hammer to the school?
Because it was the day they broke up.
- (48) What did Mickey Mouse get for Christmas?
A Dan Quayle watch.

- (49) Notice at a swimming pool: we don't swim in your toilet – please don't pee in our pool.
- (50) A man in his fifties goes to the doctor and says, “Doctor, I've got a problem. You see, when I was younger I always used to get erections that I couldn't bend with my hand. Now though, I can bend every erection I get. What I want to know is, am I getting stronger or weaker?”
- (51) God goes to the doctor and says, “Doc, I've got a problem. You see, I used to be able to make stones that were so heavy I couldn't lift them. But now I *can't* make a stone that I can't lift. The question is, am I getting more or less omnipotent?”
- (52) Why did the chicken cross the road?
It wanted to get to the other side.
- (53) Do you know the reason why the chicken decided to cross the road?
Because it wanted to get to the other side.
- (54) The reason the chicken crossed the road is that it wanted to get to the other side.
- (55) Why did the turtle cross the road?
It wanted to get to the other side.
- (56) Why did the chicken eat an octagonal-headed worm?
Because it was hungry.
- (57) Why did the chicken cross the road?
Nothing ventured, nothing gained.
- (58) Why did the chicken cross the road?
He saw a blonde hen on the other side.
- (59) How many Poles does it take to screw in a light bulb?
Five. One to take his shoes off, get on the table, and screw in the light bulb and four to wave the air deodorants to kill his foot odor.
- (60) A man walks into the front door of a bar. He is obviously drunk and staggers up to the bar, seats himself on a stool and, with a belch, asks the bartender for a drink. The bartender politely informs the man that it appears he has already had plenty to drink and that he could not be served additional liquor. The bartender offers to call a cab for him. The drunk is briefly surprised, then softly scoffs,

grumbles, climbs down from the bar stool and staggers out the front door. A few minutes later, the same drunk stumbles in the side door of the bar. He wobbles up to the bar and hollers for a drink. The bartender comes over and – still politely if not more firmly – refuses service to the man and again offers to call a cab. The drunk looks at the bartender for a moment angrily, curses, and shows himself out the side door, all the while grumbling and shaking his head. A few minutes later, the same drunk bursts in through the back door of the bar. He plops himself up on a bar stool, gathers his wits, and belligerently orders a drink. The bartender comes over and emphatically reminds the man that he is drunk and will be served no drinks. He then tells him that he can either call a cab or the police immediately. The surprised drunk looks at the bartender and in hopeless anguish cries, ‘Man! How many bars do you work at?’

(61) What’s the difference between an elephant and a watermelon?

(I don’t know.)

You’d be a fine one to send to the store for a watermelon.

(62) Russian officers in an Eastern European country go to a tavern. They order beer. The waiter places coasters on the table and serves the beer. Later they order another round. The waiter returning with the beer finds no coasters. ‘OK,’ he tells himself, ‘these are collectors,’ and puts down another set of coasters. When the third round is ordered and brought out, there are again no coasters. Angry, the waiter puts the beer down on the table, but places no more coasters. One of the Russian officers protests: ‘What’s this? No more crackers?’

(63) A blonde woman gives birth to a child with bright red hair. ‘Oh,’ says the doctor, ‘Father a redhead, is he?’ ‘No idea,’ says the woman, ‘he didn’t take his hat off.’

(64) One prostitute said to another, ‘Can you lend me ten dollars until I get back on my back?’

(65) John and his wife Mary were having a shower together in their upstairs bathroom when the doorbell rang. Mary heard the bell, got out of the shower, wrapped a towel around her, went downstairs, and opened the door. Their neighbor Charlie looked at her from the doorway, and said, ‘Oh. I see that I got you out of the shower. Sorry about that.’ ‘That’s all right,’ Mary said, ‘What do you want?’

Not too much ... my goodness you have beautiful skin. It’s so pink from the shower. Mary, if I was to give you a hundred dollars, would you remove the towel from your upper body?’ Mary thought about it

for a minute, figured why not, for a hundred bucks, and removed the towel from her breasts.

Wow,' Charlie exclaimed, 'they are truly beautiful. Listen, for another hundred bucks would you consider taking the towel all the way off?'

Why not,' Mary thought, 'that's a lot of money,' and she dropped the towel completely to the floor. Charlie had a good look, complimented her again on her fine looking body, reached into his pocket, took out two hundred dollars, gave it to her, and left. As she got back upstairs and was getting back into the shower, John asked her who was at the door. 'Just Charlie,' she said, as she started to rub his back. 'Charlie, eh,' said John, 'Did he give you the two hundred dollars he owed me?'

(66) Peter decided to go skiing with his buddy, Bob. They loaded up Peter's station wagon and headed north. After driving for a few hours, they got caught in a terrible blizzard. They pulled in to a nearby farm house and asked the attractive lady of the house if they could spend the night.

I'm recently widowed,' she explained, 'and I'm afraid the neighbors will talk if I let you stay in my house.' 'Not to worry,' Peter said, 'we'll be happy to sleep in the barn.' Nine months later, Peter got a letter from the widow's attorney. He then went up to visit his friend Bob and said,

Bob, do you remember that good-looking widow at the farm we stayed at?' 'Yes, I do.' 'Did you happen to get up in the middle of the night, go up to the house and have sex with her?' 'Yes, I have to admit that I did.' 'Did you happen to use my name instead of telling her your name?' Bob's face turned red and he said, 'Yeah, I'm afraid I did.' 'Well, thanks! She just died and left me everything!'

(67) An English bishop received the following note from the vicar of a village in his diocese:

Milord, I regret to inform you of my wife's death. Can you possibly send me a substitute for the weekend?'

(68) It so happens that if there is any institution which is not susceptible to any improvement whatsoever, it is the House of Peers.

(69) Save time and cut fingers with a parsley mincer.

(70) Imported Bedspreads by Glomar: Leaf design embroidered on fine polyester voile. Machine washable in champagne.

(71) One more drink and I'll be under the host.

(72) Isn't modern technology wonderful? I remember the excitement when we were the first family in our street to have cordless pyjamas.

- (73) A distinguished scientist was observing the heavens through the huge telescope at the Mt. Wilson Observatory. Suddenly he announced, 'It's going to rain.' 'What makes you think so?' asked his guide. 'Because,' said the astronomer, still peering through the telescope, 'my corns hurt.'
- (74) A couple give each other, as simultaneous surprise presents, tickets for two particularly sought-after concerts. The husband smugly relates the devious sequence of contacts he used to obtain his tickets, and asks how his wife acquired hers. She replies, 'I phoned the box office.'
- (75) Several centuries ago, the Pope decreed that all the Jews had to leave Italy. There was, of course, a huge outcry from the Jewish community, so the Pope offered a deal. He would have a religious debate with a leader of the Jewish community. If the Jewish leader won the debate, the Jews would be permitted to stay in Italy. If the Pope won, the Jews would have to leave. The Jewish community met and picked an aged Rabbi, Moishe, to represent them in the debate. Rabbi Moishe, however, could not speak Latin and the Pope could not speak Yiddish. So it was decided that this would be a 'silent' debate. On the day of the great debate, the Pope and Rabbi Moishe sat opposite each other for a full minute before the Pope raised his hand and showed three fingers. Rabbi Moishe looked back and raised one finger. Next, the Pope waved his finger around his head. Rabbi Moishe pointed to the ground where he sat. The Pope then brought out a communion wafer and chalice of wine. Rabbi Moishe pulled out an apple. With that, the Pope stood up and said, 'I concede the debate. This man has bested me. The Jews can stay.' Later, the Cardinals gathered around the Pope, asking him what had happened. The Pope said, 'First I held up three fingers to represent the Trinity. He responded by holding up one finger to remind me that there was still one God common to both our religions. Then I waved my finger around me to show him that God was all around us. He responded by pointing to the ground to show that God was also right here with us. I pulled out the wine and the wafer to show that God absolves us of our sins. He pulled out an apple to remind me of original sin. He had an answer for everything. What could I do?' Meanwhile, the Jewish community crowded around Rabbi Moishe, asking what happened. Well,' said Moishe, 'first he said to me, "You Jews have three days to get out of here." So I said to him, "Up yours". Then he tells me the whole city would be cleared of Jews. So I said to him, "Listen here Mr. Pope, the Jews ... we stay right here!"' 'And then?' asked a woman. 'Who knows?' said Rabbi Moishe. 'We broke for lunch.'

- (76) A pair of suburban couples who had known each other for quite some time talked it over and decided to do a little conjugal swapping. The trade was made the following evening, and the newly arranged couples retired to their respective houses. After about an hour of bedroom bliss, one of the wives propped herself up on her elbow, looked at her new partner and said 'Well, I wonder how the boys are getting along.'
- (77) How do you get two elephants in a Mini?
One in the front, one in the back.
- (78) How do you get /tʊ/ /weils/ in a Mini?
Across the Severn Bridge.
- (79) There was a boy standing on a corner selling fish. He was saying, 'Dam fish for sale, dam fish for sale.' A preacher walked up and asked why he was calling them dam fish. The kid said, 'I caught them at the dam, so they're dam fish.' The preacher bought some, took them home and asked his wife to cook the dam fish. His wife looked at him in bewilderment and said, 'Preachers aren't supposed to talk like that.' The preacher explained why they were dam fish, and she agreed to cook them. When dinner was ready and everyone was sitting down, the preacher asked his son to pass him the dam fish. His son replied, 'That's the spirit, dad. Pass the fucking potatoes!'
- (80) A rabbit goes into a butcher's shop and asks, 'Have you got any lettuce?' The butcher says, 'We don't sell lettuce here. You need the greengrocer's across the road.' The next day the rabbit comes into the shop and asks for some lettuce again. The butcher tells him, 'Look, I told you yesterday, we don't sell lettuce. You need the green-grocer.' The rabbit comes in the next day and asks the butcher again, 'Have you got any lettuce?' The butcher goes mad. He says, 'Look, I'm sick of this. How many times do I have to tell you I don't sell lettuce. If you come in here asking for lettuce, I'm going to nail your ears to the floor.' The next day the rabbit comes in and asks the butcher, 'Have you got any nails?' 'Nails? No.' 'Right,' the rabbit says, 'Have you got any lettuce?'
- (81) Why doesn't a railway engine sit down?
Because it has a tender behind.
- (82) Come forth, lazarus! And he came fifth and lost the job.
- (83) A sociology professor at the local college was fond of telling off-color jokes, a fact which embarrassed the young ladies in his class

tremendously.... Finally the girls got together and decided that at the next hint of an off-color joke they would get up *en masse* and leave the classroom. They sat there primly awaiting his next lecture. Sure enough, he began with the words, 'There is a terrible shortage of prostitutes in Singapore.' The girls looked at each other, rose, and started to leave. He called after them, 'Don't go now, girls. The boat doesn't leave for Singapore until Friday.'

(84) The tortoise has a protective shell.

(85) John found a shell on the beach.

(86) A shopper is walking along, and a leek falls from his shopping bag to the ground, unnoticed. Another shopper calls out, 'Hey! Your bag's leaking!'

(87) A shopper is standing with a bag containing groceries, including a prominently visible leek. A bystander remarks, 'You've got a leek in your bag.'

(88) Some South American stamps are un-bolivia-ble.

(89) A man walks along a city street carrying what is obviously a fender (in the British sense): the ornate metal surrounding for a hearth. The object is a square U-shape, and he is carrying it so that it surrounds him on three sides. A passer-by observes 'Home is where the hearth is.'

(90) Dorothy Stickney, invited to dinner with the President of the United States, was so nervous she quavered to him, 'What an imposing building the White House is! Who was the artichoke who designed it?'

(91) I'm only a prawn in the game.

(92) On departing for work, a person bids his family 'Olive oil.'

(93) Some people are talking about their friend, Dick. Someone says 'Dick's ears are so filthy. You could plant potatoes in them.' Another remarks, 'Hm ... Dick 'taters.'

(94) A shop selling greetings cards: 'Cardiology'.

(95) A hairdressing salon: 'Curl Up and Dye'.

- (96) The UK's *Guardian* newspaper once published an article about an ongoing period of prosperity in Albania, along with a picture of the capital, Tirana, captioned 'Tirana booms today'.
- (97) A minor football team known informally as 'Caley Thistle' (where *Caley* rhymes with *alley*) soundly defeats Celtic (then the top team in the country) in a major competition. The next day, a newspaper headline reads: 'Super Caley Go Ballistic, Celtic Are Atrocious'.
- (98) In the 1960s, a UK Labour politician was having difficulty with the trade union movement, led by Frank Cousins, and also with his own party members, who he was accustomed to addressing as 'brothers'. A political opponent remarked that his plans were being disrupted by 'his brothers and his Cousins'.
- (99) King Arthur had lots of knights who fared forth on coal-black charges to rescue beautiful maidens from dragons' clutches, but did you ever know that one of them was mounted on a St Bernard dog? His name was Sir Marmaduke, and he and the St Bernard performed many a deed of derring-do. One evening, however, they were caught in a torrential thunderstorm, and sought shelter at a nearby tavern. 'Reservation?' asked the room clerk. 'No,' admitted Sir Marmaduke. 'Sorry,' said the clerk, 'no room without a reservation.' It was at this point that he discovered that Marmaduke was sitting astride his faithful St Bernard. 'Hold on,' said the clerk, 'We'll have to find *something* for you. I wouldn't put out a knight on a dog like this.'
- (100) Once upon a time, many years ago, there was a chieftain in a remote tropical village who owned an old and battered throne of which he was very fond. One day, a visiting dignitary gave him a brand new and ornate throne, which the chieftain had to adopt immediately out of politeness.
- However, he could not bear to part with the old throne which had served him so well, so he stowed it away in the roof area of his grass hut, in case it should be useful in the future. Unfortunately the interior structure of his hut was too flimsy to support the weight of the large object, and it crashed through the grass ceiling, falling on the chieftain and killing him. The moral is that people who live in grass houses shouldn't stow thrones.
- (101) A bear walks into a bar in Baraboo, Wisconsin, and sits down. He bangs on the bar with his paw and demands a beer. The bartender approaches and says, 'We don't serve beer to bears in bars in Baraboo.' The bear, becoming angry, demands again that he be

served a beer. The bartender again tells him, 'We don't serve beer to bears in bars in Baraboo.' The bear, very angry now, says, 'If you don't serve me a beer, I'm going to eat that lady sitting at the end of the bar.' The bartender once again says 'Sorry, we don't serve beer to bears in bars in Baraboo.' The bear goes to the end of the bar, and as promised, eats the woman. He comes back to his seat and again demands a beer. The bartender says 'Sorry, we don't serve beer to bears in bars in Baraboo that are on drugs.' The bear says, 'I'm not on drugs.' The bartender says, 'Yes you are, that was a bar bitch you ate.'

(102) What kind of murderer has fibre?

A cereal killer.

(103) What do you get when you cross a murderer with a breakfast food?

A cereal killer.

(104) What do you use to colour an animal?

Hare dye.

(105) What did the python say to its victim?

I've got a crush on you.

(106) What do you call a naked bruin?

A grizzly bare.

(107) To colour an animal, you could use hare dye.

(108) When you cross a murderer with a breakfast food, you get a cereal killer.

(109) As the python said to its victim, I've got a crush on you.

(110) What do you call a depressed train?

A low-comotive.

(111) What do you call a bath tour?

A tub crawl.

(112) What does a near-sighted ghost wear?

Spooktacles.

(113) Teller: Knock, knock!

Audience: Who's there?

Teller: Noah.

Audience: Noah who?

Teller: Noah good place to eat?

- (114) A husband and wife are at a wedding reception. The husband says he needs to absent himself to perform a certain natural function'. The wife scolds him, 'You don't talk about that sort of function at this sort of function.'
- (115) It is better to be looked over than to be overlooked.
- (116) What do you call a strange market?
A bizarre bazaar.
- (117) What's the difference between leaves and a car?
One you brush and rake, the other you rush and brake.
- (118) What's the difference between a pretty glove and a silent cat?
One's a cute mitten, the other's a mute kitten.
- (119) What's the difference between a hairy dog and a painter?
One sheds his coat, the other coats his shed.
- (120) It was a lovely wedding. The cake was in tiers ... so was the bride.
- (121) The car was towing a caravan.
- (122) The committee will meet to discuss the meat quota.
- (123) The city schools teach a variety of skills.
- (124) I am interested in linguistic schools, you are interested in linguistic skills.
- (125) My car needed a repair, so I took it to a repair-shop.
- (126) My car had a dent, so I took it to a dentist.
- (127) Fox-hunting is 'the unspeakable in full pursuit of the uneatable'.
- (128) That's one small step for a man, one giant leap for mankind.
- (129) When is a door not a door?
When it is /əjær/.

(130) Why is coffee like the soil?

It is ground.

(131) Why is a goose like an icicle?

Both grow down.

(132) How's a nice girl like a sugary bird?

They're both sweet chicks.

(133) Sometimes a door is not a door. Sometimes, it is /əjær/.

(134) Coffee is like the soil. It is ground.

(135) An American soldier, serving in World War II, had just returned from several weeks of intense action on the German front lines. He had finally been granted R&R and was on a train bound for London. The train was very crowded, so the soldier walked the length of the train, looking for an empty seat. The only unoccupied seat was directly adjacent to a well dressed middle aged lady and was being used by her little dog. The war weary soldier asked, 'Please, ma'am, may I sit in that seat?' The English woman looked down her nose at the soldier, sniffed and said, 'You Americans. You are such a rude class of people. Can't you see my little Fifi is using that seat?' The soldier walked away, determined to find a place to rest, but after another trip down to the end of the train, found himself again facing the woman with the dog. Again he asked, 'Please, lady. May I sit there? I'm very tired.' The English woman wrinkled her nose and snorted, 'You Americans! Not only are you rude, you are also arrogant. Imagine!' The soldier didn't say anything else. He leaned over, picked up the little dog, tossed it out the window of the train and sat down in the empty seat. The woman shrieked and demanded that someone defend her and chastise the soldier. An English gentleman sitting across the aisle spoke up, 'You know, sir, you Americans do seem to have a penchant for doing the wrong thing. You eat holding the fork in the wrong hand. You drive your autos on the wrong side of the road. And now, sir, you've thrown the wrong bitch out of the window.'

(136) Why is coffee like the soil? When camping, the tent should be pitched on level ground.

(137) Why is a goose like an icicle? Fungus had started to grow down the front of the castle.

(138) Why is ebony like coal?

It is black.

(139) How can you spell *enemy* in 3 letters? N-M-E.

(140) Which word is always spelt wrongly?

Wrongly.

(141) A well-known wit claimed that he could make a pun on any subject. He was challenged to make a pun on 'the King'. He replied, 'The King is not a subject'.

(142) What do you call a person who puts you in touch with the spirit world?

A bartender.

(143) The medieval schoolmaster was determined to tame his unruly pupil. Therefore he threatened to flog the youth within an inch of his young life if he didn't measure up and do his work properly. He added, 'Thou hast better study thy lesson, Master Jonathan, or else methinks thou wouldst do well to get a weapon and defend thyself.' He bowed stiffly and prepared to leave.

Then the pupil made a *bow*.

(144) What do you call a blizzard forecast on April 1st?

A wind-up.

(145) Processed cheese represents the triumph of science over conscience.

(146) What makes men mean?

The letter a.

(147) Here lies Willie Longbottom, Aged 6. *Ars longa, vita brevis*.

(148) 'I'm bored,' yawned the piece of wood with a hole in it.

(149) A wife is cleaning out a grate; she pulls the metal tray out and it makes a very loud and unpleasant sound; her husband exclaims 'What a great noise!'

(150) How many Californians does it take to screw in a light bulb?

Two. One to screw it in and one to share in the experience.

(151) 'Turn up the heat,' said Tom coldly.

(152) 'I'll have the lamb,' said Tom sheepishly.

(153) 'This matter is opaque,' said Tom obscurely.

- (154) 'This object is dull,' said Tom pointedly.
(155) 'I am not a nun,' said Tom monkishly.
(156) 'I am not a flirt,' said Tom coquettishly.
(157) 'I stare at the flab,' said Tom fatuously.
(158) 'I notice the kitten,' said Tom categorically.
(159) What do you call a naked bruin?
A bare bear.
(160) What has an eye and can't see?
A potato.
(161) What kind of animal rides a catamaran?
A cat.
(162) The surgeon digs a garden. A doc yard.
(163) Joan hears wailing in the booth. Carrel singing.
(164) Joan charms a man in the sack. A male bag.
(165) Joan visits a surgeon in the garden. A doc yard.
(166) The sailor bears a stress. Pier pressure.
(167) Who broke the woman's hart?
The cruel deer-keeper.
(168) The performing lumberjack took a bough.
(169) Your mate Johnny is a hard-up deer-keeper. He really needs doe!
(170) A: I feel like a cup of tea.
B: You don't look like one!
(171) A: My ice-cream is dripping.
B: Oh that's funny I thought it was ice-cream!
(172) My father explained to me why birds fly south in winter.
It's because it's too far to walk.

- (173) I went to the zoo yesterday. I was surprised to see that the elephant was wearing red socks, but it turned out that his green ones were being washed.
- (174) A bartender is a person who puts you in touch with the spirit world.
- (175) A person who puts you in touch with the spirit world is called a bartender.
- (176) Sitting on the side of the highway waiting to catch speeding drivers, a State Police Officer sees a car puttering along at 22mph. He thinks to himself, 'This driver is just as dangerous as a speeder!' So he turns on his lights and pulls the driver over. Approaching the car, he notices that there are five old ladies – two in the front seat and three in the back – eyes wide and white as ghosts. The driver, obviously confused, says to him, 'Officer, I don't understand, I was doing exactly the speed limit! What seems to be the problem?' 'Ma'am,' the officer replies, 'You weren't speeding, but you should know that driving slower than the speed limit can also be a danger to other drivers.' 'Slower than the speed limit? No sir, I was doing the speed limit exactly twenty-two miles an hour!' the old woman says a bit proudly. The State Police Officer, trying to contain a chuckle, explains to her that '22' is the route number, not the speed limit. A bit embarrassed, the woman grins and thanks the officer for pointing out her error. 'But before I let you go, Ma'am, I have to ask ... Is everyone in this car OK ? These women seem awfully shaken and they haven't muttered a single peep this whole time,' the officer asks. 'Oh, they'll be alright in a minute officer. We just got off Route 119.'
- (177) The police have been investigating a group of people hell-bent on preventing the England cricket team from winning. But there is not much they can do about the selection committee.
- (178) In spring I am gay,
 In handsome array;
 In summer more clothing I wear;
 When colder it grows,
 I fling off my clothes;
 And in winter quite naked appear.
 (solution: a tree)
- (179) What has a mouth but cannot eat?
 A river.

- (180) What is a *cloak*?
The mating call of the Chinese frog.
- (181) What goes cluck-cluck bang?
A chicken in a minefield.
- (182) What goes ninety-nine clop?
A centipede with a wooden leg.
- (183) A chicken in a minefield goes cluck-cluck bang.
- (184) A centipede with a wooden leg goes ninety-nine clop.
- (185) An Englishman, a Scotsman and an Irishman were trying to get in to see the Olympics without tickets. The Englishman walked around the stadium and saw a pole lying on the ground and picked it up. He walked to the entrance and said, 'Peter. England. Pole vault.' The guards let him in without hesitation. The Scotsman found a manhole cover, carried it under his arm to the entrance, saying 'McGregor. Scotland. Discus throwing.' The guards let him in also. The Irishman was very frantic, since both his friends were now inside. He walked around the stadium and found a roll of barbed wire. He picked it up, walked to the entrance and said, 'Murphy. Ireland. Fencing.'
- (186) An Irishman was trying to get in to see the Olympics without tickets. He walked around the stadium and found a roll of barbed wire. He picked it up, walked to the entrance and said, 'Murphy. Ireland. Fencing.'
- (187) A man was trying to get in to see the Olympics without tickets. He noticed that competitors were allowed in if they gave their name and their chosen event, and were carrying their equipment. He walked around the stadium and found a roll of barbed wire. He picked it up, walked to the entrance and said, 'John Smith. Fencing.'
- (188) An Englishman, a Scotsman and an Irishman were looking for jobs with a construction company. The foreman decided to test them, so he set each of them to put in telegraph poles for a day. At the end of the day, he asked the Englishman how many he had put in. 'Twenty,' was the reply. Then he asked the Scotsman, who also replied 'Twenty'. Then it was the Irishman's turn and the foreman asked him how many telegraph poles *he* had been able to put up in that day. 'Three', said the Irishman. The foreman blinked, 'Three?' he repeated, puzzled. 'Only three? That's not very many. The English-

man and the Scotsman have both put in twenty each – how do you account for that?’ ‘Ah,’ replied the Irishman, ‘they may indeed have put in twenty each, but did you see how much they left sticking up out of the ground?’

(189) A company engaged in erecting telegraph poles recruited a new labourer, and the foreman set him to work on his own for the first day. At the end of the day, the foreman asked the new worker how many poles he had put in. ‘Three,’ was the reply. The foreman blinked, ‘Three?’ he repeated, puzzled. ‘Only three? That’s not very many. All the rest of the workers have put in about twenty each – how do you account for that?’ ‘Ah,’ replied the recruit, ‘they may indeed have put in twenty each, but did you see how much they left sticking up out of the ground?’

(190) Ethel Barrymore was told that a hated rival of many years’ standing had taken unto herself a new husband – and had made a full confession of past indiscretions to him. ‘What honesty! What courage!’ marvelled the reporter. ‘What a memory!’ added Miss Barrymore.

(191) Two Irishmen were digging a ditch directly across from a brothel. Suddenly, they saw a rabbi walk up to the front door, glance around and duck inside. ‘Ah, will you look at that?’ one ditch digger said, ‘What’s our world comin’ to when men of th’ cloth are visit in’ such places?’ A short time later, a Protestant minister walked up to the door and quietly slipped inside. ‘Do you believe that?’ the workman exclaimed, ‘Why, ’tis no wonder th’ young people today are so confused, what with the example clergymen set for them.’ After an hour went by, the men watched as the local parish priest quickly entered the whore house. ‘Ah, what a pity,’ the digger said, leaning on his shovel, ‘One of th’ poor lasses must be ill.’

(192) A wealthy grain merchant bought a new car and was enjoying his first ride in it no end until a motorcycle cop stopped him and suggested a visit to the local magistrate. ‘He was doing sixty,’ reported the cop. ‘Nonsense,’ declared the merchant. ‘I never got her up above forty.’ The merchant’s wife put in her two cents’ worth at this point. ‘He wasn’t going faster than twenty-five,’ she averred. A friend who had been riding in the back seat added, ‘I’d say we were virtually at a standstill when this officer came along.’ The magistrate threw up his hands and cried, ‘Stop right now – before you folks back into something.’

- (193) George S. Kaufman, after a self-made millionaire boasted, 'I was born into this world without a single penny.': When *I* was born, I owed twelve dollars.'
- (194) 'Excuse me, how long will the next bus be?'
About 12 metres.'
- (195) Guess what was on the TV last night?
A vase of flowers.

Glossary

This glossary lists some of the abbreviations and technical terms used in the main chapters. This does not include terms which are used only in the immediate vicinity of their definition, which are standard terms from disciplines such as linguistics, or which are used only within [Appendix A](#). The marker ‘(*)’ indicates a term defined, introduced, or used in a particular sense within this book; other terms are of more general currency. Terms in *italic font* indicate entries elsewhere in the glossary.

ABSURDITY (*) Used as a primitive technical term to denote a subclass of *inappropriateness* which involves a scenario which is somehow abnormal, weird, bizarre, etc. See [Chapter 5](#).

bisociation A term due to Koestler, describing a viewing of one concept from two perspectives. See [Chapter 4](#).

communicative alphabet (*) The set of basic symbols (e.g. letters of the alphabet) out of which a joke text is made. See [Section 3.2](#).

COMPATIBILITY (*) Used as a primitive technical term to denote the semantic and/or pragmatic relation between (the meaning of) a punchline and (the meaning of) a set-up which allows them to be integrated smoothly into an overall interpretation. See [Chapter 5](#).

CONFLICT (*) Used as a primitive technical term to denote the semantic and/or pragmatic relation between (the meaning of) a punchline and (the meaning of) a set-up which indicates to the audience that there is some discrepancy between the two. See [Chapter 5](#).

CONTRAST (*) Used as a primitive technical term to denote a semantic and/or pragmatic relation between two alternative interpretations of a text which leads to a humorous effect. See [Chapter 5](#).

defeasible An inference which is defeasible may be overridden by other information, and so has only a tentative or provisional status (Levinson 1983).

delivery mechanism(*) A formally definable way of presenting the information in a *propositional joke*. See Chapters [8](#) and [12](#), and [Appendix A](#).

- discourse coherence** A property possessed by a text in which consecutive segments (e.g. sentences) flow fluently. Here, given a rough formalization by defining each text segment to have *prerequisites* which must be compatible with the current context. See [Section 3.6](#).
- discourse coherence misunderstanding** (*) A subclass of the *FR delivery mechanism* in which *Conflict* is defined as lack of *discourse coherence* and *Compatibility* is defined as *discourse coherence*. See [Chapter 8](#) and [Appendix A](#).
- double entendre** (*) A type of *delivery mechanism* in which the *punchline* has more than one interpretation. See [Chapter 8](#) and [Appendix A](#).
- effect** (*) The result achieved by some linguistic device, as opposed to the *method* of achieving it. See [Section 3.8](#).
- extraneous** (*) Information or structuring within a joke which is not essential to the working of the joke. See [Section 7.3](#).
- forced reinterpretation** (*) A class of joke based on a misinterpretation in the *set-up* which is brought to light in the *punchline*. Abbreviated *FR*. See Chapters [5](#) and [8](#).
- FR** (*) Abbreviation for *forced reinterpretation*.
- frame** A complex knowledge structure embodying what is known about some familiar situation or event. See also *script*, and [Chapter 4](#).
- GTVH** The General Theory of Verbal Humour. See [Chapter 6](#).
- homophone** Two words are said to be homophones if they are pronounced in the same way, and one is then said to be a homophone of the other.
- illocutionary force** A broad classification of an utterance indicating what type of communication it is – question, statement, command, etc. (Levinson 1983).
- INAPPROPRIATENESS** (*) Used as a primitive technical term (along with ‘Inappropriate’) to denote some property of an interpretation of a text which leads to a humorous effect when the interpretation is presented in certain ways. Provisionally defined to be a disjunction of *Taboo* or *Absurd*. See [Chapter 5](#).
- inference** (*) Used as a primitive technical term to denote all the various reasoning processes that are used to make sense of linguistic meanings when a text is being interpreted. See [Section 3.9](#) and [Appendix A](#).
- jokehood** (*) The property of being a joke. Cf. *funniness*. See [Section 2.9](#).
- LA** Abbreviation within the *GTVH* for ‘Language’, the name of a knowledge resource and of a joke parameter.
- linguistic joke** (*) A joke which belongs to a formal class whose class definition crucially depends on linguistic form, such as phonetic similarity. Cf. *verbal humour*, *propositional joke*. See [Section 2.7](#).
- linguistically normal paradigmatic pun** (*) A subclass of *linguistic joke*. See [Section 9.3](#) and [Appendix A.2.2](#).

- literal meaning** (*) The semantic content of a piece of text, prior to any expansion using inference or contextual effects. See [Appendix A](#).
- LM** Abbreviation within the *GTVH* for *logical mechanism*.
- logical mechanism** A technical term within the *GTVH*, abbreviated *LM*. See [Chapter 6](#).
- malapropism** The inadvertent incorrect use of a word in place of a phonetically similar word. See [Section 9.3.2](#).
- method** (*) A mechanism used for achieving some linguistic *effect*. See [Section 3.8](#).
- narrative strategy** A technical term within the *GTVH*, abbreviated *NS*. See [Chapter 6](#).
- NS** Abbreviation within the *GTVH* for *narrative strategy*.
- OBVIOUSNESS** (*) Used as a primitive technical term (along with ‘Obvious’), to denote the property which determines which of two (or more) alternative interpretations of a text is noticed by the audience. See [Chapter 5](#).
- paradigmatic** A linguistic term involving comparison between two items which are potentially substitutable for each other in a textual context (Hudson 1971). It is also used to describe a pun where some string in the utterance is similar to some other string not in the utterance. Cf. *syntagmatic*.
- perlocutionary effect** An indication of the intended outcome of an utterance – persuasion, deception, amusement, etc. (Levinson 1983).
- phonetic string**(*) A finite sequence of symbols from a phonetic *communicative alphabet*. See [Appendix A](#).
- prerequisites** (*) Our formalization of *presuppositions* and any other assumed or required conditions for a text segment to fit fluently into the text as a whole. See [Appendix A](#).
- presupposition** Any proposition which must be true, or assumed to be true, in order for the given utterance to make sense (Levinson 1983).
- primitive** Applied to some concept which forms a building block for a theory, but which cannot be given a detailed definition in terms of other entities within the theory.
- propositional joke** (*) A joke which belongs to a formal class whose class definition depends wholly on the semantic content, and not in any way on the linguistic form. Similar to *referential humour*. See [Section 2.7](#).
- punchline** The final part of a joke, which creates the humorous effect. See *set-up*.
- punchline inference** (*) A type of *delivery mechanism* in which the *punchline* implies further consequences of the interpretation of the *set-up*. See [Chapter 8](#) and [Appendix A](#).

punchline revision (*) A subclass of the *FR delivery mechanism* in which the *punchline* implies a different rating of *Obviousness* for the interpretations of the *set-up*. See [Chapter 8](#) and [Appendix A](#).

referential humour Conventional term for jokes in which the content is important, but the detailed linguistic form is not relevant to the humour. Referential jokes are usually deemed to be translatable to other languages. Cf. *verbal humour*, *propositional joke*. See [Chapter 2](#).

script A complex knowledge structure embodying what is known about some familiar situation or event, prominent versions due to Schank and to Raskin, the latter as part of the *SSTH*. See also *frame*, and [Chapter 4](#).

script opposition A technical term within the *SSTH*, abbreviated *SO*. See [Chapter 6](#).

set-up The initial part of a joke, which prepares the way for the *punchline*.

SI Abbreviation within the *GTVH* for *situation*.

situation (1) A technical term within the *GTVH*, abbreviated *SI*. See [Chapter 6](#). (2) (*) Used as a primitive technical term to denote all the propositions held to be true or believed in a particular context. See [Appendix A](#).

SO Abbreviation within the *GTVH* for *script opposition*.

SSTH The Semantic Script Theory of Humour. See [Chapter 6](#).

story pun A joke consisting of a narrative followed by a line which is closely linked, semantically, to the story, but which is also phonetically similar to some (usually semantically unrelated) well-known phrase or saying. See [Section 9.3.4](#).

supplement (*) A class of information or structuring within a joke which, while not essential to the working of the joke, increases the humorous effect. Also known as 'supplementary information'.

syntagmatic A linguistic term involving a relationship between two items which both occur within some stretch of text (Hudson 1971). It is also used to describe a pun in which the utterance contains one or more similar parts. Cf. *paradigmatic*.

TA Abbreviation within the *GTVH* for *target*.

TABOO (*) Used as a primitive technical term to denote a subclass of *Inappropriateness* in which the scenario involves impropriety, subversive ideas, or any other matter which is not to be talked of openly and lightly. See [Chapter 5](#).

target A technical term within the *GTVH*, similar in meaning to 'butt' (of a joke), abbreviated *TA*. See [Chapter 6](#).

text (*) A written or spoken form, consisting of a finite sequence of symbols from the appropriate *communicative alphabet*. See [Appendix A](#).

theory-external (*) Applied to a construct whose definition should properly be supplied by some theory other than the one under consideration (here, that other theory might be linguistics or psychology). See [Chapter 2](#).

theory-internal (*) Applied to a construct whose definition forms part of the theory under consideration (here, a theory of jokes or of humour). See [Chapter 2](#).

type-token distinction All actual occurrences of some item (e.g. a word) are ‘tokens’, and the abstract item (or class into which it falls) is its ‘type’ (Cherry 1970).

ur-joke An abstract skeleton underlying a joke. See [Chapter 7](#).

verbal humour Usually taken to mean humour which is verbally expressed and which relies on particular properties of the language involved (e.g. homophony), therefore not being amenable to direct translation into another language. Cf. *referential humour*, *linguistic joke*. See [Chapter 2](#).

verbally expressed humour A general term for any humorous item, such as a joke, which is conveyed in written or spoken form, as opposed to a joke conveyed in some other medium, such as visually. See [Chapter 2](#).

written string (*) A finite sequence of symbols from a written *communicative alphabet*. See [Appendix A](#).

References

- Abelson, H. and Sussman, G. (1985) *Structure and Interpretation of Computer Programs*. MIT Electrical Engineering and Computer Science Series, Cambridge, MA: MIT Press.
- Alexander, R. J. (1997) *Aspects of Verbal Humour in English*. No. 13 in *Language in Performance*, Tübingen: Gunter Narr Verlag.
- Arnold, K., Gosling, J. and Holmes, D. (2000) *The Java Programming Language*, Reading, MA: Addison Wesley.
- Attardo, S. (1994) *Linguistic Theories of Humour*. No. 1 in *Humor Research*, Berlin: Mouton de Gruyter.
- (1997) 'The semantic foundations of cognitive theories of humor', *Humor: International Journal of Humor Research*, 4(10): 395–420.
- (1998) 'The analysis of humorous narratives', *Humor: International Journal of Humor Research*, 11(3): 231–60.
- (2001) *Humorous Texts: A semantic and pragmatic analysis*. No. 6 in *Humor Research*, Berlin: Mouton de Gruyter.
- (2002) 'Formalizing humor theory', in Stock, O., Strapparava, C. and Nijholt, A. (2002) (eds). and Chabanne, J.-C. (1992) 'Jokes as a text type', *Humor: International Journal of Humor Research*, 5(1/2): 165–75.
- and Mele, F. (2002) 'HAHAcronym evaluation and application opportunities', Deliverable D2, HAHAcronym: Humorous Agent for Humorous Acronyms, IST Programme, Future and Emerging Technologies, European Commission.
- and Raskin, V. (1991) 'Script theory revis(it)ed: joke similarity and joke representation model', *Humor: International Journal of Humor Research*, 4(3): 293–347.
- Attardo, D. H., Baltes, P. and Petray, M. J. (1994) 'The linear organization of jokes: analysis of two thousand texts', *Humor: International Journal of Humor Research*, 1(7): 27–54.
- Hempelmann, C. and Maio, S. D. (2002) 'Script oppositions and logical mechanisms: modeling incongruities and their resolutions', *Humor: International Journal of Humor Research*, 15(1): 3–46.
- Beattie, J. (1776) 'An essay on laughter, and ludicrous composition', in *Essays*, Edinburgh: William Creech. Reprinted by Garland, New York, 1971.
- Ben-Amos, D. (1976) 'Solutions to riddles', *Journal of American Folklore*, 89: 249–54.
- Berger, A. A. (1998) *An Anatomy of Humor*, Brunswick, NJ: Transaction Publishers. First published 1993.
- Bergsen, H. (1911) *Laughter: An essay on the meaning of the comic*, New York: Macmillan.

- Binsted, K. (1995) 'Using humour to make natural language interfaces more friendly', in Kitano, H. (ed.) *Proceedings of the IJCAI workshop on AI and Entertainment*, Montreal.
- (1996) *Machine Humour: An implemented model of puns*. Ph.D. thesis, University of Edinburgh, Scotland.
- and Ritchie, G. (1994) 'An implemented model of punning riddles', in *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, Seattle.
- and — (1996) 'Speculations on story puns', in Hulstijn, J. and Nijholt, A. (eds) (1996) *Proceedings of the International Workshop on Computational Humor*. No. 12 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.
- and — (1997) 'Computational rules for generating punning riddles', *Humor: International Journal of Humor Research*, 10(1): 25–76.
- and — (2001) 'Towards a model of story puns', *Humor: International Journal of Humor Research*, 14(3): 275–92.
- , Pain, H. and Ritchie, G. (1997) 'Children's evaluation of computer-generated punning riddles', *Pragmatics and Cognition*, 5(2): 309–58.
- Bobrow, D. G. and Winograd, T. (1977) 'An overview of KRL, a knowledge representation language', *Cognitive Science*, 1(1): 3–46.
- , Kaplan, R. M., Kay, M., Norman, D. A., Thompson, H. and Winograd, T. (1977) 'GUS, a frame-driven dialog system', *Artificial Intelligence*, 8: 155–73.
- Cerf, B. (1946) *Anything for a Laugh*, New York: Grosset and Dunlap.
- (1964) *Bennett Cerf's Vest Pocket Book of Jokes*, London: Hammond, Hammond and Company.
- Chapman, A. J. and Foot, H. C. (eds) (1976) *Humour and Laughter: Theory, research and applications* (1st edn), London: Transaction Publishers.
- and — (eds) (1977) *It's a Funny Thing, Humour*, Oxford: Pergamon Press.
- Charniak, E. and McDermott, D. (1985) *Introduction to Artificial Intelligence*, Reading, MA: Addison-Wesley.
- Cherry, C. (1970) *On Human Communication*, Cambridge, MA: MIT Press. First published 1957.
- Chiaro, D. (1992) *The Language of Jokes*, London and New York: Routledge.
- Chomsky, N. (1957) *Syntactic Structures*, The Hague: Mouton.
- (1965) *Aspects of The Theory of Syntax*, Cambridge, MA: MIT Press.
- (1975) *The Logical Structure of Linguistic Theory*, New York: Plenum Press. First published 1955.
- Clark, D. A. (1968) *Jokes, Puns and Riddles*, New York: Doubleday.
- Clark, M. (1970) 'Humor and incongruity', in Morreall, J. (ed.) (1987). *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press. Reprinted from *Philosophy*, 45: 20–32 (1970).
- Coltheart, M. (1981) 'The MRC psycholinguistic database', *The Quarterly Journal of Experimental Psychology*, 33(A): 497–505.
- Cornwell, D. and Hobbs, S. (1991) 'Defining the pun', in Bennett, G. (ed.) *Spoken in Jest*. Vol. 21 of Folklore Society Mistletoe Series, Sheffield, UK: Sheffield Academic Press.
- Crain, S. and Steedman, M. (1985) 'On not being led up the garden path: the use of context by the psychological syntax processor', in Dowty, D., Karttunen, L. and Zwicky, A. (eds) *Natural Language Parsing*, London: Cambridge University Press.
- Crystal, D. (1995) *The Cambridge Encyclopaedia of the English Language*, Cambridge: Cambridge University Press.

- Culler, J. (ed.) (1988) *On Puns: The foundations of letters*, New York: Basil Blackwell.
- Dascal, M. (1985) 'Language use in jokes and dreams: sociopragmatics vs. psycho-pragmatics', *Language and Communication*, 5(2): 95–106.
- De Palma, P. and Weiner, E. J. (1992) 'Riddles: accessibility and knowledge representation', in *Proceedings of COLING-92 (15th International Conference on Computational Linguistics)*, Nantes.
- Deckers, L. and Avery, P. (1994) 'Altered joke endings and a joke structure schema', *Humor: International Journal of Humor Research*, 7(4): 313–21.
- Dienhart, J. M. (1999) 'A linguistic look at riddles', *Journal of Pragmatics*, 31(1): 95–125.
- Dolitsky, M. (1983) 'Humor and the unsaid', *Journal of Pragmatics*, 7: 39–48.
- (1992) 'Aspects of the unsaid in humor', *Humor: International Journal of Humor Research*, 5(1/2): 33–43.
- Ertner, J. D. (1993) *Super Silly Animal Riddles*, New York: Sterling Publishing Co. Inc.
- Fauconnier, G. (1985) *Mental Spaces*, Cambridge, MA: MIT Press.
- (1996) 'Analogical counterfactuals', in *Spaces, Worlds, and Grammar*, Chicago: Chicago University Press.
- Fave, L. L., Haddad, J. and Maesen, W. A. (1976) 'Superiority, enhanced self-esteem, and perceived incongruity humour theory', in Chapman, A. J. and Foot, H. C. (eds) (1976). *Humour and Laughter: Theory, research and applications* (1st edn), London: Transaction Publishers.
- Fodor, J. D. and Ferreira, F. (eds) (1998) *Reanalysis in Sentence Processing*, Dordrecht: Kluwer Academic.
- Forabosco, G. (1992) 'Cognitive aspects of the humor process: the concept of incongruity', *Humor: International Journal of Humor Research*, 5(1/2): 45–68.
- Freud, S. (1966) *Jokes and their Relation to the Unconscious*, London: Routledge & Kegan Paul. First published 1905.
- Friedman, J. (1971) *A Mathematical Model of Transformational Grammar*, New York: American Elsevier.
- (1972) 'Mathematical and computational models of transformational grammar', in Meltzer, B. and Michie, D. (eds) *Machine Intelligence 7*, Edinburgh: Edinburgh University Press.
- Gazdar, G., Klein, E., Pullum, G. and Sag, I. (1985) *Generalized Phrase Structure Grammar*, Oxford: Blackwell.
- Giora, R. (1988) 'Victor Raskin: semantic mechanisms of humor', *Philosophia*, 18(4/7): 409–15.
- (1991) 'On the cognitive aspects of the joke', *Journal of Pragmatics*, 16(5): 465–85.
- Godkewitsch, M. (1976) 'Physiological and verbal indices of arousal in rated humour', in Chapman, A. J. and Foot, H. C. (eds) (1976) *Humour and Laughter: Theory, research and applications* (1st edn), London: Transaction Publishers.
- Goldstein, J. H. and McGhee, P. E. (eds) (1972) *The Psychology of Humor*, New York: Academic Press.
- Gruner, C. (1978) *Understanding Laughter: The workings of wit and humor*, Chicago: Nelson-Hall.
- (1997) *The Game of Humor*, New Brunswick, NJ: Transaction Publishers.
- Hartley, D. (1810) 'Excerpt from "Observations on Man"', in Morreall, J. (ed) (1987) *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press.
- Hayes, P. J. (1979) 'The logic of frames', in Metzing, D. (ed.) *Frame Conceptions and Text Understanding*, Berlin: Walter de Gruyter.
- (1987) 'A critique of pure treason', *Computational Intelligence*, 3(3): 179–85.

- Hazlitt, W. (1819) 'Excerpt from "Lectures on the English Comic Writers"', in Morreall, J. (ed) (1987) *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press.
- Heller, L. G. (1974) 'Toward a general typology of the pun', *Language and Style*, 7: 271–82.
- Hetzron, R. (1991) 'On the structure of punchlines', *Humor: International Journal of Humor Research*, 4(1): 61–108.
- Hobbes, T. (1840) 'Human nature', in Molesworth, W. (ed.) *The English Works of Thomas Hobbes of Malmesbury*, Vol. IV, London: John Bohn.
- Hobbs, J. (1990) *Literature and Cognition*. No. 21 in Lecture Notes, Stanford, CA: Centre for the Study of Language and Information.
- , Stickel, M., Appelt, D. and Martin, P. (1993) 'Interpretation as abduction', *Artificial Intelligence*, 63: 69–142.
- Hockett, C. F. (1977) 'Jokes', in *The View from Language*, Athens, GA: University of Georgia Press.
- Hofstadter, D. and Gabor, L. (1989) 'Synopsis of the workshop on humor and cognition', *Humor: International Journal of Humor Research*, 2(4): 417–40.
- Hudson, R. (1971) *English Complex Sentences*, Amsterdam: North-Holland.
- Hulstijn, J. and Nijholt, A. (eds) (1996) *Proceedings of the International Workshop on Computational Humor*. No. 12 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.
- Hutcheson, F. (1750) 'Excerpt from "Reflections Upon Laughter"', in Morreall, J. (ed) (1987) *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press.
- Jespersen, O. (1933) 'Punning – or allusive phrases in English', in *Selected Writings*, Tokyo: Sanyo. First published 1900.
- Johnson, R. (1976) 'Two realms and a joke: bisociation theories of joking', *Semiotica*, 16(3): 195–221.
- Kagan, J. (1967) 'On the need for relativism', *American Psychologist*, 22: 131–47.
- Kant, I. (1892) 'Excerpt from "Critique of Judgement"', in Morreall, J. (ed) (1987) *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press.
- Katz, B. (1993) 'A neural resolution of the incongruity and incongruity-resolution theories of humour', *Connection Science*, 5: 59–75.
- (1996) 'A neural invariant of humour', in Hulstijn, J. and Nijholt, A. (eds) *Proceedings of the International Workshop on Computational Humor*. No. 12 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.
- Keith-Spiegel, P. (1972) 'Early conceptions of humor: varieties and issues', in Goldstein, J. H. and McGhee, P. E. (eds) (1972) *The Psychology of Humor*, New York: Academic Press.
- Kelly, L. G. (1971) 'Punning and the linguistic sign', *Linguistics*, 66: 5–11.
- Kempson, R., Meyer-Viol, W. and Gabbay, D. (2001) *Dynamic Syntax: The flow of language understanding*, Oxford: Blackwell.
- Kimball, J. P. (1973) *The Formal Theory of Grammar*. Foundations of Modern Linguistics, Englewood Cliffs, NJ: Prentice-Hall.
- Kitagaki, I. (1990) 'A fuzzy determination method of generating a laugh and popularity/inferiority: "students' holiday" and "the lowest in running"', *Journal of Japan Society for Fuzzy Theory and Systems*, 2(1): 100–4. In Japanese.
- (1993) 'Extraction of identity on pronunciation concerning play-on-word and evaluation of a tentative software: aiming a word processor of human friendliness', Tech. rep. HC92–65, Institute of Electronics, Information and Communication Engineers of Japan. In Japanese.

- Koestler, A. (1970) *The Act of Creation*, London: Pan Books. First published by Hutchinson & Co, 1964.
- Langevin, R. and Day, H. I. (1972) 'Physiological correlates of humor', in Goldstein, J. H. and McGhee, P. E. (eds) (1972) *The Psychology of Humor*, New York: Academic Press.
- Larson, G. (1982) *The Far Side*, Kansas City, KA: Andrews and McMeel.
- (1983) *Beyond The Far Side*, Kansas City, KA: Andrews and McMeel.
- (1984) *In Search of The Far Side*, Kansas City, KA: Andrews and McMeel.
- (1985) *Bride of The Far Side*, Kansas City, KA: Andrews and McMeel.
- Lascarides, A. and Asher, N. (1993) 'Temporal interpretation, discourse structure and commonsense entailment', *Linguistics and Philosophy*, 16: 437–93.
- Latta, R. L. (1999) *The Basic Humor Process: A cognitive-shift theory and the case against incongruity*. No. 5 in *Humor Research*, Berlin: Mouton de Gruyter.
- Laurel, B. (1993) *Computers as Theatre*, Reading, MA: Addison Wesley.
- Lessard, G. and Levison, M. (1992) 'Computational modelling of linguistic humour: Tom Swifties', in *ALLC/ACH Joint Annual Conference*, Oxford.
- and — (1993) 'Computational modelling of riddle strategies', in *ALLC/ACH Joint Annual Conference*, Georgetown University, Washington, DC.
- and — (1995) 'Linguistic and cognitive under pinnings of verbal humour', in *International Cognitive Linguistics Association Conference*, Albuquerque. Available at <http://www.cs.queensu.ca/CompLing/publish.html> (accessed June 2002).
- and — (1997) 'Rule-governed word play and creativity', in *Proceedings of Mind II: Computational Models of Creative Cognition*, Dublin City University. Available at <http://www.cs.queensu.ca/CompLing/publish.html> (accessed June 2002).
- , and Venour, C. (2002) 'Cleverness versus funniness', in Stock, O., Strapparava, C. and Nijholt, A. (eds) (2002) *Proceedings of the April Fools' Day Workshop on Computational Humor*. No. 20 in *Twente Workshops of Language Technology*, Enschede, Netherlands: University of Twente.
- Levinson, S. (1983) *Pragmatics*, Cambridge: Cambridge University Press.
- Levison, M. and Lessard, G. (1992) 'A system for natural language generation', *Computers and the Humanities*, 26: 43–58.
- Lippman, L. G. and Dunn, M. L. (1995) "'You might summarize the pun research", said Tom abstractly'. Talk at 13th International Humor Conference, Birmingham.
- and — (2000) 'Contextual connections within puns: effects on perceived humor and memory', *Journal of General Psychology*, 127: 185–97.
- Loehr, D. (1996) 'An integration of a pun generator with a natural language robot', in Hulstijn, J. and Nijholt, A. (eds) (1996) *Proceedings of the International Workshop on Computational Humor*. No. 12 in *Twente Workshops on Language Technology*, Enschede, Netherlands: University of Twente.
- Lyons, J. (1977a) *Semantics: Volume 1*, Cambridge: Cambridge University Press.
- (1977b) *Semantics: Volume 2*, Cambridge: Cambridge University Press.
- McGhee, P. E. (1979) *Humor: Its origin and development*, San Francisco: W. H. Freeman and Company.
- and Goldstein, J. H. (eds) (1983) *Handbook of Humor Research*, New York: Springer-Verlag.
- , Ruch, W. and Hehl, F.-J. (1990) 'A personality-based model of humor development during adulthood', *Humor: International Journal of Humor Research*, 3: 119–46.

- McKay, J. (2000) 'Generation of idiom-based witticisms to aid second language learning'. Master's thesis, Division of Informatics, University of Edinburgh, Scotland.
- Maier, N. R. F. (1932) 'A gestalt theory of humour', *British Journal of Psychology*, 23: 69–74.
- Maranda, E. K. (1971) 'The logic of riddles', in Maranda, P. and Maranda, E. K. (eds) *Structural Analysis of Oral Tradition*, Philadelphia: University of Pennsylvania Press.
- Matthews, T. S. (1974) *Great Tom: Notes towards the definition of T. S. Eliot*, London: Weidenfeld & Nicolson.
- (2002) 'Generation of idiom-based witticisms to aid second language learning', in Stock, O., Strapparava, C. and Nijholt, A. (eds) (2002) *Proceedings of the April Fools' Day Workshop on Computational Humor*. No. 20 in Twente Workshops of Language Technology, Enschede, Netherlands: University of Twente.
- Metcalf, F. (1994) *The Penguin Dictionary of Jokes*, London: Penguin Books.
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., Miller, K. and Teng, R. (1990) 'Five papers on WordNet', *International Journal of Lexicography*, 3(4). Revised March 1993.
- Minsky, M. (1975) 'A framework for representing knowledge', in Winston, P. H. (ed.) *The Psychology of Computer Vision*, New York: McGraw-Hill.
- (1980) 'Jokes and the logic of the cognitive unconscious'. AI Memo 603, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, MA.
- (1986) *The Society of Mind*, London: Heinemann.
- Morkes, J., Kernal, H. K. and Nass, C. (1999) 'Effects of humor in task-oriented human-computer interaction and computer-mediated communication: a direct test of SRCT theory', *Human-Computer Interaction*, 14(4): 395–435.
- Morreall, J. (1983) *Taking Laughter Seriously*, Albany, NY: SUNY Press.
- (ed.) (1987) *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press.
- (1996) 'Review of "Taking Humour Seriously"', *Humor: International Journal of Humor Research*, 9: 401–6.
- Muir, F. and Norden, D. (1991) *The Utterly Ultimate 'My Word' Collection*, London: Mandarin Humour Classics.
- Nash, W. (1985) *The Language of Humour*. No. 16 in English Language, Harlow: Longman.
- Nerhardt, G. (1970) 'Humor and inclination to laugh: emotional reactions to stimuli of different divergence from a range of expectancy', *Scandinavian Journal of Psychology*, 11: 185–95.
- (1976) 'Incongruity and funniness', in Chapman, A. J. and Foot, H. C. (eds) (1976) *Humour and Laughter: Theory, Research and Applications* (1st edn), London: Transaction Publishers.
- Nijholt, A. (2002) 'Embodied agents: a new impetus to humor research', in Stock, O., Strapparava, C. and Nijholt, A. (eds) (2002) *Proceedings of the April Fools' Day Workshop on Computational Humor*. No. 20 in Twente Workshops of Language Technology, Enschede, Netherlands: University of Twente.
- Nobbs, D. (1986) *A Bit of a Do*, London: Methuen.
- Norrick, N. R. (1986) 'A frame-theoretical analysis of verbal humour: bisociation as schema conflict', *Semiotica*, 60(3/4): 225–45.
- (1989) 'Intertextuality in humor', *Humor: International Journal of Humor Research*, 2(2): 117–39.

- (2001) 'On the conversational performance of narrative jokes: toward an account of timing', *Humor: International Journal of Humor Research*, 14(3): 255–74.
- Oaks, D. D. (1994) 'Creating structural ambiguities in humor: getting English grammar to cooperate', *Humor: International Journal of Humor Research*, 7(4): 377–401.
- O'Mara, D. A. and Waller, A. (2003) 'What do you get when you cross a communication aid with a riddle?', *The Psychologist*, 16(2): 78–80.
- _____, and Todman, J. (2002) 'Linguistic-humour and the development of language skills in AAC', in *Proceedings of 10th Biennial Conference of the International Society for Augmentative and Alternative Communication*, Odense, Denmark.
- Palmer, J. (1991) *Potboilers: Methods, concepts and case studies in popular fiction*, London: Routledge.
- _____. (1994) *Taking Humour Seriously*, London: Routledge.
- Paolillo, J. C. (1998) 'Gary Larson's *Far Side*: nonsense? nonsense!', *Humor: International Journal of Humor Research*, 3: 261–90.
- Parsons, D. (ed.) (1953) *Can It Be True?*, London: Macdonald.
- Paulos, J. A. (1980) *Mathematics and Humor*, Chicago and London: University of Chicago Press.
- Pearce, M. T. and Wiggins, G. A. (2001) 'Towards a framework for the evaluation of machine compositions', in *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in the Arts and Sciences*, York.
- Pease, A., Winterstein, D. and Colton, S. (2001) 'Evaluating machine creativity', in Weber, R. and von Wangenheim, C. G. (eds) *Case-Based Reasoning: Papers from the Workshop Programme at ICCBR 01*, Vancouver.
- Pepicello, W. J. and Green, T. A. (1984) *The Language of Riddles*, Columbus: Ohio State University Press.
- _____, and Weisberg, R. W. (1983) 'Linguistics and humor', in McGhee, P. E. and Goldstein, J. H. (eds) (1983) *Handbook of Humor Research*, New York: Springer-Verlag.
- Raju, A. (1991) 'Decoding the comic', in Bennett, G. (ed.) *Spoken in Jest*. Vol. 21 of Folklore Society Mistletree Series, Sheffield: Sheffield Academic Press.
- Random House (1988) *Random House College Dictionary* (revised edn), New York: Random House.
- Raskin, V. (1985) *Semantic Mechanisms of Humor*, Dordrecht: Reidel.
- _____. (1987) 'The interdisciplinary field of humor research', *Semiotica*, 66(4): 441–9.
- _____. (1995) 'Logical mechanisms: the elusive knowledge resource of the joke'. Talk at 13th International Humor Conference, Birmingham.
- _____. (1998) 'From the editor', *Humor: International Journal of Humor Research*, 11(1): 1–4.
- _____, and Attardo, S. (1994) 'Non-literality and non-bona-fide in language: approaches to formal and computational treatments of humor', *Pragmatics and Cognition*, 2(1): 31–69.
- Redfern, W. (1984) *Puns*, New York: Blackwell.
- Reiter, R. (1978) 'On reasoning by default', in *Theoretical Issues in Natural Language Processing 2*, Urbana-Champaign: University of Illinois.
- _____. (1980) 'A logic for default reasoning', *Artificial Intelligence*, 13(1–2): 81–132.
- Ritchie, G. (1999) 'Developing the incongruity-resolution theory', in *Proceedings of the AISB Symposium on Creative Language: Stories and Humour*, Edinburgh.
- _____. (2000) 'Describing verbally expressed humour', in *Proceedings of the AISB Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*, Birmingham.

- (2001) 'Assessing creativity', in *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in Arts and Science*, York.
- (2003) 'The JAPE riddle generator: technical specification'. Informatics research report EDI-INF-RR-0158, School of Informatics, University of Edinburgh.
- Roberts, R. B. and Goldstein, I. P. (1977) 'The FRL manual'. AI Memo 409, AI Laboratory, MIT, Cambridge, MA.
- Rosch, E. and Mervis, C. (1975) 'Family resemblances: studies in the internal structure of categories', *Cognitive Psychology*, 7: 573–605.
- Ross, A. (1998) *The Language of Humour*, London: Routledge.
- Rothbart, M. K. (1976) 'Incongruity, problem-solving and laughter', in Chapman, A. J. and Foot, H. C. (eds) (1976). *Humour and Laughter: Theory, research and applications* (1st edn), London: Transaction Publishers.
- (1977) 'Psychological approaches to the study of humour', in Chapman, A. J. and Foot, H. C. (eds) (1977) *It's a Funny Thing, Humour*, Oxford: Pergamon Press.
- and Pien, D. (1977) 'Elephants and marshmallows: a theoretical synthesis of incongruity-resolution and arousal theories of humour', in Chapman, A. J. and Foot, H. C. (eds) (1977) *It's a Funny Thing, Humour*, Oxford: Pergamon Press.
- Ruch, W. (1992) 'Assessment of appreciation of humor: studies with the 3WD humor test', in Spielberger, C. D. and Butcher, J. N. (eds) *Advances in Personality Assessment: Volume 9*, Hillsdale, NJ: Lawrence Erlbaum.
- (ed.) (1996) 'Special issue: measurement approaches to sense of humor', *Humor: International Journal of Humor Research*, 9 (3/4).
- , Attardo, S. and Raskin, V. (1993) 'Toward an empirical verification of the general theory of verbal humor', *Humor: International Journal of Humor Research*, 6(2): 123–36.
- Russell, B. (1905) 'On denoting', *Mind*, 14: 479–93.
- Russell, S. and Norvig, P. (1995) *Artificial Intelligence: A modern approach*, Englewood Cliffs, NJ: Prentice-Hall.
- Rutter, J. F. (1997) *Stand Up as Interaction: Performance and audience in comedy venues*. Ph.D. thesis, Dept of Sociology, Salford University.
- Schank, R. and Abelson, R. (1977) *Scripts, Plans, Goals and Understanding*, Hillsdale, NJ: Lawrence Erlbaum.
- Schopenhauer, A. (1819) 'Excerpt from "The World as Will and Idea"', in Morreall, J. (ed) (1987) *The Philosophy of Laughter and Humor*, Albany, NY: SUNY Press.
- (1883) *The World as Will and Idea*, London: Trübner.
- Shapiro, S. C. (ed.) (1992) *Encyclopaedia of Artificial Intelligence* (2nd edn), New York: John Wiley.
- Shultz, T. R. (1972) 'The role of incongruity and resolution in children's appreciation of jokes and cartoons: an information-processing analysis', *Journal of Experimental Child Psychology*, 13: 456–77.
- (1974) 'Development of the appreciation of riddles', *Child Development*, 45: 100–5.
- (1976) 'A cognitive-developmental analysis of humour', in Chapman, A. J. and Foot, H. C. (eds) (1976). *Humour and Laughter: Theory, research and applications* (1st edn), London: Transaction Publishers.
- Sobkowiak, W. (1991) *Metaphonology of English Paronomasic Puns*. Vol. 26 of University of Bamberg Studies in English Linguistics, Frankfurt: Peter Lang.
- Steedman, M. (1973) *The Formal Description of Musical Perception*. Ph.D. thesis, University of Edinburgh.

- Steele, G. L. (1990) *Common Lisp: The language* (2nd edn), Maynard, MA: Digital Press.
- Stock, O. (1996) “‘Password Swordfish’: verbal humour in the interface”, in Hulstijn, J. and Nijholt, A. (eds) (1996) *Proceedings of the International Workshop on Computational Humor*. No. 12 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.
- and Strapparava, C. (2002) ‘Humorous agent for humorous acronyms: the HAHAAcronym project’, in Stock, O., Strapparava, C. and Nijholt, A. (eds) (2002) *Proceedings of the April Fools’ Day Workshop on Computational Humor*. No. 20 in Twente Workshops of Language Technology, Enschede, Netherlands: University of Twente.
- , and Nijholt, A. (eds) (2002) *Proceedings of the April Fools’ Day Workshop on Computational Humor*. No. 20 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.
- Suls, J. M. (1972) ‘A two-stage model for the appreciation of jokes and cartoons: an information-processing analysis’, in Goldstein, J. H. and McGhee, P. E. (eds) (1972) *The Psychology of Humor*, New York: Academic Press.
- (1977) ‘Cognitive and disparagement theories of humour: a theoretical and empirical synthesis’, in Chapman, A. J. and Foot, H. C. (eds) (1977) *It’s a Funny Thing, Humour*: Oxford: Pergamon Press.
- (1983) ‘Cognitive processes in humor appreciation’, in McGhee, P. E. and Goldstein, J. H. (eds) (1983) *Handbook of Humor Research*, New York: Springer-Verlag.
- Svebak, S. (1977) ‘Some characteristics of resting respiration as predictors of laughter’, in Chapman, A. J. and Foot, H. C. (eds) (1977) *It’s a Funny Thing, Humour*, Oxford: Pergamon Press.
- Takizawa, O., Yanagida, M., Ito, A. and Isahara, H. (1996) ‘On computational processing of rhetorical expressions – puns, ironies and tautologies’, in Hulstijn, J. and Nijholt, A. (eds) (1996) *Proceedings of the International Workshop on Computational Humor*. No. 12 in Twente Workshops on Language Technology, Enschede, Netherlands: University of Twente.
- Turner, R. (1984) *Logics for Artificial Intelligence*, Chichester: Ellis Horwood.
- van der Sandt, R. (1988) *Context and Presupposition*, London: Croom Helm.
- Veatch, T. (1998) ‘A theory of humour’, *Humor: International Journal of Humor Research*, 11(2): 161–215.
- Venour, C. (1999) ‘The computational generation of a class of puns’. Master’s thesis, Queen’s University, Kingston, Ontario.
- Webb, K. (ed.) (1978) *The Crack-a-Joke Book*, London: Puffin.
- Wilde, O. (1966) *Complete Works of Oscar Wilde*, London: Collins.
- Wilson, C. P. (1979) *Jokes: Form, content, use and function*, London: Academic Press.
- Wilson, R. C. and Kidd, F. C. (eds) (1999) *The MIT Encyclopaedia of the Cognitive Sciences*, Cambridge, MA: MIT Press.
- Winograd, T. (1983) *Language as a Cognitive Process. Volume I: Syntax*, Reading, MA: Addison-Wesley.
- Wittgenstein, L. (1967) *Philosophical Investigations* (3rd edn), Oxford: Blackwell. First published 1953.
- Wright, E. L. (1979) ‘Illusion and truth’, *Philosophy and Phenomenology*, 39: 402–32.

Index

- 3WD test [158](#)
 see also INC-RES component
- abbreviations *see* acronyms
- abduction [93](#), [94](#), [95](#)
- Abelson, H. [29](#)
- Abelson, R. [51](#), [52](#), [93](#)
- abstract syntax of jokes [9](#), [24](#), [108](#), [141](#), [185](#)
- absurd humour [48](#), [182](#)
- ABSURDITY [61](#)–2, [92](#), [93](#), [100](#), [102](#), [165](#), [177](#), [178](#)–9
- accessibility [63](#)
- accidental humour [100](#);
 see also malapropism
- accommodation [39](#), [184](#), [191](#), [194](#)–5, [196](#)
- acronyms [154](#)–5
- aggression theory of humour [7](#)–8, [94](#)
- Alexander, R. J. [2](#), [27](#), [63](#), [118](#), [134](#), [138](#), [139](#), [140](#)
- Alice's Adventures in Wonderland* [139](#)
- Allen, W. [41](#)
- alphabet:
 communicative [33](#), [123](#)–4, [192](#);
 orthographic [189](#);
 phonetic [189](#)
- ambiguity [7](#), [64](#), [91](#), [94](#), [168](#);
 absence from two-stage model [65](#), [67](#);
 in assumed model [45](#), [190](#);
 basis for misinterpretation [14](#), [30](#), [31](#), [33](#), [44](#), [61](#), [90](#), [92](#), [95](#), [99](#), [105](#), [106](#), [113](#), [160](#), [164](#), [165](#);
 classificatory criterion [83](#), [88](#), [130](#);
 in double entendre [99](#), [100](#);
 effect or method [43](#)–4;
 not necessary for a pun [115](#);
 in pun [68](#), [111](#)–12, [114](#), [115](#), [122](#), [131](#)–3, [137](#)
- ambiguity types [39](#)–42;
 focus [82](#), [90](#), [162](#);
 lexical [126](#), [192](#);
 lexical-syntactic [90](#);
 orthographic-lexical [43](#);
 orthographic-syntactic [43](#);
 phonetic-lexical [114](#), [115](#), [122](#), [132](#), [137](#), [161](#), [165](#);
 phonetic-syntactic [43](#);
 pronunciation [138](#);
 semantic-pragmatic [91](#), [160](#);
 semantic-referential [165](#)
- analogy [53](#);
 false [71](#)
- antonyms [154](#)–5;
 local [70](#), [73](#), [74](#), [112](#)
- Appelt, D. [95](#)
- Armstrong, N. [129](#)
- Arnold, K. [51](#)
- artificial intelligence [1](#), [4](#), [6](#), [7](#), [11](#)–12, [20](#), [73](#)
- Asher, N. [93](#)
- association, semantic [45](#), [124](#), [131](#), [135](#), [139](#), [152](#), [191](#)
- Attardo, D. H. [17](#), [35](#), [39](#)
- Attardo, S. [3](#), [4](#), [7](#), [10](#), [13](#), [14](#), [16](#), [17](#), [26](#), [29](#), [30](#), [35](#), [46](#), [47](#), [50](#), [55](#)–6, [60](#), [63](#), [67](#), [69](#)–80, [84](#), [85](#)–8, [99](#), [109](#), [110](#), [112](#), [113](#), [116](#), [139](#), [143](#)–4, [180](#), [181](#), [187](#), [188](#)
- Avery, P. [104](#)

background knowledge [38](#)–9, [100](#), [169](#), [183](#)–4, [193](#)

Baltes, P. [17](#), [35](#)

BBC [115](#), [165](#)

Beattie, J. [46](#), [47](#), [49](#)

Beckwith, R. [149](#), [154](#)

beliefs:

effect of audience's [39](#), [169](#), [180](#), [193](#);

speaker's [42](#);

systems of [64](#);

formal model [193](#)–4

Ben-Amos, D. [40](#)

Berger, A. A. [1](#), [182](#)
 Bergson, H. [49](#)
 binary arithmetic [31](#)–[2](#), [38](#)
 Binsted, K. [3](#), [22](#), [23](#), [40](#), [120](#)–[1](#), [122](#)–[4](#), [125](#), [128](#), [131](#), [138](#), [143](#), [147](#), [149](#), [150](#), [157](#)–[8](#)
 bisociation [51](#)–[4](#)
 Bobrow, D. G. [51](#)
 bottom-up analysis [3](#)
 Brown, A. [101](#)
 butts of jokes [71](#), [180](#);
 see also Target

 cartoons [8](#), [13](#), [57](#)–[8](#), [59](#), [84](#);
 Far Side [75](#)
 catch-phrases [181](#)
 Cerf, B. [31](#), [34](#), [102](#), [116](#), [120](#), [172](#), [173](#)
 Chabanne, J.-C. [16](#)
 Chapman, A. J. [4](#)
 Charniak, E. [12](#), [93](#)
 Cherry, C. [30](#)
 Chiaro, D. [4](#)
 Chomsky, N. [11](#), [12](#), [40](#)
 Cicero, M. T. [14](#)
 Clark, D. A. [42](#), [43](#)
 Clark, M. [48](#), [49](#)
 cleverness [153](#);
 see also ingenuity;
 intricacy
 cognitive rules [59](#), [65](#)–[7](#)
 cognitive science [1](#), [7](#)
 cognitive theory of humour [7](#), [19](#)
 Coltheart, M. [149](#)
 Colton, S. [158](#)
 Common Lisp [51](#)
 COMPATIBILITY [61](#)–[4](#), [95](#)–[6](#)
 computers:
 applications [6](#), [22](#)–[3](#);
 evaluation of models *see* evaluation of computer models;
 generation of jokes [20](#)–[33](#), [28](#), [133](#), [141](#), [143](#)–[58](#);
 modelling [3](#), [5](#), [11](#)–[12](#), [20](#)–[3](#), [186](#), [187](#)–[8](#);
 testing [20](#)–[2](#), [43](#), [141](#), [188](#);
 understanding of jokes [21](#)–[3](#)
 concepts:
 units of knowledge representation [45](#), [114](#), [119](#), [123](#), [128](#), [134](#), [141](#), [154](#), [156](#), [191](#), [192](#)
 CONFLICT [61](#)–[4](#), [94](#)–[6](#), [99](#)
 Conjecture [92](#), [95](#), [159](#), [164](#), [177](#)–[8](#)
 connectors [35](#)
 consistency, logical [95](#), [190](#)–[1](#)
 constructed items:
 definition [5](#);
 examples [37](#), [78](#), [81](#), [82](#), [114](#), [123](#), [129](#), [131](#), [133](#), [162](#), [168](#), [170](#), [171](#), [172](#)
 context linkage *see* linkage to context
 contexts for jokes [24](#), [109](#), [114](#)–[15](#), [117](#), [139](#), [184](#), [193](#);
 textual [118](#)–[24](#)
 contextual appropriateness [45](#), [124](#)–[5](#), [126](#), [131](#)–[5](#), [141](#), [194](#), [199](#)
 CONTRAST [61](#)–[4](#), [70](#), [101](#), [102](#), [103](#), [107](#), [141](#), [165](#), [173](#), [176](#), [177](#), [178](#), [183](#)–[4](#)
 controlled testing [20](#), [23](#), [82](#), [149](#), [150](#)–[1](#), [153](#), [155](#), [157](#)
 conundrums *see* riddles as puzzles
 core meanings [191](#), [194](#), [195](#)
 Cornwell, D. [111](#), [112](#), [160](#)–[1](#)
 Crain, S. [101](#)
 Crystal, D. [31](#), [32](#), [37](#)
 Culler, J. [110](#)
 cultural factors [16](#), [38](#), [114](#), [193](#)

Dartmouth Conference [11–12](#)
 Dascal, M. [60](#)
 data abstraction [29](#);
 see also encapsulation
 data structures [51–2](#), [72–3](#)
 Day, H. I. [11](#)
 De Maio, S. [69](#), [70](#), [72](#), [74](#), [75](#)
 De Palma, P. [63](#), [167](#)
 Deckers, L. [104](#)
 declarative accounts *see* procedural accounts vs. declarative
 delivery mechanisms [67](#), [176–8](#), [183–4](#), [186](#), [198](#)
 dependencies:
 discontinuous [44](#);
 unbounded [44](#)
 description-based jokes [167–9](#)
 descriptive theory of humour [10](#), [20](#), [24–7](#), [43](#)
 descriptivism [10](#)
 dialect [38](#), [140](#)
 dictionaries *see* lexicons
 Dienhart, J. M. [29](#), [30](#), [31](#), [32](#), [34](#), [37](#), [40](#), [81](#), [91](#), [116](#), [122](#), [136](#), [166–7](#)
 directed graphs [72](#), [73](#)
 discourse [159–74](#);
 coherence [39](#), [94](#), [95](#), [96–7](#);
 understanding [64](#), [194–5](#)
 discourse coherence misunderstanding [96](#), [176](#), [178](#), [184](#), [185](#), [197](#), [198](#)

discovery methods [35](#)
 disparagement [180](#)
 Dolitsky, M. [60](#), [104](#), [105](#)
 double entendre [98–100](#), [176](#), [198](#)
 dramatic tension [172](#), [181](#)
 Dunn, M. L. [136](#)
 dynamic syntax [17](#)
 effects (within an analysis) [42–4](#)

embarrassment [182](#)
 embellishment [181](#)
 empirical evidence [2](#), [18](#), [78](#), [79–80](#)
 encapsulation [157](#)
 epigrams [13](#)
 Ertner, J. D. [128](#)
 essential factors in humour [61](#), [180](#)
 essentialist theory of humour [10](#)
 evaluation of computer models [20–1](#), [23](#), [144](#), [145](#), [146](#), [149](#), [150–1](#)
 evolution of humour [2](#), [188](#)
 exaggeration [173](#), [182](#)
 expectation [17](#), [50](#), [65–7](#);
 clash [198](#);
 thwarted [94](#), [180](#)
 experimental design [157](#)
 experiments, possible [19–20](#), [162](#), [164](#), [187](#)
 explanation of humour [1](#), [25](#)
 extraneous information [88–9](#), [173](#), [179](#), [182](#), [183–4](#)

facilitation [171](#), [181](#)
 falsification [1](#), [8–9](#), [12](#), [18](#), [26](#), [69](#), [79–80](#), [187](#)
 Fauconnier, G. [93](#)
 Fave, L. L. [7](#)
 Fellbaum, C. [149](#), [154](#)
 Ferreira, F. [101](#)
 figures of speech [53](#)
 figure–ground reversal [74](#)
 focus (within script) *see* foregrounded elements
 Fodor, J. D. [101](#)
 Foot, H. C. [4](#)
 Forabosco, G. [48](#)
 forced reinterpretation:
 compared to two-stage [62–8](#);
 distinguished from puns [113](#), [136–7](#), [140](#);
 misleading description [165](#), [168–9](#);
 models [59–64](#), [90–7](#);
 need for other factors [107](#);
 potential equivalence class [83](#);
 question forms [159–63](#);
 SSTH as [60](#), [63](#), [70](#), [74](#);
 widely discussed [176](#)
 foregrounded elements [70](#), [73](#), [74](#)
 formal definitions [189–99](#)
 formality [5](#), [11–12](#), [182](#), [187](#), [188](#)
 FR *see* forced reinterpretation
 frames [51–4](#)
Frasier (TV show) [102](#)
 Freud, S. [4](#), [7](#), [47](#), [63](#)
 Friedman, J. [20](#)
 funniness:
 compared to jokehood [16](#), [18](#);
 computer jokes [21](#), [149](#), [153](#), [155](#), [157–8](#);
 effect of questions [164](#);
 effect of supplements [179](#), [180](#), [182](#);

human ratings [57](#), [163](#)–4;
 audience factors [100](#);
 not considered here [24](#), [39](#), [159](#), [175](#)–6, [188](#)
 funny stories *see* narrative jokes
 funny words [182](#)

Gabbay, D. [17](#)
 Gabor, L. [84](#)–5, [179](#)
 garden path LM [60](#), [75](#)
 Gazdar, G. [44](#)
 general knowledge *see* world knowledge
 General Theory of Verbal Humour
 (GTVH) [3](#), [5](#);
 defining concepts [72](#)–6;
 empirical aspects [79](#)–[80](#);
 HCPP viewed as [156](#);
 joke identity [84](#);
 joke similarity [79](#), [85](#)–8;
 knowledge resource [70](#)–1, [76](#), [77](#), [79](#), [85](#);
 Language [71](#)–2, [76](#), [77](#), [78](#), [86](#)–[88](#);
 Logical Mechanism [56](#), [71](#)–2, [74](#)–6, [84](#), [86](#), [88](#);
 Narrative Strategy [71](#)–2, [76](#), [78](#), [86](#)–7;
 outline [70](#)–2;
 parameters [71](#), [76](#)–7, [84](#), [86](#)–8;
 relation to theory of language [78](#)–9;
 scripts [72](#)–4, [78](#)–9;
 Script Opposition [71](#)–4, [76](#), [84](#), [86](#)–8;
 Situation [71](#)–2, [76](#), [78](#), [86](#)–8;
 summary of critique [80](#);
 Target [71](#)–2, [76](#), [78](#), [86](#)–7
 generative capacity/power [38](#)
 Giora, R. [63](#), [70](#)
 Godkewitsch, M. [11](#), [59](#)
 Goldstein, I. P. [51](#)
 Goldstein, J. H. [4](#)
 Gosling, J. [51](#)
 grammar testers [20](#)–1

grammaticality:

assumption [9](#), [37](#)–8, [94](#), [124](#);
 formal definition [190](#), [191](#);
 judgement [20](#), [187](#)
 Green, T. A. [40](#), [88](#), [130](#), [166](#)
 Greimas, A. J. [60](#)
 Gricean maxims [88](#)
 groaning [109](#)
 Gross, D. [149](#), [154](#)
 Gruner, C. [7](#), [8](#), [48](#), [49](#)
 GTVH *see* General Theory of Verbal Humour
The Guardian [119](#)

Haddad, J. [7](#)
 HAHAAcronym [154](#)–5, [157](#)–8
 Hartley, D. [46](#)
 Hayes, P. J. [12](#), [52](#)
 Hazlitt, W. [47](#)
 HCPP [149](#)–51, [157](#)
 Hehl, F.-J. [57](#)
 Heller, L. G. [59](#), [112](#)–13, [138](#)–9
 Hempelmann, C. [69](#), [70](#), [72](#), [74](#), [75](#)
 Hetzron, R. [4](#), [91](#)
 history of humour research [4](#)
 Hobbes, T. [180](#)
 Hobbs, J. [53](#), [95](#)
 Hobbs, S. [111](#), [112](#), [160](#)–1
 Hockett, C. F. [35](#)–6
 Hofstadter, D. [84](#)–5, [179](#)
 Holmes, D. [51](#)
 homonyms [152](#)
 homophony [32](#), [115](#), [122](#), [131](#), [139](#);
 see also homonyms
 hostility theory of humour *see* aggression
 theory of humour
 Hudson, R. [110](#)
 human behaviour [1](#), [6](#), [11](#), [188](#)
 humour theory [1](#)–2, [6](#), [11](#), [24](#), [64](#), [156](#);
 different types [7](#);
 linguistic analysis and [18](#)–20
 Hutcheson, F. [46](#)

identity of jokes *see* jokes, identity

idiom [152](#)–3

illocutionary force [91](#), [160](#), [164](#)

importance of humour [1](#)

INAPPROPRIATENESS [61](#)–4, [99](#)–100, [107](#), [141](#), [165](#), [176](#)–7, [178](#), [179](#), [182](#), [183](#)–4, [187](#)

INC-RES component [57](#)–8

incongruity [4](#), [46](#)–58, [62](#)–4, [92](#), [97](#), [168](#), [170](#), [172](#), [191](#);

dynamic [49](#)–50;

inherent [50](#);

licensed [49](#)

(*see also* suspension of disbelief);

presentational [50](#);

in secure setting [49](#);

static [49](#)–50;

theory of humour [7](#), [18](#), [70](#), [156](#)

incongruity-resolution [54](#)–8, [59](#)–68, [70](#), [75](#), [107](#)

inconsistency, logical *see* consistency, logical

inference [45](#), [52](#)–3, [57](#), [93](#)–4, [97](#), [188](#), [191](#), [193](#);

default [52](#)–3, [93](#);

defeasible [52](#);

non-monotonic [52](#);

plausible [92](#)
 information update [191](#), [194](#)–5
 informativeness [63](#)
 ingenuity [182](#);
 see also intricacy;
 cleverness
 initiator jokes [111](#), [160](#)–1
 innuendo [36](#), [178](#)
 insults [180](#), [182](#);
 see also butts of jokes
 intelligence:
 artificial *see* artificial intelligence;
 human [6](#)
 intelligent agents [21](#), [22](#)
 interactive jokes [159](#)–61, [163](#), [164](#), [174](#), [199](#)
 International Phonetic Alphabet [5](#), [30](#)
 interpretation [61](#)–8, [70](#), [72](#), [90](#)–[107](#), [137](#), [176](#)–7, [183](#), [196](#);
 most obvious [196](#), [197](#), [198](#)
 interpretation sequences [195](#)
 intertextuality [181](#)
 intonation [31](#), [154](#)
 intricacy [136](#);
 see also ingenuity;
 cleverness
 intuitive judgement [20](#), [187](#)
 IR *see* incongruity-resolution
 Isahara, H. [22](#), [142](#)
 isotopy disjunction model [60](#)
It's a Mad, Mad, Mad, Mad World [139](#)
 Ito, A. [22](#), [142](#)

JAPE [83](#)–4, [123](#), [128](#), [131](#), [147](#)–9, [150](#), [157](#)
 Java [51](#)
 Jespersen, O. [110](#)–11
 Johnson, R. [51](#)
 Johnson, S. [135](#)

jokes:

- abstract concept [15–16](#)
(*see also* jokehood);
- appreciation [16](#), [18](#), [22](#);
- books [5](#), [15](#), [160](#);
- class [19–20](#), [24–6](#), [29](#), [33](#), [67](#), [79](#), [83](#), [107](#), [110](#), [137](#), [175–6](#), [186–8](#);
- components [175–85](#), [186](#), [188](#);
- comprehension [16](#), [18](#);
- generation *see* computer generation of jokes;
- identity [5](#), [81–9](#);
- similarity [5](#), [79](#), [81–9](#);
- tellers [159–61](#), [182](#);
- telling [30](#), [104](#), [109](#), [159](#), [178](#);
- underlying core [83–5](#)

jokehood [16](#), [18](#), [20](#), [39](#), [100](#), [149](#), [153](#), [164](#), [182](#), [188](#)

Jonson, B. [135](#)

Joyce, J. [113](#)

juxtaposition [75](#)

Kagan, J. [65](#)

Kant, I. [46](#), [47](#), [48](#), [49](#), [50](#), [66](#)

Kaplan, R. M. [51](#)

Katz, B. [47](#), [64](#)

Kay, M. [51](#)

Keith-Speigel, P. [7](#), [46](#), [47](#), [48](#), [49](#)

Kelly, L. G. [99](#)

Kempson, R. [17](#)

Kernal, H. K. [23](#)

Kidd, F. C. [1](#)

Kimball, J. P. [38](#)

Kitagaki, I. [22–3](#)

Klein, E. [44](#)

knock-knock jokes [126–7](#), [159](#), [199](#)

knowledge representation [51](#), [141](#), [164](#)

knowledge resource (in GTVH) [70–1](#), [76](#), [78–9](#)

Koestler, A. [47](#), [48](#), [51](#), [53](#), [54](#), [66](#)

LA *see* General Theory of Verbal Humour, Language

Langevin, R. [11](#)

Language (in GTVH) *see* General Theory of Verbal Humour

language:

- computer production [141–2](#), [144](#), [148](#), [152](#), [157](#), [188](#)
(*see also* computer generation of jokes);
- generation [132](#);
- learning [151](#);
- understanding [21](#), [53](#), [45](#), [64](#), [66–7](#), [101](#), [188](#), [194–6](#)
(*see also* left-right processing)

Larson, G. [75](#)

Lascarides, A. [93](#)

Latta, R. L. [1](#), [2](#), [18](#), [49](#), [100](#)

Laurel, B. [164](#)

left-to-right processing [17](#), [45](#)

Lessard, G. [143–7](#)

Levinson, S. [39](#), [91](#), [160](#)

Levison, M. [143–7](#)

lexemes [148](#), [192](#);
see also lexical entry

lexical analysis [118](#), [125](#), [192](#), [199](#)

lexical entries [73](#), [147–8](#), [150](#), [155](#), [192–3](#)

lexically definable predicates [192](#)

lexicons [45](#), [145](#), [150](#), [152](#), [192–3](#), [199](#);
see also lexemes

LIBJOG [143–4](#)

light bulb jokes [71](#)–[2](#), [76](#), [143](#), [159](#)
linguistic theory:
 assumed by GTVH [78](#);
 relationship to humour theory [29](#), [78](#)–[9](#), [189](#);
 universal [9](#)–[10](#), [24](#)
linguistic:
 joke [14](#)–[15](#), [31](#), [176](#), [182](#)–[4](#);
 knowledge in joke comprehension [30](#)–[3](#), [38](#)–[9](#), [63](#), [139](#), [193](#);
 levels [34](#), [40](#)–[2](#)
linguistic appropriateness *see* contextual appropriateness
linguistics:
 analogy with [9](#), [175](#);
 Chomskyan [9](#);
 contributing discipline [25](#), [73](#), [141](#);
 formal [4](#);
 generative [2](#), [9](#)–[10](#), [11](#)–[12](#), [20](#), [187](#)
linkage to context [135](#)–[6](#), [191](#), [192](#)
Lippman, L. G. [136](#)
literal meanings [92](#), [94](#), [97](#), [190](#), [191](#), [194](#)–[5](#)
LM *see* General Theory of Verbal Humour, Logical Mechanism
Loehr, D. [23](#)
logic [52](#), [190](#), [194](#);
 humorous [7](#), [67](#), [68](#), [84](#)
logical implication [45](#), [93](#), [97](#), [194](#)
Logical Mechanism (in GTVH) *see* General Theory of Verbal Humour
Lyons, J. [110](#), [191](#)

McDermott, D. [12](#), [93](#)
McGhee, P. E. [4](#), [49](#), [57](#)
McKay, J. [23](#), [143](#), [151](#), [153](#), [158](#)
Maesen, W. A. [7](#)
Maier, N. R. F. [49](#), [66](#)

- malapropisms [116](#)–17, [124](#), [125](#), [135](#), [141](#)
 Maranda, E. K. [167](#)
 markedness [63](#)
 Martin, P. [95](#)
 Matthews, T. S. [31](#)
 Mele, F. [143](#)
 mental spaces [93](#)
 Mervis, C. [10](#)
 meta-level explanation of context [120](#)
 metaphor [53](#)–4, [168](#)
 metathesis [122](#), [127](#)
 Metcalf, F. [13](#)
 method (within an analysis) [42](#)–4
 methodology [2](#)–4, [6](#)–27, [28](#)–9, [69](#), [80](#), [153](#), [156](#)–8, [186](#), [187](#)
 Meyer-Viol, W. [17](#)
 Miller, G. A. [149](#), [154](#)
 Miller, K. [149](#), [154](#)
 mimetic phrases [118](#)
 Minsky, M. [7](#), [51](#), [52](#), [59](#)
 misdirection, deliberate [104](#)–5
 misinterpretation [14](#), [113](#), [91](#), [98](#), [106](#), [113](#), [159](#)–60, [161](#);
 see also misunderstanding
 misunderstanding [53](#), [67](#), [79](#), [99](#), [107](#)
 model adopted [17](#), [18](#), [29](#), [44](#)–5, [105](#), [189](#), [196](#)
Monty Python's Flying Circus [54](#)
 Morkes, J. [23](#)
 morphemes [34](#)–5, [40](#)
 morphology [193](#)
 Morreall, J. [4](#), [7](#), [34](#), [46](#), [47](#), [48](#), [49](#), [50](#), [79](#), [93](#)
 MRC psycholinguistic database [149](#)
 Muir, F. [120](#), [136](#)
 Murray, C. [128](#)
- narrative:
 repetitive *see* repetition;
 simple [104](#), [106](#), [178](#), [198](#)
 narrative jokes [92](#), [98](#), [109](#), [121](#), [161](#), [164](#), [170](#)
 Narrative Strategy (in GTVH) *see* General Theory of Verbal Humour
 Nash, W. [1](#), [140](#)
 Nass, C. [23](#)
 necessary and sufficient conditions [10](#), [21](#), [26](#), [53](#)–4, [76](#), [129](#), [178](#)
 Nerhardt, G. [48](#)–9, [50](#)
 neurological effects [2](#)
 newspaper headlines [118](#)–19
 Nijholt, A. [23](#)
 Nobbs, D. [127](#)
 NON component [57](#)–8
 non-humorous text [12](#), [35](#)–6, [38](#), [57](#), [64](#), [87](#), [164](#), [179](#), [181](#);
 GTVH position on [76](#), [78](#), [80](#)
 non-jokes [16](#), [21](#), [53](#), [82](#), [88](#), [93](#), [149](#);
 see also non-humorous text
 non-linguistic factors *see* social/cultural setting of humour;
 cultural factors;
 visual presentation
 non-monotonic reasoning [52](#), [196](#)
 Norden, D. [120](#), [136](#)
 Norman, D. A. [51](#)
 Norrick, N. R. [51](#), [52](#), [60](#), [120](#), [181](#)
 Norvig, P. [12](#)
 NS *see* General Theory of Verbal Humour, Narrative Strategy
- O'Mara, D. [23](#), [82](#)
 Oaks, D. D. [42](#), [43](#)

OBVIOUSNESS [61](#)–4, [92](#), [94](#)–5, [97](#), [99](#)–100, [104](#)–5, [176](#)–7, [190](#), [193](#), [194](#)–5
 one-liner [99](#)
 orthographic puns [134](#), [138](#)–9, [186](#)
 orthography [14](#), [30](#)–3, [34](#)–5, [40](#)–1, [113](#), [140](#), [189](#)

Pain, H. [143](#), [149](#), [157](#)–8

Palmer, J. [93](#), [164](#)

Paolillo, J. C. [75](#)

paradigmatic puns *see* puns,paradigmatic

parallelism *see* intertextuality

parameters:

 of computer program [156](#), [157](#);

 of formal theory [83](#);

 in GTVH 71–2, [76](#)–8, [80](#), [84](#)–8

paraphony [116](#);

see also similarity, phonetic

parity errors [38](#)–9

Parker, D. [101](#)

parody [181](#)

paronymy [112](#);

see also similarity, phonetic

Parsons, D. [100](#)

parts of speech [45](#), [190](#)

Paulos, J. A. [47](#), [48](#), [59](#)
 Pearce, M. T. [158](#)
 Pease, A. [158](#)
 Pepicello, W. J. [30](#), [40](#), [55](#), [88](#), [130](#), [166](#), [167](#)
 perlocutionary effects [91](#), [160](#)
 Petray, M. J. [17](#), [35](#)
 philosophy of language [4](#)
 phonetic:
 form [5](#), [30](#)–3, [44](#)–5, [110](#), [126](#), [127](#);
 string [189](#)–90
 phonetics [24](#), [33](#), [34](#), [40](#), [111](#)
 phonology *see* phonetics
 physiological effects [2](#)
 Pien, D. [55](#), [63](#)
 Plato [4](#)
 Pocheptzov [98](#)
 poetic devices [53](#)
 poetic story jokes [120](#)
 practical jokes [160](#), [174](#), [181](#)
 pragmatic well-formedness [37](#), [124](#), [128](#), [131](#)
 pragmatics [5](#), [159](#)–74
 pranks [181](#);
 see also practical jokes
 prediction:
 meta-level [104](#);
 during text processing [101](#)–5, [194](#), [196](#)
 (*see also* expectation);
 of a theory [1](#), [2](#), [8](#), [43](#), [69](#), [79](#)–80, [175](#), [184](#), [187](#)
 prerequisites [39](#), [191](#), [194](#)–5
 presupposition [39](#), [45](#), [60](#), [93](#), [94](#), [95](#), [97](#), [191](#)
 priming [101](#), [104](#)–5
 primitive constructs [24](#)–5, [187](#), [189](#)
 procedural accounts:
 characterization of model [68](#), [156](#);
 vs. declarative [17](#)–18, [21](#), [112](#)
 processing:
 language [28](#)–9, [64](#), [156](#);
 psychological [18](#)–20
 pronunciation [31](#), [32](#), [34](#), [45](#), [134](#), [190](#)
 proper names [111](#)
 propositions, logical units of meaning [39](#), [45](#), [53](#), [67](#), [92](#)–4, [119](#), [179](#), [183](#), [190](#)–1, [194](#), [196](#)–8
 propositional jokes [14](#)–15, [31](#), [107](#), [141](#), [176](#), [182](#)–4, [186](#)
 prosody *see* intonation
 prototypical definition [10](#)
 psychoanalytical theory of humour [7](#)
 psychological research [4](#), [11](#), [57](#), [64](#), [158](#), [187](#)
 Pullum, G. [44](#)
 puns [5](#), [13](#), [14](#), [30](#), [33](#), [59](#), [99](#), [108](#)–42, [161](#), [168](#), [175](#);
 bilingual [140](#);
 computer generated [144](#)–53, [193](#);
 in Japanese [142](#);
 linguistically abnormal paradigmatic [130](#)–4, [186](#);
 linguistically normal paradigmatic [125](#)–6, [127](#), [131](#), [135](#), [137](#), [142](#), [156](#), [182](#), [184](#), [185](#), [199](#);
 multi-way [140](#);
 in other media [138](#)–9;
 paradigmatic [101](#), [110](#), [114](#)–27, [130](#)–7, [125](#)–6, [127](#), [139](#);
 syntagmatic [110](#), [127](#)–30, [133](#), [135](#), [138](#), [141](#), [142](#)
 punchline conflict [107](#), [196](#)–7, [198](#)
 punchline inference [98](#), [176](#), [197](#)–8
 punchline resolution [197](#), [198](#)
 punchline revision [97](#), [176](#), [197](#), [198](#)
 punchline:
 answering puzzle [163](#)–4, [179](#);
 in delivery mechanism [183](#), [184](#), [196](#)–8;

empty [36](#);
 Hockett definition [35–6](#);
 pre- and post- [176–7](#), [179](#), [183](#), [184](#);
 in propositional jokes [59–68](#), [91–2](#), [94–9](#), [105–7](#);
 in puns [120–1](#), [123–4](#), [131–2](#), [133](#), [137](#), [149](#);
 in repetitive joke [170](#), [171](#), [173](#), [181](#);
 in responder joke [161](#);
 in riddle [166](#);
 thwarting expectations [180](#), [181](#)
 punctuation [139](#)

question–answer form [123](#), [159–69](#), [184](#)
 questions raised in set-up [163–4](#), [179](#), [181](#)

Raju, A. [171](#)
 Raskin, V. [1](#), [3](#), [4](#), [7](#), [13](#), [27](#), [28](#), [36](#), [39](#), [40](#), [46](#), [47](#), [49](#), [51](#), [53](#), [54](#), [60](#), [63](#), [69–73](#), [75–9](#), [84](#), [85–8](#), [98](#), [109](#), [110](#), [112](#), [134](#),
[143–4](#), [166](#), [167](#), [180](#)
 readability [149](#)
 reasoning *see* inference
 Redfern, W. [27](#), [110](#)
 referential jokes [14–15](#), [30](#)
 Reiter, R. [52](#), [93](#)
 relief/release theory of humour [7](#)
 repetition [170–4](#), [182](#), [181](#)
 research subproblems [63–4](#), [141](#), [187](#)
 resolution of ambiguity [61](#), [112](#)
 resolution of incongruity *see* incongruity-resolution

resolution preceding incongruity [56](#)
 responder jokes [111](#), [160](#)–1
 riddles [13](#), [22](#), [31](#), [55](#), [78](#), [105](#), [159](#)–60, [165](#)–9, [186](#);
 as NS in GTVH [71](#), [76](#), [77](#);
 punning [111](#), [122](#)–4, [125](#), [130](#), [133](#)–4, [146](#)–9;
 as puzzles [166](#)–9, [174](#)
 Ritchie, G. [3](#), [24](#), [61](#), [94](#), [120](#)–1, [143](#), [148](#), [149](#), [150](#), [157](#)–8
The Rivals [116](#)
 Roberts, R. B. [51](#)
 Rosch, E. [10](#)
 Ross, A. [5](#), [107](#), [126](#)
 Rothbart, M. K. [46](#), [49](#), [55](#), [63](#)
 Ruch, W. [11](#), [54](#), [57](#), [71](#), [76](#), [85](#)–7, [158](#)
 Russell, B. [39](#)
 Russell, S. [12](#)
 Rutter, J. F. [69](#)

 Sag, I. [44](#)
 salience in context [114](#), [119](#), [191](#), [197](#)
 Sandt, R. van der [39](#)
 scales of measurement [149](#), [150](#), [153](#), [155](#), [157](#)–8
 Schank, R. [51](#), [52](#), [93](#)
 schemata [83](#)–4, [148](#), [150](#)
 Schopenhauer, A. [46](#), [47](#), [48](#), [49](#), [66](#)
 scripts [51](#)–4, [70](#)–5, [79](#)–80, [93](#);
 see also frames
 Script Opposition (in GTVH) *see* General Theory of Verbal Humour
 selectional restrictions [95](#)
 semantic concept *see* concepts, units of knowledge representation
 semantic fields [154](#)–5, [191](#)
 semantic representation [45](#), [101](#), [190](#)–1
 Semantic Script Theory (of language) [78](#)
 Semantic Script-based Theory of Humour (SSTH) [3](#), [25](#), [51](#), [54](#), [60](#), [63](#), [69](#)–70, [72](#)–5, [78](#), [79](#), [80](#), [99](#), [112](#)
 semantic well-formedness [94](#), [95](#), [124](#), [128](#), [131](#)
 sense of humour [11](#), [21](#), [57](#), [158](#)
 set-ups:
 containing puzzles [163](#)–4, [179](#);
 in propositional jokes [59](#)–68, [91](#)–2, [94](#)–9, [104](#)–7;
 in puns [120](#), [123](#), [137](#), [149](#);
 repetitive [171](#), [173](#);
 in riddles [165](#)–6;
 with supplements [181](#);
 supplied by audience [161](#)
 sets of beliefs [180](#), [193](#)
 SEX component [57](#)–8
 sexual factors [61](#), [70](#), [93](#), [154](#), [179](#)
 Shapiro, S. C. [1](#)
 Sheridan, R. [116](#)
 Shultz, T. R. [30](#), [46](#), [48](#), [49](#), [50](#), [54](#)–5, [59](#), [63](#), [66](#)
 SI *see* General Theory of Verbal Humour, Situation
 similarity of stimulus [48](#)–9
 similarity, jokes *see* jokes, similarity
 similarity, orthographic [113](#), [145](#)
 similarity, phonetic:
 assumed constructs [44](#);
 as factor in puns [110](#), [115](#)–16, [118](#)–19, [122](#), [124](#), [125](#), [128](#), [130](#), [137](#), [156](#), [183](#), [184](#);
 formal constructs [189](#);
 across languages [140](#);
 as a type of ambiguity [40](#), [111](#)–12
 similarity, textual [44](#), [82](#), [87](#), [89](#), [122](#)
 simile [53](#)–4
 Situation (in GTVH) *see* General Theory of Verbal Humour
 situation, representation of context [39](#), [45](#), [183](#), [191](#), [193](#), [199](#)

slot-filler structures [72](#), [73](#)
SO *see* General Theory of Verbal Humour, Script Opposition
Sobkowiak, W. [110](#), [111](#), [116](#)
social theory of humour [7](#)
social/cultural setting of humour [1](#), [2](#), [38](#), [193](#)
sociological research [64](#)
speech acts [160](#)
speech vs. writing [29–34](#), [37](#), [134](#), [138](#)
spelling [31](#), [32](#), [45](#), [134](#), [190](#)
spontaneous humour [109](#), [114](#)
spoonerisms [128](#)
SSTH *see* Semantic Script-based Theory of Humour
statistical analysis [57](#), [158](#)
statistical significance [149](#), [153](#)
Steedman, M. [101](#), [170](#)
Steele, G. L. [51](#)
Stevenson, R. L. [38](#)
Stickel, M. [95](#)
stimulus-side theory of humour [18–19](#)
Stock, O. [22](#), [143](#), [154](#)
story puns [120–2](#)
Strapparava, C. [143](#), [154](#)

structural description of jokes [175](#)–6;
 notation for [24](#)
 sufficient conditions for humour [182](#)
 Suls, J. M. [17](#)–[18](#), [46](#), [54](#), [56](#), [59](#), [65](#)–8
The Sun [119](#)
 superiority [180](#)–2
 superiority theory of humour [7](#)–[8](#), [94](#), [180](#)
 supplementary factors [101](#), [102](#), [174](#), [179](#)–82, [183](#)–4
 surprise [13](#), [17](#), [55](#)–6, [57](#), [61](#), [66](#)–7, [70](#), [180](#);
 theory of humour [49](#);
 see also expectation
 suspension of disbelief [49](#), [64](#)
 Sussman, G. [29](#)
 Svebak, S. [11](#)
 syllables [34](#)–5
 syntax [40](#), [190](#), [192](#)
 syntax trees [44](#)

TA *see* General Theory of Verbal Humour, Target
 TABOO [61](#), [92](#), [93](#), [101](#), [177](#), [178](#)
 Takizawa, O. [22](#), [142](#)
 Target (in GTVH) *see* General Theory of Verbal Humour
 taxonomies:
 of humour [50](#), [57](#);
 of jokes [26](#), [44](#), [110](#), [137](#), [140](#);
 of theories [7](#)
 templates [150](#)
 Tengi, R. [149](#), [154](#)
 text generation *see* language, computer production;
 language generation
 text processing *see* language generation;
 language understanding;
 processing, language
 text strings [33](#), [35](#)–7, [44](#)–5, [189](#), [195](#), [199](#)
 text-interpreter [193](#)–5
 text-meaning mapping [45](#), [190](#), [191](#)
 theory-external definitions [24](#)–5
 theory-internal definitions [24](#)–5
 theory:
 avoidance of [27](#);
 of grammar [9](#);
 of humour *see* humour theory;
 of jokes [9](#), [17](#)–[18](#), [37](#), [168](#), [175](#), [184](#);
 of language [9](#)–[10](#), [12](#), [19](#), [28](#), [54](#), [78](#);
 premature [2](#)
 Thompson, H. [51](#)
 timing [163](#)
 Todman, J. [23](#)
 Tom Swifty [136](#), [144](#)–5
 top-down theorising [3](#)
Treasure Island [38](#)
 Turner, R. [52](#)
 two-stage model [59](#)–[60](#), [65](#)–8, [101](#)
 type/token distinction [30](#)
 typography [32](#), [139](#)

universal theory of humour [2](#), [8](#), [16](#), [26](#)
 ur-jokes [84](#)–5
 use–mention distinction [139](#)
 user-interfaces [22](#)–3

Veatch, T. [1](#)
 VEH *see* verbally expressed humour

Venour, C. [140](#), [143](#), [149](#)–51, [157](#)
 verbal jokes [14](#)–[15](#), [106](#)
 verbalized humour [13](#)
 verbally expressed humour:
 ambiguity in [49](#)–[50](#);
 addressed by others [4](#);
 basic idea [2](#), [13](#)–[14](#);
 coverage here [5](#), [186](#)–[7](#);
 incongruity in [47](#), [49](#);
 IR models of [59](#), [60](#);
 methodology for exploring [8](#), [12](#), [24](#), [25](#)–[6](#), [33](#), [35](#), [36](#), [43](#), [44](#), [78](#);
 odd poems [31](#);
 topic of GTVH [3](#), [69](#)
 VINCI [144](#)–[6](#), [150](#), [157](#)
 visual presentation [13](#), [139](#)
 visual puns [139](#)

Waller, A. [23](#), [82](#)
 Watson, C. [31](#)
 Webb, K. [125](#)
 Weiner, E. J. [63](#), [167](#)
 Weisberg, R. W. [30](#), [40](#), [55](#), [167](#)
 West, M. [127](#)
 Wiggins, G. A. [139](#), [158](#)
 Wilde, O. [13](#), [50](#), [129](#)
 Wilson, C. P. [4](#), [7](#), [46](#), [49](#), [51](#), [163](#)
 Wilson, R. C. [1](#)
 Winograd, T. [40](#), [51](#)
 Winterstein, D. [158](#)
 WisCraic [151](#)–[3](#), [157](#)
 with holding information [104](#)–[5](#)
 Wittgenstein, L. [10](#)
 Wode house, P. G. [34](#)

WordNet [149](#), [154](#)
working assumptions [10](#), [14](#), [16](#), [19](#), [35](#), [37](#)
world knowledge [32](#), [38–9](#), [53](#), [63](#), [155](#), [188](#)
World Wide Web, jokes on [5](#), [15](#)
Wright, E. L. [111](#), [140](#)
written forms [30](#), [44–5](#), [110](#), [134](#)
written strings [189–90](#)

Yamaguchi, H. [105](#)
Yanagida, M. [22](#), [142](#)