Behavioural, Electrophysiological and Connectionist Studies in Inflectional Morphology

A thesis submitted for the degree of Doctor of Philosophy

by Giles Shilson, Balliol College, University of Oxford

Trinity Term MM



ACKNOWLEDGMENTS

Professor Kim Plunkett and all in the Cognitive Science Unit, in particular Dr Patrick Juola, Dr Ramin Nakisa and Asma Siddiki; Dr Anna Nobre and all in the Brain and Cognition Laboratory, especially Dr Anling Rao, Dr Ed Wilding and Rebecca Saxe; all at Balliol College for reasons too numerous and obvious to mention; also Dr George Boys-Stones, Dr Thomas Dixon, Stephen Fry, Dr Benny Lee, Professor Anthony Levi, Clarissa Levi, Dr Virginia Marchman, Professor William Marslen-Wilson, Dr Stephen Tucker, Dr Jennifer Utman, the Medical Research Council for funding this Doctorate, and finally my family for all their support.

RETIACVLIS SCIENTIA FICTIS
CAPITVLA QVAE SEQVVNTVR
PERFECI PRO GRADV
DOCTORIS IN PHILOSOPHIA
SHILSON BALLIOLENSIS

Behavioural, Electrophysiological and Connectionist Studies in Inflectional Morphology

A thesis submitted for the degree of Doctor of Philosophy
by Giles Shilson, Balliol College, University of Oxford
Trinity Term MM

Abstract

Theories of generative linguistics hold that language processing occurs by means of the manipulation of symbols by explicit rules. The past 15 years have seen the development of a radical challenge to this view, derived from a form of computational modelling called Parallel Distributed Processing or connectionism. Connectionist linguistic theory holds that language processing takes place at a subsymbolic level, and that the appearance of rule-driven behaviour is formed by the abstraction of patterns from the environment.

The English past tense has become a critical arena of dissent between symbolic and subsymbolic theories of linguistics. There are two current dominant theories of past tense inflection: hybrid dual-route theory, championed by Pinker, posits a symbolic, explicit, default rule for the regularisation process, and an associative memory component for irregular exceptions; single-route theory maintains that regular and irregular inflectional morphology may both be accounted for within a single, subsymbolic, associative system which contains no explicit linguistic rules.

This doctoral thesis describes a new classification of phonological similarity between verbs ('neighbourhood density'), which is used to develop mutually-exclusive and empirically-testable hypotheses from the two dominant theoretical perspectives of English past tense inflectional morphology. Empirical research is conducted in the domains of experimental psychology, electrophysiology and connectionist modelling: four novel verb elicitation tasks are performed on adults; two ERP studies investigate brain activity for regular, irregular and novel verb inflection; and five neural network simulations are built in order to compare human data with network performance. Data are reported which have implications for single- and dual-route theories of past tense processing.

Behavioural, Electrophysiological and Connectionist Studies in Inflectional Morphology

A thesis submitted for the degree of Doctor of Philosophy

by Giles Shilson, Balliol College, University of Oxford

Trinity Term MM

Full Abstract

This thesis explores key issues in the debate between single- and dual-route theories of inflectional morphology, from the perspective of three different fields of cognitive science: experimental psychology, neurophysiology and artificial neural network modelling. Within these fields, factors of regular and irregular neighbourhood density are controlled using novel and real verb stimuli, and found to have implications for the two dominant theories of past tense processing.

Chapter 1: General Introduction

Theories of generative linguistics hold that language processing occurs by means of the manipulation of symbols by explicit rules. The past 15 years have seen the development of a radical challenge to this view, derived from a form of computational modelling called Parallel Distributed Processing or connectionism. Connectionist

linguistic theory holds that language processing takes place at a subsymbolic level, and that the appearance of rule-driven behaviour is formed by the abstraction of patterns from the environment.

The English past tense has become a critical arena of dissent between symbolic and subsymbolic theories of linguistics. There are two current dominant theories of past tense inflection: hybrid dual-route theory, championed by Pinker, posits a symbolic, explicit, default rule for the regularisation process, and an associative memory component for irregular exceptions; single-route theory maintains that regular and irregular inflectional morphology may both be accounted for within a single, subsymbolic, associative system which contains no explicit linguistic rules.

Chapter 2: What is Connectionism?

The basic principles of connectionist structure are described and it is shown how assumptions about the behaviour of neurons are reflected in the principles underlying neural network modelling. The basic theory and terminology behind connectionism are explained, including node function, network structure and learning principles, with which psychologists have been able to develop new, testable theories in many areas of cognition.

Chapter 3: The Single-Route/Dual-Route Debate

Key papers in the single-route/dual-route debate are rehearsed, beginning with the first connectionist paper on the English past tense. A background to the field

up to the mid-1990s is provided: during that period, symbolic theories have been refined, and connectionist models have become more powerful, more accurate and better implemented.

Chapter 4: Neighbourhoods

One critical issue in the single-route/dual-route debate is that each theory makes different predictions about the impact that the number of regular verbs in the rhyming neighbourhood of a novel stimulus may have on its inflection. Steven Pinker's hybrid dual-route theory proposes two processing routes, one for regulars and one for irregulars. According to this theory, verbs are processed by an explicit, symbolic, default suffixation rule unless blocked by a match (or sometimes a partial match) with a known irregular in a separate associative mechanism. The model predicts, therefore, that irregular neighbourhood density will affect the inflection of a novel verb, but that the number of similar regular verbs will not have an effect. Conversely, connectionist single-route theories predict that, since verbs are processed by analogy to all known items, both irregular *and* regular neighbourhood density will have an impact on novel verb inflection.

In this chapter it is shown how the findings of certain empirical studies of inflectional morphology are undermined by unsound methodology, specifically the poor classification of stimuli according to their rhyming neighbourhood. The chapter focuses on two papers: Prasada & Pinker (1993) and a riposte by Lee (1995). The systems of neighbourhood classification used by Prasada & Pinker and Lee are dis-

cussed, analysed and found to be flawed. A new way of classifying neighbourhood density is proposed, and it is argued that this quantitative classification, in which a verb is assigned a measure of irregular and regular density in comparison with *all* verbs in the language, provides a more objective and methodologically sound basis for stimulus selection and control. This new classification forms the basis for the three major experimental components of the thesis in subsequent chapters.

Chapter 5: Behavioural Studies

A set of four past tense elicitation experiments is described, involving the application of the new neighbourhood classification to a set of novel verb stimuli, and the results are found have implications for the single-route/dual-route debate in inflectional morphology. Adult subjects were required to inflect stem forms of novel verbs to appropriate past tense forms. Stimuli were controlled for irregular and regular neighbourhood density, and it was found that novel stimuli were regularised or irregularised by adults at different rates, depending on their neighbourhood structure. Specifically, effects of irregular and regular neighbourhood density were found for both pseudoregular and pseudoirregular novel verbs: in both cases, novel verbs with an above-median number of irregulars in their rhyming neighbourhoods were more likely to be inflected by an irregular paradigm than stimuli sharing their rime with few irregulars (or none at all). Furthermore, and critically, novel verbs with an above-median number of *regulars* in their neighbourhoods were more likely to be regularised than those with few regular neighbours. It is shown that this finding conflicts with

the dual-route theory of inflectional morphology in which, since regular inflections are generated 'on-the-fly' (rather than being stored—like irregulars—in the lexicon), regular neighbourhood density should have no effect on irregularisation rate.

This set of experiments therefore argues strongly against the hybrid dual-route theory and supports alternative connectionist accounts, which predict that irregular and regular neighbourhood density should *both* affect the rate of regularisation of novel verbs. Moreover, the direct conflict between the results of these experiments and the findings reported by Prasada & Pinker suggests that their data may—as suggested in chapter 4—have resulted from methodological flaws in the design of their stimulus categories.

Chapter 6: Electrophysiological Studies

Two ERP studies are described. The first used real verbs as stimuli to be inflected to the past tense. Statistical analysis found substantial differences in activation between responses to regular and irregular stem forms, thereby replicating Breyer (1997), although in this case more differences were found. This may be due to the use of a larger number of stimuli, with neighbourhood density carefully controlled and balanced, and an improved presentation procedure. Subset analyses found that differences in regular and irregular neighbourhood density each had noticeable effects on the inflection of both regular and irregular verbs.

The second ERP experiment recorded EEG from subjects while they were inflecting *novel* stimuli to appropriate past tense forms. Small frontal and lateral effects

were found when the waveforms for regularised and irregularised stimuli were compared statistically, but otherwise the elicited waveforms were remarkably similar. This suggests that a similar process may be taking place regardless of the inflectional paradigm used. However, when subset analyses were conducted on the regularised novel stimuli, it was found that the number of irregulars in the rime neighbourhood of a regularised stimulus had a marked effect on its waveforms but, conversely, that the number of *regulars* in its neighbourhood had almost no impact at all. These findings would seem to be more easily accounted for by dual-route than single-route theories.

Waveforms for regularised novel verbs were then compared with those for real regular and irregular verbs. Substantial differences were found after the main language-related potential peak, showing that novel stimuli tended to evoke less positive potentials than regular or irregular real verbs. Particularly significant effects were found in posterior areas.

In sum, the two ERP studies yielded a considerable amount of new data which provide a detailed analysis of the effects of regular and irregular neighbourhood structure on the potentials elicited during real and novel past tense verb processing.

Chapter 7: Connectionist Studies

In this chapter, it is shown how simple feedforward connectionist networks were designed and constructed to map phonological representations of monosyllabic English verb stem forms onto phonological representations of their past tense inflected

forms. These networks successfully learnt to map verb stems onto their past tense forms, and successfully modelled the 'regularity by frequency' interaction found in human behavioural data. A regression analysis was performed on the system's internal representations to find out whether internal function for regular and irregular stimuli could be dissociated: significant partial correlations were found between the networks' internal state and the type of verb being processed. The networks were tested on 48 novel verb stimuli (used in the behavioural experiments from chapter 5), and the influence of neighbourhood density on network performance was compared with the human data. The networks' reactions to novel verb stimuli broadly paralleled earlier behavioural findings for humans, although significant differences were found. Despite these differences, the networks accounted for much of the human data, and were particularly competent at modelling findings which have been used in the literature to argue *against* subsymbolic single-system theories of English past tense inflectional morphology.

Chapter 8: General Conclusion

The principal arguments and findings of the thesis are summarised, and suggestions for future work made.

Collectively, these three sets of studies—behavioural, electrophysiological and connectionist—have produced a considerable quantity of data, which provide detailed findings with regard to regular and irregular neighbourhood density. In particular, several new findings for novel verbs are reported. Critically, some of these results pro-

vide evidence against Pinker's hybrid dual-route theory, which predicts that regular neighbourhood density should have no effect on past tense inflection; on the contrary, these data, which show that irregular *and* regular neighbourhood effects were found when inflecting verbs to the past tense, tend to support the single-route perspective of past tense morphology.

References

Breyer, N. (1997). *Neurophysiological studies of semantic organisation in the human brain.* M.Sc. thesis, University of Oxford.

Lee, B. (1995). On the processing of regular and irregular inflections: The symbolic-connectionist debate revisited. In C. Koster & F. Wijnen (Eds.), *Proceedings of the Groningen Assembly on Language Acquisition.* The Netherlands: Centre for Language and Cognition, Groningen.

Pinker, S. (1991). Rules of language. *Science, 253*, 530–534.

Prasada, S., & Pinker, S. (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes, 8*, 1–56.

Contents

1 General Introduction **1**

1.0 Introduction 2

1.1 Regular and Irregular Verbs 5

1.2 Chomsky, Halle and Mohanan 9

1.3 Competing Post-Chomskian Theories 11

1.4 Conclusion 16

2 What is Connectionism? **18**

2.0 Introduction 19

2.1 Principles of Network Architecture 21

2.2 Learning 28

2.2.1 The Learning Algorithm 29

2.2.2 Bias 30

2.3 Three Boolean Problems and Their Connectionist Solutions 31

2.3.1 The AND Function 33

2.3.2 The OR Function 35

2.3.3 The XOR Function 37

2.4 Advantages and Disadvantages 42

2.4.1 A Sample Problem 43

2.5 A Brief Note about Rules 45

2.6 Conclusion 46

3 The Single-Route / Dual-Route Debate **47**

3.0 Introduction 48

3.1 Rumelhart & McClelland 48

3.1.1 The English Past Tense Model 50

3.1.2 Comparison with Data from Bybee & Slobin 52

3.1.3 Comparison with Data from Kuczaj 55

3.2 Critiques of R&M's Model 57

3.3 New Connectionist Models 61

3.4 A New Dual-Route Theory, and More Debate 66

3.5 Conclusion 68

4 Neighbourhoods **70**

4.0 Introduction 71

4.0.1 Prasada & Pinker (1993) 71

4.0.2 Lee (1995) 75

4.1 Analysis of Stimuli from Prasada & Pinker and Lee 78

4.1.1 Nearest Neighbour Analysis 79

4.1.2 Partword Analysis 83

4.2 Glushko (1979) and Marchman (1997) 88

4.3 A New Classification of Neighbourhood Density 92

4.3.1 Re-Analysis of Prasada & Pinker's and Lee's Stimuli 99

4.4 Conclusion 105

5 Behavioural Studies **106**

5.0 Introduction 107

5.1 Method 109

5.1.1 Subjects 109

5.1.2 Materials 110

5.1.3 Procedure 113

5.2 Results 115

5.2.1 Results for Pseudoirregular Stimuli 116

5.2.2 Results for Pseudoregular Stimuli 118

5.3 Discussion 121

5.3.1 Neighbourhood Density Effects 121

5.3.2 Effects of Presentation Modality and Context 123

5.4 Conclusion 123

6 Electrophysiological Studies **125**

6.0 Introduction 126

6.0.1 Jaeger et al. (1996) 127

6.0.2 Criticisms of Jaeger's Paper 134

6.0.3 Other Recent Studies 140

6.1 Experiment One — Real Verbs 147

6.1.1 Method 148

6.1.2 Results 156

6.1.3 Discussion 183

6.2 Experiment Two — Novel Verbs 188

6.2.1 Method 188

6.2.2 Results 191

6.2.3 Discussion 201

6.3 Comparisons Between Real Verb and Novel Verb Findings 204

6.3.1 Real Regular Verbs vs. Regularised Novel Verbs 205

6.3.2 Real Irregular Verbs vs. Regularised Novel Verbs 206

6.3.3 Discussion 210

6.4 Conclusion 211

7 Connectionist Studies **213**

7.0 Introduction 214

7.1 Method 216

7.1.1 Network Configuration 216

7.1.2 Training Corpus 216

7.1.3 Representation 219

7.1.4 Training 222

7.2 Results 224

7.2.1 Error Rates for Training Set 224

7.2.2 Generalisation to Unseen Stimuli 228

7.2.3 Regularity by Frequency Interaction 229

7.2.4 Internal Representations 238

7.2.5 Generalisation to Novel Forms 242

7.3 Discussion 252

7.4 Conclusion 254

8 General Conclusion **255**

8.1 Summary 256

8.2 Future Work 260

A Irregular Verbs — List α **263**

B Irregular Verbs — List β **265**

C Test Passage Used in Behavioural Studies **268**

D Cloze Passage Used in Behavioural Studies **269**

E Real Verb ERP Stimuli **274**

F Novel Verb ERP Stimuli **279**

G Familiarity Data for ERP Stimuli 284

H Nvalue Data for ERP Stimuli 286

I ERP Presentation Program 288

J Nvalue Data for Novel ERP Stimuli 291

K Irregular Verbs in Network Training Set 293

L Regular Verbs in Network Training Set 295

M PlaNet Network Training Script 308

N Regression Data on Hidden Units 310

References 315

List of Figures

1.1 A dual-route theory of past tense morphology 11

2.1 A fully connected 3-layer network 21

2.2 Generic representation of a single node 22

2.3 A sample node with inputs, outputs and weights shown 23

2.4 The sigmoid activation function 26

2.5 A simple single-layer network 33

2.6 A binary threshold activation function (where $\theta = 1$) 34

2.7 The AND problem 36

2.8 The OR problem 37

2.9 The XOR problem 38

2.10 A simple multi-layer network 39

2.11 A sample solution to the XOR problem 39

2.12 The XOR problem rendered linearly separable 41

4.1 Venn diagram with samples of each cluster in the /əʊ/ neighbourhood 93

5.1 Overall results for pseudoirregular stimuli across all paradigms 117

5.2 Modality × context interaction for pseudoirregular stimuli 118

5.3 Overall results for pseudoregular stimuli across all paradigms 120

5.4 Irregular × regular interaction for pseudoregular stimuli 120

6.1 56-channel electrode montage used in ERP experiments 155

6.2 Lateral verb type × hemisphere interaction between 240–280msec. . . 162

6.3 Frontal verb type × hemisphere interaction between 360–400msec. . . 163

6.4 Lateral verb type × hemisphere interaction between 920–960msec. . . 165

6.5 Midlateral verb type × hemisphere interaction between 920–960msec. 166

6.6 Posterior verb type × hemisphere interaction between 920–960msec. . 167

6.7 HighI vs LowI condition × hemisphere interaction for irregulars . . . 171

6.8 Rdensity × hemisphere interaction for irregulars 173

6.9 Rdensity × hemisphere interaction between 400–440msec. 174

6.10 Rdensity × hemisphere interaction between 440–480msec. 175

6.11 Rdensity × hemisphere interaction between 480–520msec. 175

6.12 Posterior Rdensity × hemisphere interaction between 600–640msec. . 176

6.13 Idensity × hemisphere interaction for regulars 178

6.14 Lateral Idensity × hemisphere interaction for regulars 179

6.15 Frontal Rdensity × hemisphere interaction for regulars 182

6.16 Late midlateral Rdensity × hemisphere interaction for regulars 183

6.17 Frontal condition × hemisphere interaction (120–160 msec.) 193

6.18 Frontal condition × hemisphere interaction (160–200 msec.) 194

6.19 Frontal Idensity × hemisphere interaction 197

6.20 Lateral Rdensity × hemisphere interaction 199

6.21 Frontal Idensity × hemisphere interaction 200

6.22 Lateral verbclass × hemisphere interaction 208

6.23 Frontal verbclass × hemisphere interaction 208

6.24 Posterior verbclass × hemisphere interaction 209

7.1 Composition of 16-bit binary sequence representation 222

7.2 Percentage of tokens correct by verb type 225

7.3 Regularity by frequency interaction across all five simulations 236

7.4 Total results for pseudoirregular verbs 243

7.5 Results for pseudoirregular verbs by network 244

7.6 Total results for pseudoregular verbs 246

7.7 Results for pseudoregular verbs by network 247

7.8 Idensity × subject class interaction for pseudoirregulars 250

7.9 Idensity × Rdensity interaction for pseudoregulars 251

List of Tables

1.1 The 10 most frequently-occurring English verbs 6

2.1 Boolean functions AND, OR and XOR 32

2.2 Sample network solution to the AND problem 34

2.3 Sample network solution to the OR problem 35

2.4 Sample network solution to XOR problem 40

2.5 Sample training set for palindrome network 43

4.1 NN analysis of Prasada & Pinker's stimuli 80

4.2 NN analysis of Lee's stimuli 81

4.3 Partword analysis of Prasada & Pinker's stimuli 84

4.4 Partword analysis of Lee's stimuli 85

4.5 Neighbourhood density analysis of English irregular verbs 97

4.6 Neighbourhood density analysis of English regular verbs - Part 1 . . . 100

4.7 Neighbourhood density analysis of English regular verbs - Part 2 . . . 101

4.8 Neighbourhood analysis of Prasada & Pinker's stimuli 102

4.9 Neighbourhood analysis of Lee's stimuli 103

5.1 Neighbourhood analysis of novel verb stimuli 112

6.1 Brain regions activated in Jaeger et al. study 129

6.2 Past tense tasks compared to reading tasks 132

6.3 Means of stimulus types used in experiment 1 149

6.4 Means of stimulus types used in experiment 2 189

7.1 IPA, DiSC and Moby alphabetic equivalents 220

7.2 Binary encoding based on Moby Pronunciator 223

7.3 Number (N) of verbs per category in regularity × frequency analysis . 234

A.1 British English irregular verbs—list α 264

B.1 British English irregular verbs—list β (Part 1) 266

B.2 British English irregular verbs—list β (Part 2) 267

E.1 High friend / low enemy real stimuli 275

E.2 High friend / high enemy real stimuli 276

E.3 Low friend / low enemy real stimuli 277

E.4 Low friend / high enemy real stimuli 278

F.1 High friend / low enemy novel stimuli 280

F.2 High friend / high enemy novel stimuli 281

F.3 Low friend / low enemy novel stimuli 282

F.4 Low friend / high enemy novel stimuli 283

G.1 Familiarity values for ERP real verb stimuli 285

H.1 Nvalues for ERP real verb stimuli 287

J.1 Nvalues for ERP novel verb stimuli 292

K.1 Irregular verbs used in network training 294

L.1 Regular verbs used in network training 296

L.2 Regular verbs used in network training 297

L.3 Regular verbs used in network training 298

L.4 Regular verbs used in network training 299

L.5 Regular verbs used in network training 300

L.6 Regular verbs used in network training 301

L.7 Regular verbs used in network training 302

L.8 Regular verbs used in network training 303

L.9 Regular verbs used in network training 304

L.10 Regular verbs used in network training 305

L.11 Regular verbs used in network training 306

L.12 Regular verbs used in network training 307

N.1 Coefficients and correlations for hidden units 311

N.2 Coefficients and correlations for hidden units 312

N.3 Coefficients and correlations for hidden units 313

N.4 Coefficients and correlations for hidden units 314

Chapter 1

General Introduction

Language is the dress of thought.

Johnson

1.0 Introduction

What is thought? How are mental processes enacted in the brain? For over three hundred years, two dominant schools of opinion on these questions have been Continental Rationalism (founded by Descartes and developed by—among others—Leibniz, Malebranche and Spinoza) and the British Empiricism of the great Englishman, Irishman and Scotsman: John Locke, Bishop Berkeley and David Hume. Rationalism considered, among other doctrines, that the nature of thought is the manipulation of symbols by rules (to dress the concepts in modern terminology), while Empiricism held that thought arises from the association of concepts due to 'resemblance', 'contiguity' or 'cause and effect'.

Psycholinguistics is the study of the mental mechanisms of language; its goal is an understanding of the way in which language is acquired, comprehended and produced. Within this field, the views posited by the Rationalists and Empiricists are still being applied and assessed in a linguistic framework.

Noam Chomsky, undoubtedly the most influential linguist of the 20^{th} century, formalised the Rationalist approach to language in his theory of 'generative linguistics' in the 1960s, in which all linguistic processing occurs by the use of symbolic rules. This theoretical perspective was very much in keeping with the 'If p then q' symbolic functionalism of computation during the same period. But the approach raises many difficult questions: how are rules implemented? Where do they come from? What *is* a symbol?

Developments in technology over the past half-century have led to a move away from symbol manipulation and towards subsymbolic connectionist approaches to computation and—by extension—thinking processes, which owe a debt to the Empiricist associationism of the 18^{th} century. A key issue arising from this view is whether rules are a necessary component of cognitive processes; connectionist networks, their proponents argue, provide the illusion of rule-driven behaviour without containing explicit rules. This approach solves the problems associated with symbols by simply doing away with the need for symbol manipulation. In the connectionist framework, learning and thought result from the exposure of a genetically-constrained system to a vast experience, by which (through predisposition) the system can extract 'rules' from its environment (see Elman et al., 1996, for a thorough review of these issues).

Language provides a meeting-point for these symbolic and subsymbolic perspectives, and one particular question has become a highly contentious and hard-fought issue for the different approaches: how do children and adults inflect verbs to produce appropriate past tense forms? An ever-increasing literature of behavioural data, combined with new technology, has led to the emergence of two dominant theories: on the one hand, Steven Pinker and his colleagues have developed a hybrid approach to verb inflection, with a Chomskian rule component for regulars and a connectionist-style lexicon for irregulars; on the other hand, connectionists (initiated by James McClelland and David Rumelhart) have argued that regular and irregular verbs can be processed within a single subsymbolic associative system.

Each theory makes different predictions about how people behave under certain

conditions. By designing experiments to test such predictions one can find evidence in support of—or against—each theoretical perspective.

This thesis contributes to the debate by means of experimental studies, focusing particularly on 'neighbourhood effects' in verb processing: that is, the differences in processing that result from phonological similarities between a verb and other verbs in the language. Each theory of inflectional morphology makes mutually exclusive predictions about precisely what effects (if any) such similarities might produce, and these hypotheses are explained and tested in the course of this thesis.

The experiments are conducted in three different areas of cognitive science: experimental psychology, electrophysiology and computational modelling. By means of experimental psychology, one can discover how adults behave; by electrophysiological studies one can investigate what is happening in the brain during processing; and by building connectionist models, data from humans may be compared with the behaviour of computational models trained to perform similar tasks.

But why the past tense? The reasons are perhaps best summed up by Pinker:

> The English past tense is the perfect site. The phenomenon is circumscribed and therefore tractable to study. Its history, acquisition, and patterns of use have been abundantly documented. It has obvious rule-like and memory-like features that serve as hurdles that any theory must clear. And it has brought out the best in contemporary theoreticians, with each side devising a clever, elegant, detailed, and surprising model. The past tense is the only case I know in which two great systems of Western thought may be tested and compared on a single rich set of data, just like ordinary scientific hypotheses. (1999, page 90)

In the remainder of this chapter, the patterns of inflectional morphology (*i.e.* how

the stem of a verb is changed to produce its past tense form) will be described, before the contrasting theories which attempt to explain these processes are introduced.

1.1 Regular and Irregular Verbs

There are several ways in which an English verb can be inflected to make a past tense form and, according to the pattern used, verbs may be broadly categorised as regular or irregular. Regular verbs undergo a process of suffixation whereby the stem form and one of the three allomorphs[1] /-ɪd/, /-d/ or /-t/ are concatenated to produce the past tense, or product (*e.g.* respectively 'paint' $\mapsto$ 'painted', 'bore' $\mapsto$ 'bored', 'steep' $\mapsto$ 'steeped'). The choice of suffix is determined according to three criteria, dependent on the final phoneme of the stem. These are (in order):

1. If the final phoneme is an alveolar stop[2], the suffix /-ɪd/ is used;
2. If the final phoneme is a voiceless consonant, use /-t/;
3. Otherwise use /-d/.

An irregular verb, on the other hand, undergoes one of several other inflectional processes to create its past tense form:

- no change (*e.g.* beat $\mapsto$ beat);
- vowel change (*e.g.* throw $\mapsto$ threw);
- final consonant change (*e.g.* build $\mapsto$ built);

[1] That is, phonologically distinct variants of a single morpheme.
[2] *I.e.* the plosives /t/ or /d/.

- final vowel / consonant cluster change (*e.g.* think $\mapsto$ thought);
- blends of regular and irregular inflection (*e.g.* keep $\mapsto$ kept, hear $\mapsto$ heard);
- total change (*e.g.* go $\mapsto$ went).

Irregular verbs are much less common in English than regulars: there are only about 160 simple irregular verbs in English[3], compared to many thousands of regulars. Despite having such a low 'type frequency' (as this phenomenon is known), irregulars tend to have relatively high 'token frequency' in the language; that is to say, although there are very few irregular verbs, they occur very frequently in the language. Indeed, the ten most commonly-occurring English verbs[4] are all irregular, as shown in Table 1.1 (where OPM stands for occurrences per million words of text).

Verb	OPM
to be	39,175
to have	12,458
to do	4,367
to say	2,765
to make	2,312
to go	1,844
to take	1,575
to come	1,561
to see	1,513
to get	1,486

Table 1.1: The 10 most frequently-occurring English verbs

Any model of past tense processing must account for certain grammar-level facts about verbs. Pinker (1999) isolates three such facts:

[3]There are approximately 280 simple *and compound* irregulars, of which about 80 are zero-frequency.

[4]According to Francis & Kučera (1982).

Stem–Past Similarity Regular past tenses are identical to their stem forms except for the addition of a short suffix. Irregulars (with only two exceptions[5]) also share much of their phonological material between their stem and past tense forms.

Change–Change Similarity When a stem is altered to create an irregular past tense form, the inflection often follows a pattern common to other irregular verbs. For instance, the vowel-change pattern /ɪ/↦/æ/ is a common irregular inflection (*e.g.* 'begin', 'sing', 'drink'); further, the /ɪ/↦/ʌ/ pattern of past participle inflection is also found in many of these verbs (*e.g.* 'begun', 'sung', 'drunk').

Stem–Stem Similarity Verbs which follow a single vowel-change process are often similar in their stem forms. For example, /ɪ/↦/æ/ verbs (such as 'drink', 'shrink', 'spring', etc.) tend to begin with a consonant cluster, and the sound following the vowel tends to be /ŋ/.

A model of past tense inflection must also be able to account for developmental findings about past tense acquisition, and behavioural differences between regular and irregular verbs.

For instance, it has been argued that a 'U-shaped' learning curve can be seen in the pattern of children's inflections during the process of acquisition (see, for example, Ervin, 1964; Brown, 1973). This learning curve appears to reflect an acquisition process which can be divided into three chronological and developmental stages, as follows.

[5] *Viz.* 'be'↦'was' and 'go'↦'went'.

In stage 1, children use past tense forms of only a small number of verbs. These verbs tend to be mostly high-frequency and irregular. On the occasions when children use past tenses at all, they tend to do so correctly. At this stage, it is suggested that children are storing and using separate lexical items, and show no evidence for rule-based inflection being used, or indeed anything other than rote learning of stem-past pairs.

In stage 2, however, various differences in inflection give the appearance of a linguistic rule having been learnt. The onset of stage 2 is marked by the appearance of overregularisation errors on verbs which had previously (in stage 1) been inflected correctly. In other words, when a child is presented with /sɪŋ/ he or she may now overregularise the stem to produce */sɪŋd/ (*i.e.* stem + regular suffix) or */sæŋd/ (inflected form + regular suffix), even though during stage 1 the child may have correctly said /sæŋ/ (Ervin, 1964). These errors only constitute a small minority of all inflections, and verbs may be inflected incorrectly and then correctly (or vice versa) within a single sentence (Marcus et al., 1992).

During this stage, children use a much larger number of verbs in the past tense than in stage 1, of which the majority are regular items. A linguistically important development in this stage is that a child can generate a past tense form from a novel word stem (*e.g.*, to use the classic example, /rɪk/ $\mapsto$ /rɪkt/). This cannot be due to memorisation of the stem–product pair, and therefore shows the ability to generalise the inflectional process to novel forms (Berko, 1958).

In stage 3, the acquisitional problems are solved, and regular and irregular verbs

are consistently used and inflected correctly.

This developmental profile is known as U-shaped learning as, if one were to plot correct verb use (on the y axis) against subject age (on the x axis), the curve would appear to be U-shaped.

1.2 Chomsky, Halle and Mohanan

Chomsky & Halle (1968) developed a theory of 'generative morphology' (refined by Halle & Mohanan, 1985) which attempted to account for the facts of regular and irregular inflectional morphology by positing rules for regular and irregular verbs. Such rules could incorporate the three types of similarity described above with ease. An example rule might be, for instance:

$$\text{If stem} = \text{consonants} + /\text{ɪ}/ + /\text{ŋ}/ \text{ then change } /\text{ɪ}/ \text{ to } /\text{ʌ}/$$

thus accounting for stem–past similarity, change–change similarity and stem–stem similarity in one rule. This rule-based theory has the advantage of being very concise; indeed, Chomsky & Halle managed to account for almost all irregular morphological patterns with only three rules.

With this theory, Chomsky and his colleagues argued that verbs constitute "a continuum of productivity and generality that extends from [regular] affixation of the '-ed' suffix...to total suppletion"[6], with irregular paradigms such as /ɪ/↦/ʌ/ between the two, whereby both regular and irregular inflection are rule-driven.

[6]Halle & Mohanan (1985), page 104.

In order to constrain grammar into very few rules, they had to posit rules which operated at the level of phonemic features (such as voicing) rather than at the level of whole consonants or vowels. To account for these rules of grammar and phonology, Chomsky proposed the principle of 'deep structure', whereby a word is stored in an abstract (and perhaps unpronounceable) form which is processed and then transformed into recognisable 'surface structure' for use. This means that a word can be stored in a compact form, and that one 'deep' representation can account for the different pronunciations of a single word depending on its sentential context: the final consonant cluster of the word 'hand', for example, may be pronounced as /nd/ out of context, as /ndʒ/ in 'hand you', as /m/ in 'hand me' or as /ŋ/ in 'hand care' (Lahiri & Marslen-Wilson, 1991).

However, it has been argued that the 'generative morphology' theory makes unreasonable claims about the development of linguistic skills in children. In the first place, the rules are based on a deep structure which requires, for instance, that children hear the word 'fight' and infer the underlying form /fiːXt/, having been presented with the surface form and extrapolated the rules which had generated it. Secondly, irregular verbs form 'fuzzy categories' which cannot be strictly delineated into different groups according to the rule by which they are inflected (see Bybee & Slobin, 1982, on page 12 below). As Pinker (1999) pointed out:

> Chomsky and Halle concede the problem and say that their grammar is only what children *would* construct if, hypothetically, they could hear the entire vocabulary in one sitting before figuring out the rules, rather than learning the everyday words first. But then it's not clear what their

> theory is a theory *of*—it is not, by their concession, a theory of how children acquire words or how real adults represent them. (Page 101)

1.3 Competing Post-Chomskian Theories

Early classical and functionalist explanations of the psychological processes involved in inflectional morphology held that there is a symbolic, default suffixation rule which produces the regular inflection of a verb stem unless the process is blocked by a list containing irregular verbs which have been learnt by rote (*e.g.* Kuczaj, 1977). This is depicted in Figure 1.1. Note the inhibitory link between the irregular lookup system and the rule component.

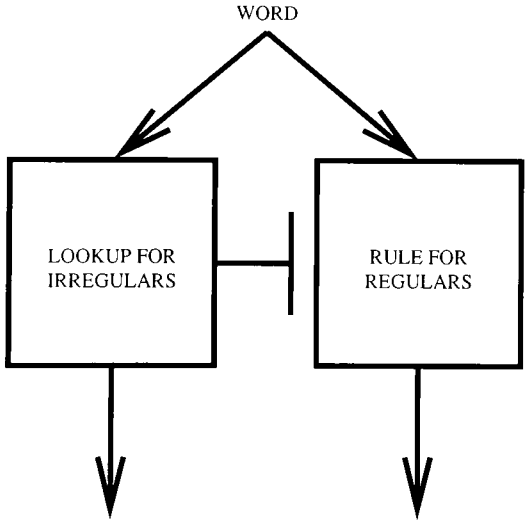


Figure 1.1: A dual-route theory of past tense morphology

When a verb stem is presented to the processing system, the list of irregulars is scanned for the target verb at the same time as the stem is presented to the rule component. If a match between the input verb and a verb in the list of irregulars is

found, the rule component is blocked and the verb is irregularised accordingly; if, on the other hand, no match is found, then the verb is processed by the (default) suffixation rule and regularised. This theory has a great strength in terms of the memory and processing power required: only the irregular stem/past tense pairs (which constitute only a very small minority of English verbs in terms of type frequency) need to be stored in memory; regulars need no representation, as their absence from the lexicon is sufficient for the correct inflectional process to take place. Besides using only a fraction of the memory required to store all verbs in the lexicon, inflection by this process would also be significantly faster than if all verbs had to be looked up by serial search in a list containing many thousands of word representations.

However, the theory fails to take into account experimental evidence both from developmental psychology and from psycholinguistics. For instance, it has been shown that young children can inflect novel verbs by irregular paradigms (Berko, 1958); that is to say, when presented with a hitherto unseen verb (such as 'to bling'), children do not always inflect it to the past tense by the regular paradigm (saying 'blinged') but can also generate past tense forms according to irregular paradigms (*e.g.* 'bling' $\mapsto$ 'blung', by analogy to verbs such as 'swing', 'fling', 'cling' and so on). Since the novel verb is not present in the rote-learnt list of irregulars, the theory would suppose that it would be regularised by the default suffixation process. Similarly, it has also been shown that phonological similarity to an irregular verb can lead to a real regular verb being incorrectly irregularised (Bybee & Slobin, 1982).

A logical adaptation of the classic dual-route theory was proposed by Bybee &

Moder (1983), who argued that a novel verb stem may be inflected by analogy to known irregular stems, and that the probability of a novel verb undergoing inflection according to an irregular paradigm is a function of its similarity to known irregular verbs. So, for instance, an inflection of /blɪŋ/ ↦ /blʌŋ/ is relatively likely, as around nine real verbs are inflected in this way, whereas a novel verb inflection such as /nɪst/ ↦ /nʌst/ is relatively unlikely, as all verbs with /-ɪst/ stems are regular. Between these two extremes, it is argued that there is a gradual decrease in the probability of irregular inflection (a 'similarity gradient'), and this approach may account for the historicophilological finding that regular verbs which are phonologically similar to irregulars have a tendency—over time—to become irregularised[7].

Pinker (1991) formalised this new approach by positing a modified dual-route model (also known as a hybrid model) in which, while regular verbs are still processed by a default, explicit rule-driven component, irregular verbs are stored in a pattern associator rather than simply in a rote-learnt list. Regulars are assembled as required, rather than being stored in memory, and the irregular component allows for a verb to be recognised as being irregular if a match is made between the input pattern and a pattern already stored in memory. *Partial* matches may also be triggered if, for example, a regular verb is presented to the pattern associator and found to be similar to a known irregular. Such partial matches may result in a false irregularisation of a regular verb. The model predicts that regulars which are similar

[7]For example, the irregular /niːl/ ↦ /nɛlt/ was, it is thought, entirely regular until the 19th century.

to known irregulars will be inflected more slowly than those which are dissimilar to irregulars, owing to partial matches in the pattern associator, even if they are eventually inflected correctly, and experimental work has found that this is indeed the case (Seidenberg & Bruck, 1990). Prasada, Pinker & Snyder (1990) also found support for the theory, since it predicts that high-frequency irregular past tense forms will be inflected more quickly than low-frequency irregulars (being stronger in memory and thus quicker to retrieve); this effect will not be noticeable for regulars if they are generated afresh rather than being stored in memory. Prasada et al. tested subjects on irregular and regular stimuli with high- and low-frequency past tense forms (controlling for stem frequency), and found that a frequency effect was indeed found for irregulars, but not for regular verbs. This study has been replicated several times, and has yielded consistently robust findings for irregulars, but variable results for regulars: low-frequency regulars have been found to be inflected more slowly than high-frequency regulars (although with less of a differential than for irregulars), at the same speed, and even faster (Seidenberg & Bruck, 1990; Ullman, 1993; Beck, 1997)[8].

The hybrid dual-route model also makes predictions about priming effects. In particular, regular past tense forms should prime their stem counterparts better than irregular past tenses should, since irregular stems and pasts are stored as separate (albeit connected) entities, and regular past tenses simply *are* concatenated stems. This has been borne out by behavioural studies (Stanners, Neiser, Hernon & Hall,

[8]See Beck (1997) for discussion.

1979). Furthermore, the effects have been replicated using cross-modal and subliminal priming methodologies (Marslen-Wilson, 1998)[9].

This words-and-rules theory therefore seems to account for much of the data in *irregular* inflectional morphology, particularly the similarity effects for irregulars found by Bybee & Moder. However, certain questions regarding *regular* verbs are left unanswered, perhaps the most important of these questions being: is the 'regular route' of the inflectional system necessarily rule-driven, or could regular as well as irregular inflection take place by analogy to known verbs? This is important because, if analogy can be shown to influence the inflection of regular as well as irregular verbs, the issue then arises as to whether it is necessary to have two processing routes at all.

Traditional, symbolic rule-based accounts of past tense inflection have—in the past 15 years—been joined by a quite different strain of models based on Parallel Distributed Processing (PDP or 'connectionist') theory, by which models have been developed which—their creators have argued—can account for the data on inflection of regular and irregular verbs in an economical, single-route system.

Each approach seems to have advantages over the other: symbolic theory can handle issues such as generativity with little difficulty, but relies on assumptions (such as symbol passing) which have not been found to have empirical basis in psychological reality. Connectionist models, on the other hand, attempt to model functions in a 'biologically plausible' manner, by using self-organising networks of 'neurons' with

[9]Also described in Marslen-Wilson, Hare & Older (1995), Orsolini & Marslen-Wilson (1997).

some 'genetic' architectural constraints, but with few hard-wired explicit rules.

Verbal past tense acquisition is one of the most thoroughly researched areas of cognitive science in recent years, and symbolic and connectionist theorists have taken different approaches. Symbolic theories of past tense inflectional morphology involve two mechanisms, explicit rules (for regular verbs) and memory for irregulars; connectionists, on the other hand, have conducted research using artificial neural networks to investigate whether past tense acquisition of English verbs can be modelled satisfactorily using a single system, and without the incorporation of explicit linguistic rules.

1.4 Conclusion

Verbs may be regular or irregular in their past tense forms. Each type of verb has properties which must be explained or accounted for by any successful psychological model of past tense processing. Early models posited rules for regulars and irregulars; subsequently, the two main schools of thought have become:

- dual-route theory, proposing separate processing: rules for regulars and an associative lexicon for irregular verbs; and
- single-route theory, proposing that all verbs can be processed within the constraints of a single subsymbolic system.

Each perspective has wide implications not just for past tense processing, but for the very nature of the implementation of thought.

Before the developments in this area can be discussed, however, it will be necessary to understand precisely what connectionism *is*. For this reason, the next chapter will explain the principles underlying connectionist modelling, and chapter 3 will look in detail at the influence of connectionism on research into the psychology of past tense processing in English.

Chapter 2

What is Connectionism?

2.0 Introduction

Connectionism is often described as being 'neurally inspired'. That is to say, its principles rest on knowledge of—or plausible assumptions about—the ways in which neurons work in the brain. McLeod, Plunkett & Rolls (1998) make five such assumptions, which will be treated here:

1. A neuron integrates and passes information about the set of neurons acting upon it as input. One or more neurons at a previous level of processing 'fire', producing excitatory or inhibitory signals which are passed across synapses to the dendrites of a receiving neuron as input. This neuron then acts upon the information it has received: if the input exceeds a certain threshold, the neuron will fire and, in doing so, will pass a signal along its axon and across synapses to act as input to other neurons in the chain.

2. A neuron also passes information about the level of its input: the amount of stimulation it receives from other neurons is passed on by being represented in its firing rate.

3. Neurons are layered. It has already been mentioned that neurons act in a chain, each one receiving information from a set of input neurons and converting that information into output to act on a subsequent set. But at a macrocosmic level, bundles of neurons are layered into different cytoarchitectonic structures[1]. In the process of reading aloud, for example, millions of neurons transmit and

[1]That is to say, physical brain structures.

transform information between the moment of light hitting the retina and the moment of utterance, but this process may be more helpfully understood in terms of physical brain structures: information from patterns of light falling on the retina is carried via the optic nerve to the lateral geniculate body, and thence via the optic radiation to the cortical visual areas in the occipital lobe. The activity from the visual cortex then stimulates other cortical areas, eventually reaching motor cortex, where it is transformed into the necessary stimulation to produce the appropriate vocal utterance by signals to the nose, mouth, throat and diaphragm. Each structure in this process will transform the signal it receives as input into a signal which is appropriate to be used as input for the next structure.

4. The influence of one neuron on another depends on the strength of the connection between them. More exactly, when a neuron fires, it increases or decreases the probability that a subsequent neuron will fire; the amount by which the probability is altered depends on the strength of the synaptic connection between them. A single neuron may itself produce input for many hundreds or thousands of others, but as the synaptic strengths between the neurons may differ, its activation may have a different effect on the different neurons receiving its activation as input.

5. Learning results from altering interneuronal connection strengths. This is a more contentious assumption, as neural learning systems are still not well un-

derstood. At a higher level, however, it is clear that an organism can learn—by experience—to change its behavioural response to a stimulus, and McLeod et al. (1998) argue that this may be the case at the level of individual neurons.

This chapter describes the basic principles of connectionist structure and shows how these five assumptions about the behaviour of neurons are reflected in the principles underlying neural network modelling.

2.1 Principles of Network Architecture

Connectionist models are constructed from nodes[2], with connections between these nodes; and as the housebrick has given rise to some of the most varied and complex architecture, so the complexity of connectionist models derives from their fundamental simplicity.

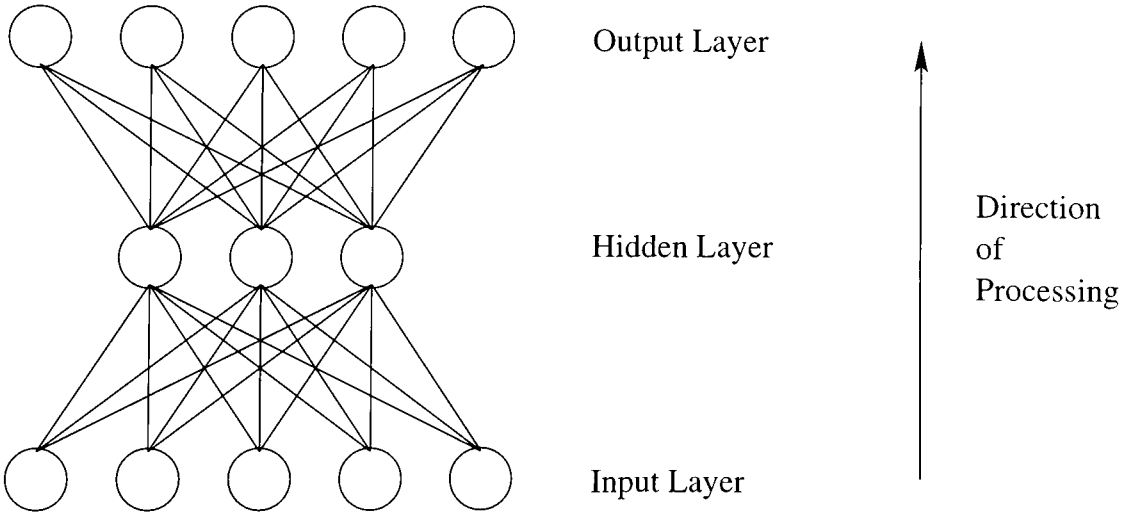


Figure 2.1: A fully connected 3-layer network

[2]The terms node and unit, in this context, are interchangeable in the literature.

Figure 2.1 shows a sample network, in this case a 3-layer feedforward system with 5 input nodes, 3 hidden nodes (*i.e.* nodes in an internal layer) and 5 output nodes. The nodes are represented by circles and the connections are represented by lines. All of the nodes in the input layer are connected to all of the hidden nodes; similarly, all nodes in the hidden layer are connected to all of the output nodes[3]. The input nodes receive information, in numerical form, from some external source. This is then transformed and passed to the hidden layer, which in turn transforms the information and passes it to the output layer, where it is again transformed and produced as output. The nature of the transformations and the manner of information transfer can be seen by looking closely at a single node.

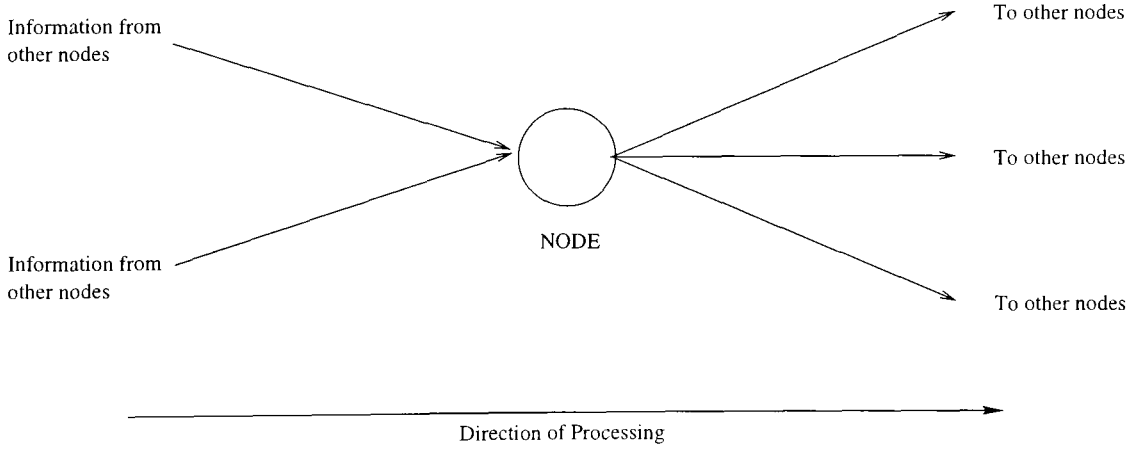


Figure 2.2: Generic representation of a single node

Figure 2.2 shows a single node, together with some inputs from—and outputs to—other nodes. Two important processes occur, even within this simple structure. First, the node transforms the information it receives by performing an activation function

[3]Hence the term 'fully connected'.

on the (usually summed) input. This function yields the activation of the node, which is the information to be passed on to other nodes in the chain. Secondly, the connections themselves have inherent strengths, also represented by numbers: these are generally known as weights. The activation of a node is multiplied by the weight of the connection it passes along, to produce input for the next node. We can see these two processes in action in Figure 2.3.

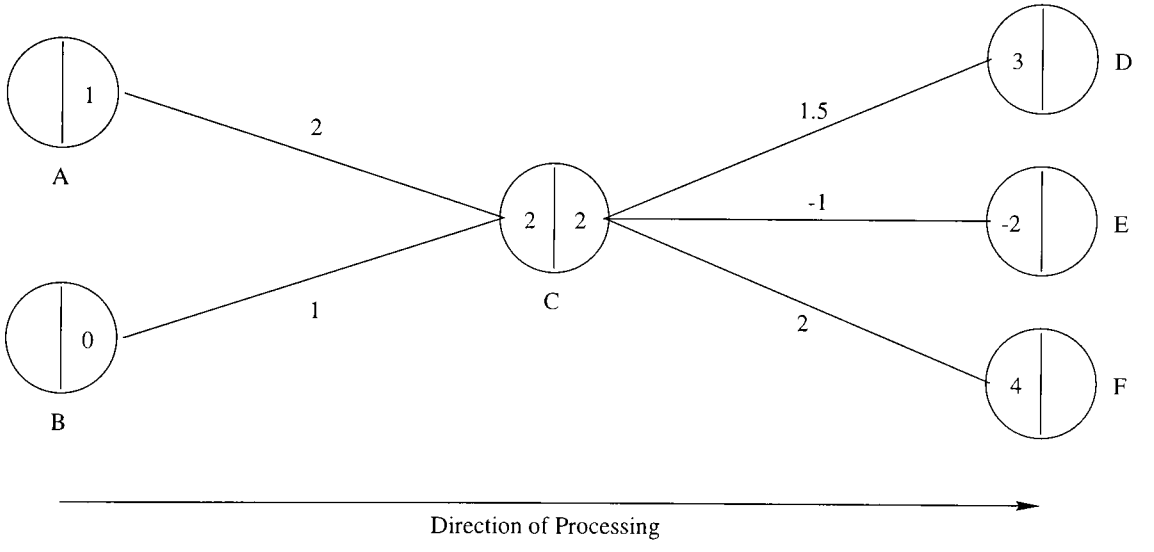


Figure 2.3: A sample node with inputs, outputs and weights shown

From Figure 2.3, we see that node C receives information from two nodes, (nodes A and B). The activation of node A is 1, and the activation of node B is 0. These activations are multiplied by the weights of the connections between them and node C, and added together to give the summed input to node C. In this example, the summed input is 2 (being $(1 \times 2) + (0 \times 1)$, and depicted in the left half of node C). Node C then applies an activation function to its input; in this case, an identity function is being used for the sake of simplicity, and the activation (depicted in the

right half of node C) is therefore the same as the input. This activation can then be passed on to nodes D, E and F. Note that in this case, the input to nodes D, E and F from node C will be 3, -2 and 4 respectively (being (2×1.5), (2×-1) and (2×2)). From this example we can see that node C produces negative (inhibitory) input to node E, and positive (excitatory) input to nodes D and F. Clearly, the greater the weight of a connection between two nodes, the greater the impact of the first node's activation on the second.

A node therefore performs two operations: it sums the inputs from the nodes acting upon it, and then applies an activation function to that summed input in order to produce its activation.

The summed input to receiving $node_i$ from $sendingnodes_j$ may be defined as

$$summedinput_i = \sum_j a_j w_{ij}$$

where a_j refers to the activation of the $sendingnodes_j$ and w_{ij} refers to the weights of the connections between (sending) $nodes_j$ and (receiving) $node_i$. Note that in the expression w_{ij}, the *first* subscript refers to the receiving node and the *second* subscript to the sending nodes[4].

A node may perform many different types of activation function on the summed input to produce its activation. In the case above (Figure 2.3), a simple identity function was used, where

$$a_i = summedinput_i$$

[4]This is somewhat counterintuitive, but standard in the literature.

However, in many neural networks[5], including those constructed and analysed in this thesis, a nonlinear activation function is used, specifically the sigmoid (*i.e.* S-shaped) function

$$a_i = \frac{1}{1 + e^{-summedinput_i}}$$

where e is the exponential.

The sigmoid activation function has great benefits in neural network modelling, for instance by restricting activation to levels between fixed points (0 and 1). No matter how much excitatory or inhibitory input such a node receives, its activation is constrained within clear boundaries. This reflects neural behaviour in that, obviously, a neuron can only fire as fast as its maximum firing rate, and cannot fire less than not at all.

Another advantage of the sigmoid activation function is that the gradient of the function is steepest at the point where the summed input is 0, and almost flat where the summed input is very large or very small. In practice, this means that most learning (*i.e.* the largest weight changes) takes place in 'uncommitted' neurons, while heavily-committed neurons are relatively unaffected.

This function may be represented graphically, as shown in Figure 2.4.

As Figure 2.4 shows, with an input above approximately 4, the node will be fully 'on' (that is, producing a '1' as output), and will be fully 'off' below approximately -4. But between those values, the node shows great sensitivity to different input. Around the most sensitive input level, where the gradient of the curve is steepest

[5]For which read *artificial* neural networks throughout.

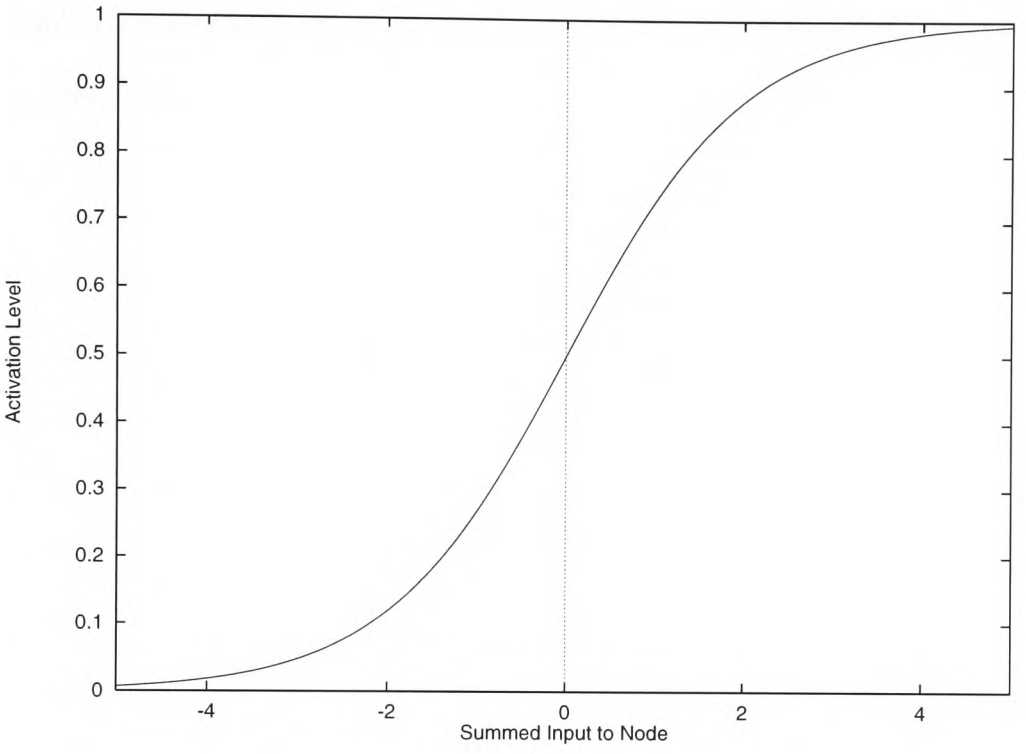


Figure 2.4: The sigmoid activation function

(that is, where the summed input is around zero), the node can discriminate between very slight differences in input. For example, inputs of -0.5 and +0.5 would produce activation of 0.378 and 0.622 respectively, a difference of 0.244 (nearly a quarter of the output range); however, inputs of +4.0 and +5.0 would elicit activations of 0.982 and 0.993, a difference of only 0.011 (just over 1% of the output range), even though the two pairs of inputs are exactly the same distance apart. This nonlinear transformation is responsible for much of the subtlety of connectionist network learning.

A node, then is merely a simple arithmetic operator combined with a transformation function. Alone, it is of little value. But when nodes are interconnected in a layered structure (a network), they can demonstrate a sometimes surprising ability

to solve complex problems. The key to a network solution lies in the strengths of the connections between nodes, and the power and flexibility of neural networks rests in their capacity to alter those weights by learning, as described in detail below.

Having explained the principles underlying connectionist modelling, we can now relate them to McLeod et al.'s assumptions about the behaviour of neurons in the human brain, described earlier in the chapter. Just as neurons integrate and pass information from 'sending' neurons (assumption 1), so a connectionist node sums the inputs acting upon it and passes information about that summed input to other nodes. Similarly, as neurons pass information about the *level* of their input (assumption 2), so connectionist units transform their input into activation by means of an activation function. In this case, of course, the activation is a number rather than the more physical activity of the firing rate of a neuron, but as firing rate appears to be one-dimensional, the numerical model is a valid one. We have seen that, like the brain, neural networks can be multi-layered[6] (assumption 3), and that the influence of the strength of synaptic connections between neurons (assumption 4) is also directly represented in neural networks by connection weights.

The fifth assumption cited above from McLeod et al. (1998) was that learning results from altering interneuronal connection strengths. But how does learning take place in a connectionist network?

[6]This will be considered in more detail on pages 37 *et seq.*

2.2 Learning

It is sometimes said that one of the great advantages of neural networks over other systems of cognitive modelling is that they have no need for explicit rules: Pinker & Prince (1988), for example, specifically stated that Rumelhart & McClelland's 1986 model of English past tense acquisition "contains no explicit rules" (p.173). This is not strictly true. Although connectionist models do circumvent the need for many rules, and still produce the appearance of rule-driven behaviour, many are in fact utterly dependent on some form of explicit generic rule for learning: a learning algorithm[7]. A model may begin with its connection weights set to random values and would thus, when presented with a pattern at input, produce a meaningless output pattern. It therefore needs to learn, in order to find a solution to the problem being modelled. In practice, the process of training a network to find a solution to a problem thus falls into two phases: presentation and learning. First, the model is presented with a training pattern at the input layer. This information is passed through the network and a response produced at output[8]. Secondly, the model needs to be shown whether its solution was correct or incorrect (and by what amount) in order for it to learn.

This learning takes place in three stages: an external system can provide the network with the correct (that is, desired) output for a pattern, and then the model can compare that with its own (produced) output, and adjust its connection weights

[7] See page 45 below for a fuller discussion of the nature of rules.

[8] A complete set of presentations—one of each input pattern—is known as an epoch or cycle.

such that it would be more likely to produce the correct output, were it to be presented with the same input pattern again. This system is known as 'supervised learning', and requires a learning algorithm which enables the model to calculate (and thence minimise) the amount of error it is producing. By implementing such an algorithm, a neural network can therefore learn by experience, giving it great power as a tool for modelling both the solution to a problem and the process by which that solution is reached.

2.2.1 The Learning Algorithm

One of the most common learning algorithms is known as the Delta (Δ) rule[9] (Rumelhart & McClelland, 1986). We have already seen how the activation of a unit is calculated by applying an activation function to the summed input to that unit, which is in turn determined by multiplying the activation of each of its input nodes by the weight of the connection between the input node and the receiving node, and summing those products. Therefore, if the receiving node produces an inappropriate (incorrect) level of activation, that error can be rectified by altering the connection weights acting upon it as input. If a weight is reduced between two nodes (say, node A and node B), the level of activation of node A will have less of an effect on node B; conversely, if the weight is increased, the activation of node A will have more of an effect on node B. This procedure may be expressed algebraically as

$$\Delta w_{ij} = (a_i(target) - a_i(actual))a_j\eta$$

[9]The term 'Delta rule' actually covers several variants of learning algorithm, of which the simplest is explained here.

where η (eta) is a constant representing the learning rate of the system[10]. The equation above shows that the amount of change to the weight of the connection between $node_j$ and $node_i$ is determined by calculating the difference between the desired activation of $node_i$ ($a_i(target)$) and the actual activation ($a_i(actual)$), and multiplying the result by the activation of the input node (a_j) and a constant parameter (η). Thus the Delta rule apportions 'blame' for the error onto the input units according to their activation levels, and distributes the weight changes accordingly throughout the network. Clearly, the greater the error and the greater the influence of the input unit, the greater will be the weight change of the connection. By means of this learning algorithm, the model attempts to find a set of weights which will produce the correct output for any of the input patterns on which it is trained. Of course, it is theoretically possible that the randomly-assigned set of weights provided to the network before training could produce correct solutions without the need for training at all; but in practice the model will probably require many presentations of the various input patterns in order to begin to match the desired outputs.

2.2.2 Bias

It should, perhaps, be noted at this point that many connectionist networks have an extra unit connected to hidden and/or output nodes, known as a 'bias' unit. The unit is unique in a model since it receives no input, either from other nodes or from an external source, and its activation is predetermined and immutable at +1 (in other

[10]The effect of η is thus to prevent overcorrection during network learning.

words, it is always 'on'). The strengths of its connections to other units, however, are fully flexible (just like all the others) and can change during learning. The bias node therefore serves two main purposes: first, if connected with a negative (inhibitory) weight to another unit, that unit must receive more input than it would otherwise need, in order to fire. The bias is thus effectively raising the threshold of the node on which it is acting negatively, rendering the node less active than it would otherwise be. Conversely, if connected to a unit by a positive (excitatory) connection, that node is made *more* active than it would have been.

2.3 Three Boolean Problems and Their Connectionist Solutions

Simple problems may require only simple networks and little training in order to be modelled successfully; in general, the more complex the problem, the more sophisticated the network must be (and the more training necessary) to find a solution. This may be demonstrated clearly if we look at how connectionism can be used to model three basic logic functions, known as AND, OR and XOR[11].

These functions may be expressed in Boolean terms[12] as in Table 2.1.

Columns 1 and 2 of Table 2.1 show the four possible combinations of 2 binary propositions (labelled in the Table as $\mathcal{P}_1$ and $\mathcal{P}_2$), one combination in each row: both propositions false; only $\mathcal{P}_2$ true; only $\mathcal{P}_1$ true; and both propositions true. Columns 3,

[11] 'Exclusive or'.

[12] Boolean notation uses the binary digits 0 and 1 to represent logical propositions, with 0 representing 'false', and 1 'true'.

$\mathcal{P}_1$	$\mathcal{P}_2$	AND	OR	XOR
0	0	0	0	0
0	1	0	1	1
1	0	0	1	1
1	1	1	1	0

Table 2.1: Boolean functions AND, OR and XOR

4 and 5 show the logical outcomes of the propositional pairs when a certain (labelled) function is performed on them. Looking at column 3, for example, we see that when the AND function is applied to the propositions, the function only returns a true (*i.e.* 1) value when both $\mathcal{P}_1$ *and* $\mathcal{P}_2$ are true—hence the name of the function. Similarly, the OR function returns a true value if either $\mathcal{P}_1$ or $\mathcal{P}_2$ (or both) is true, while the XOR function returns a true value if either $\mathcal{P}_1$ or $\mathcal{P}_2$ (but *not* both $\mathcal{P}_1$ *and* $\mathcal{P}_2$) is true. In formal notation, the three functions may be defined thus (AND by line 1, OR by line 2 and XOR by line 3):

$$(\mathcal{P}_1 \wedge \mathcal{P}_2) \tag{2.1}$$

$$(\neg\mathcal{P}_1 \wedge \mathcal{P}_2) \vee (\mathcal{P}_1 \wedge \neg\mathcal{P}_2) \vee (\mathcal{P}_1 \wedge \mathcal{P}_2) \tag{2.2}$$

$$(\neg\mathcal{P}_1 \wedge \mathcal{P}_2) \vee (\mathcal{P}_1 \wedge \neg\mathcal{P}_2) \tag{2.3}$$

To give examples of practical situations which may logically be reduced to these Boolean functions: a door with two bolts can only be opened if both bolt 1 *and* bolt 2 are open (AND function, column 3 in Table 2.1); a normally-sighted person can see if either the left eye *or* the right eye (*or* both) is open (OR function, column 4); and I, as a short-sighted person, can see clearly when I am wearing my spectacles *or* when I am wearing my contact lenses but *not*, importantly, when I am wearing both (XOR

function, column 5).

It is important that these functions, though simple, are thoroughly understood in order that they may be drawn on as example problems in neural network modelling.

2.3.1 The AND Function

Only a very simple, single-layer[13] network is required to model the AND function, with 2 input units and 1 output unit.

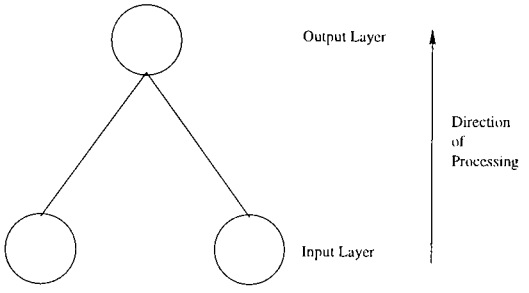


Figure 2.5: A simple single-layer network

Figure 2.5 depicts such a network. For the sake of simplicity, in this example it is assumed that the units use a binary threshold activation function, whereby a unit will fire if and only if the summed input to that unit is greater than or equal to the threshold. In this example, the threshold (θ) of all units is 1; a unit will therefore produce an activation of 1 if the summed input is 1 (or more), and will otherwise produce an activation of zero. The binary threshold activation function is graphed in Figure 2.6.

[13]To avoid confusion, it should be pointed out that although there are two layers of nodes in a so-called 'single-layer' network, the input layer and the output layer, this is nonetheless known as a single-layer network because there is only one layer of connections. This entirely contradicts the fact that a network with two layers of connections is known as a three-layer network.

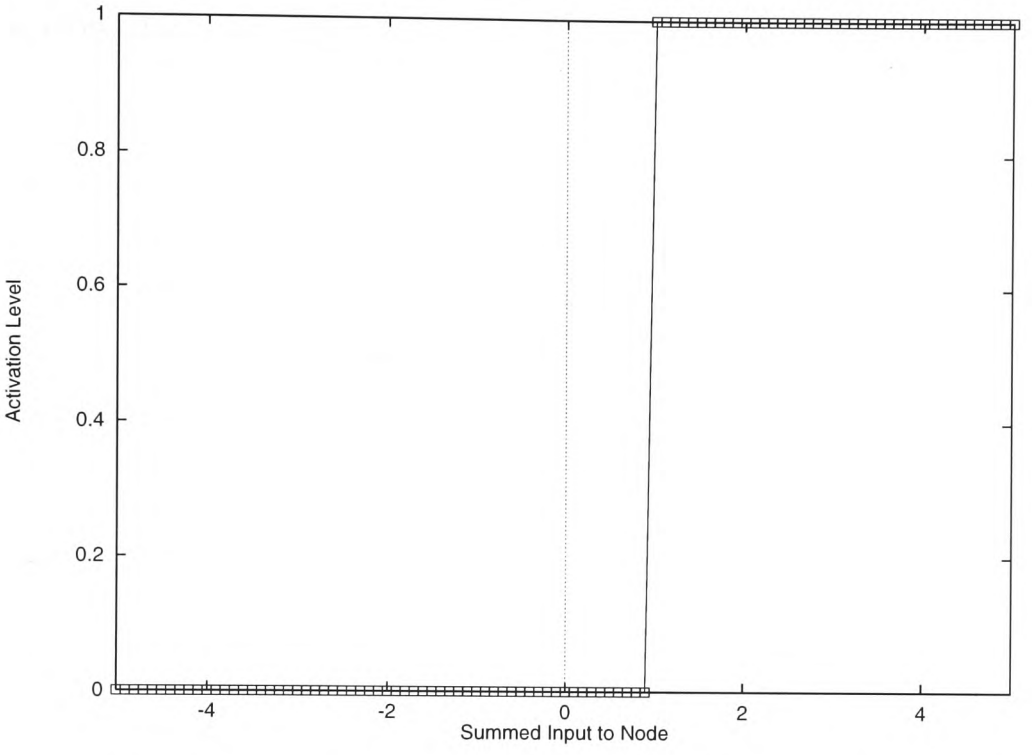


Figure 2.6: A binary threshold activation function (where $\theta = 1$)

In this example network, the connection weights are restricted to values between 0 and 1[14]. A possible solution to the AND problem might be for both of the connection weights to be 0.6; Table 2.2 shows why this network state satisfies the constraints of the problem.

a_1	a_2	$summedinput_3$	a_3
0	0	$(0 \times 0.6) + (0 \times 0.6) = 0.0$	0
0	1	$(0 \times 0.6) + (1 \times 0.6) = 0.6$	0
1	0	$(1 \times 0.6) + (0 \times 0.6) = 0.6$	0
1	1	$(1 \times 0.6) + (1 \times 0.6) = 1.2$	1

Table 2.2: Sample network solution to the AND problem

Table 2.2 shows the performance of the network when tested on the four input

[14]although other limits are often used in connectionist models, most commonly $-0.5 \leq w_{ij} \leq 0.5$ or $-1 \leq w_{ij} \leq 1$, where the gradient of the sigmoid function is steepest.

patterns, shown in the first two columns. Column 3 shows the summed input to the output unit, and column 4 shows the activation of the output unit for each pattern. As is shown in the Table, the output unit (unit 3) only 'fires' when both input units are on, and thus the network may be said to have successfully modelled the AND function. This example (where $w_{31} = w_{32} = 0.6$) is only one of a theoretically infinite number of potential solutions to this problem, as the model will be considered to have found a solution in all cases where the values of the two connection weights are individually less than, but collectively more than or equal to, the firing threshold of the output unit. To put the point algebraically, the following three criteria must be satisfied: $w_{31} < \theta$; $w_{32} < \theta$; and $(w_{31} + w_{32}) \geq \theta$.

2.3.2 The OR Function

Similarly, a structurally identical network may be said to have modelled the OR function successfully in all cases where $w_{31} \geq \theta$ and $w_{32} \geq \theta$, as illustrated in Table 2.3 below, where again $\theta = 1$.

a_1	a_2	$summedinput_3$	a_3
0	0	$(0 \times 1.1) + (0 \times 2.0) = 0.0$	0
0	1	$(0 \times 1.1) + (1 \times 2.0) = 2.0$	1
1	0	$(1 \times 1.1) + (0 \times 2.0) = 1.1$	1
1	1	$(1 \times 1.1) + (1 \times 2.0) = 3.1$	1

Table 2.3: Sample network solution to the OR problem

In this solution, sample values of $w_{31} = 1.1$ and $w_{32} = 2.0$ are used. Since both connections have weights above 1, if either input unit is activated (or if both are), the output unit will produce an activation of 1.

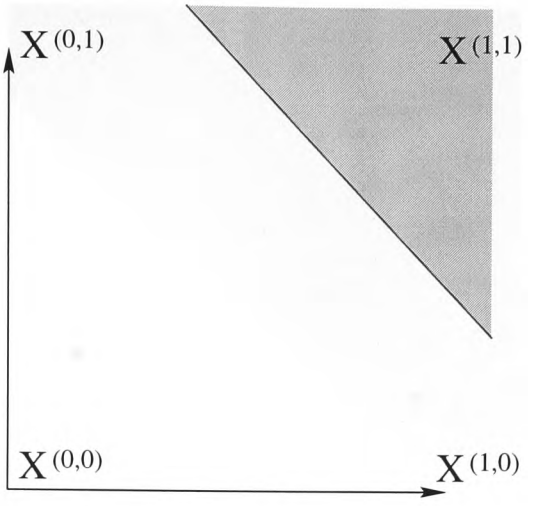


Figure 2.7: The AND problem

The AND and OR problems and their solutions may be depicted in graphical form, using the values of the two input units as axes; the four input patterns may therefore be represented as co-ordinates in two-dimensional problem space.

In Figures 2.7 and 2.8, values for the two input units are plotted on the x and y axes, and the four crosses in each figure represent the four possible input configurations to the network. The line in each graph separates the problem space into two areas, where on one side of the line the desired output is 0, and on the other (shaded) side the desired output is 1.

Because a straight line can be drawn across the problem space in each graph in such a way as to separate the two appropriate outputs (1 or 0) for each of the functions described above, they are both said to be linearly separable problems. The term 'linear separability' does not only apply to two-dimensional problems, however. In cases where there are 3 input units (plottable on a three-dimensional graph) the

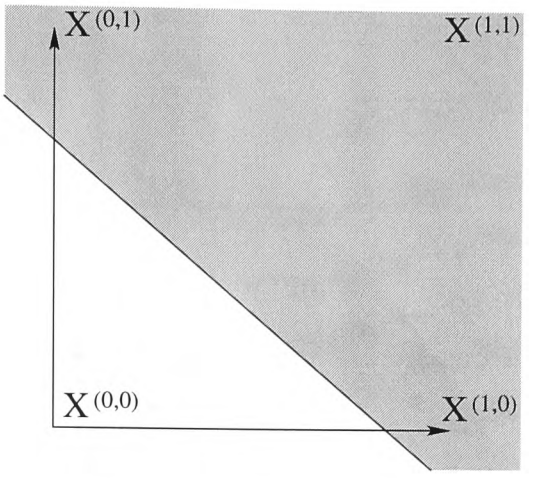


Figure 2.8: The OR problem

problem would be said to be linearly separable if the problem space could be separated into two parts by a plane, rather than by a line. Similarly, a hyperdimensional problem is linearly separable if the problem space can be partitioned by a hyperplane. This condition of linear separability is important since a single-layer, feedforward network (such as the one used in the two problems described above) can model any linearly separable function but, conversely, cannot solve any problem which is linearly inseparable (see, for example, Minsky & Papert (1969) for further discussion of this point).

2.3.3 The XOR Function

With reference to Table 2.1 on page 32 above, the problem space for the XOR function can be plotted as shown in Figure 2.9.

Unlike the first two Boolean functions dealt with in this section, it is clear that the XOR function is not linearly separable: that is, there is no straight line which can

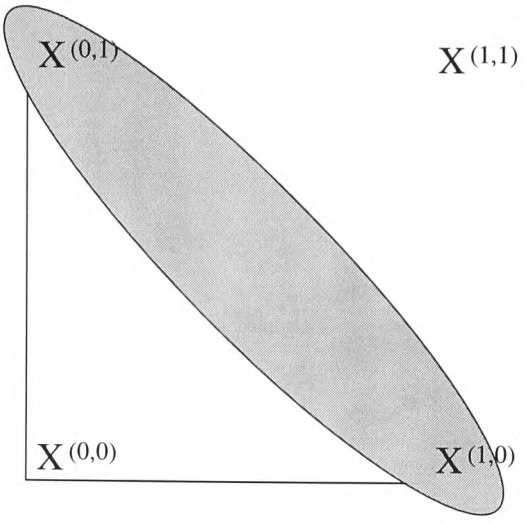


Figure 2.9: The XOR problem

be drawn on the graph to separate those parts of the graph which should yield an output of 1 from those parts which should produce an output of 0.

The problem can be solved, however, by using a multi-layer network, as shown in Figure 2.10.

Here we see that an internal or hidden layer (comprising unit 3 and unit 4) has been added to the simpler networks used in the earlier models. This hidden layer enables the problem space to be 'folded' (or re-represented) in order to render the problem linearly separable.

If it is assumed for the sake of example that the units in the hidden layer have the same binary threshold activation function as the output unit (unit 5), where $\theta = 1$, a sample solution may be found as in Figure 2.11. Table 2.4 shows the mathematics behind this sample solution.

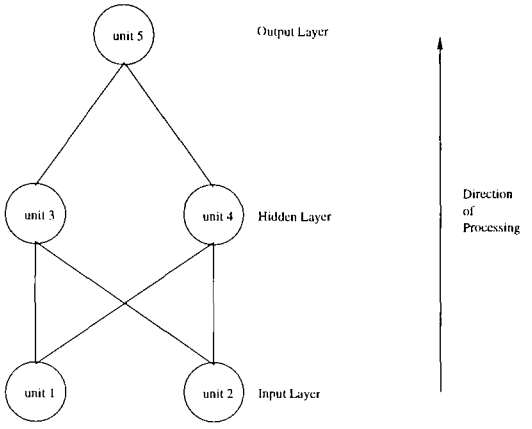


Figure 2.10: A simple multi-layer network

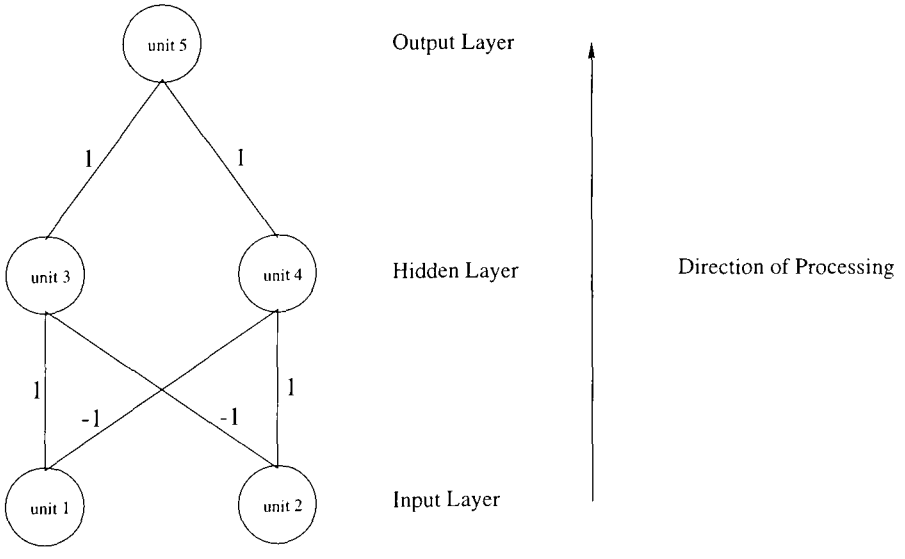


Figure 2.11: A sample solution to the XOR problem

a_1	a_2	$summedinput_3$	$summedinput_4$	a_3	a_4	$summedinput_5$	a_5
0	0	$(0 \times 1) + (0 \times -1) = 0$	$(0 \times -1) + (0 \times 1) = 0$	0	0	$(0 \times 1) + (0 \times 1) = 0$	0
0	1	$(0 \times 1) + (1 \times -1) = -1$	$(0 \times -1) + (1 \times 1) = 1$	0	1	$(0 \times 1) + (1 \times 1) = 1$	1
1	0	$(1 \times 1) + (0 \times -1) = 1$	$(1 \times -1) + (0 \times 1) = -1$	1	0	$(1 \times 1) + (0 \times 1) = 1$	1
1	1	$(1 \times 1) + (1 \times -1) = 0$	$(1 \times -1) + (1 \times 1) = 0$	0	0	$(0 \times 1) + (0 \times 1) = 0$	0

Table 2.4: Sample network solution to XOR problem

As we can see from Table 2.4, the first 3 input patterns are mapped identically onto the activation in the hidden nodes: $00 \mapsto 00$, $01 \mapsto 01$ and $10 \mapsto 10$. However, the fourth pattern, 11, produces an activation of 00 in the hidden nodes. In terms of the graph shown in Figure 2.9 above, it is as though the problem space has been folded so that the corner with input pattern (1,1) has been folded down over the diametrically opposite corner (0,0). If one imagines the graph drawn onto transparent paper and then folded, we can see how a straight line could then be drawn to partition the resultant problem space, as shown in Figure 2.12.

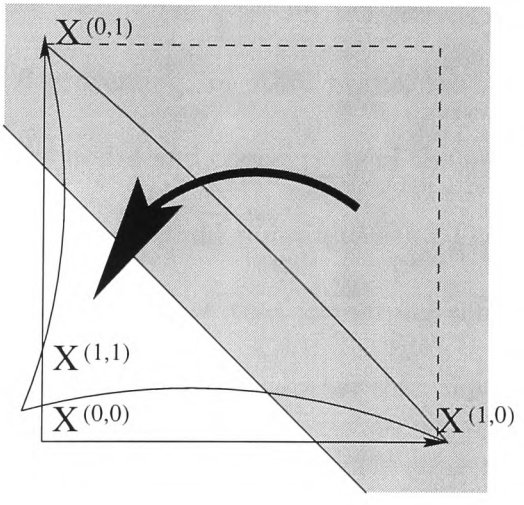


Figure 2.12: The XOR problem rendered linearly separable

Intuitively, since XOR is logically defined as $(\neg\mathcal{P}_1 \wedge \mathcal{P}_2) \vee (\mathcal{P}_1 \wedge \neg\mathcal{P}_2)$, the logic flow is captured by Figure 2.11: the hidden layer allows for the creation of $(\neg\mathcal{P}_1 \wedge \mathcal{P}_2)$ and $(\mathcal{P}_1 \wedge \neg\mathcal{P}_2)$ in the hidden nodes, and the connections between the hidden nodes and the output node replicate the OR function by each having a value $\geq \theta$.

From this graph, we see an example of a linearly inseparable problem being con-

verted to a linearly separable one, by the use of a layer of hidden nodes which change the nature of the problem by developing novel internal representations of the input patterns in order to discover a solution. It is this ability to create internal representations which gives multi-layer networks the power to model and solve highly complex problems.

2.4 Advantages and Disadvantages

Two points should be noted here about neural networks, by means of which they differ radically from more traditional tools for psychological modelling. The first is that they use distributed processing: in other words, the solution to the problem is distributed throughout the nodes and connections of the network. There is no single part of the network to which one could point and say "There is the specific rule for solving the problem", despite the fact that the model can give the epiphenomenal appearance of explicit rule-driven behaviour (other than the learning algorithm itself). In this important respect, connectionist systems are quite unlike symbolic models.

Secondly, a model is not taught at any stage what problem it is intended to solve. It is merely presented many times with inputs, calculates outputs, and is told what it should have produced. Note that these are not the same thing: being taught solutions to examples of a problem is not the same as being taught what the problem is. The network achieves a solution by trial and error, and by abstracting some underlying essence of the link between the presented inputs and the corresponding desired outputs.

Connectionist networks thus have a great strength: they can solve a problem for which the process of solution is not known to the modeller. In theory, for example, a neural network could design a glider (if it were able to test how good or bad its creations were) despite possessing no information about the laws of aerodynamics. However, this strength can also be a great weakness: since a network is not 'aware' (in any sense) of the problem itself, but only of exemplars which fit the problem, it may appear to find a solution, but actually have solved a different problem. To clarify this point, an example problem will be considered.

2.4.1 A Sample Problem

Say one wanted to create a network to recognise palindromic 8-bit input; that is, input which is the same read backwards or forwards. An hypothetical network such as this might be trained on the input/output patterns shown in Table 2.5.

Input Pattern	Desired Output
1 0 1 0 1 0 1 0	0
0 1 1 0 0 1 1 0	1
1 1 0 0 0 0 1 1	1
0 0 1 0 0 1 0 1	0
1 1 1 1 1 1 1 1	1
0 1 0 1 1 0 1 0	1
1 0 0 0 0 0 1 1	0
0 0 0 1 0 0 0 1	0

Table 2.5: Sample training set for palindrome network

Note that the network should produce an output of 1 when the input pattern is palindromic (*e.g.* the patterns on lines 2, 3, 5 and 6 of the training set shown above),

otherwise 0.

After training, the network may have successfully learnt these eight pattern pairs. But has it actually learnt to recognise a palindrome? We can test the network's ability to generalise to novel input by presenting it with a hitherto unseen pattern, say 1 0 0 1 1 0 0 1. This is, of course, palindromic, and our hypothetical network should therefore produce a 1 as output. But let us suppose that it actually produces a 0, thereby demonstrating that it has failed to learn the solution to the correct problem. So, what problem *has* it solved? Given the training set in Table 2.5 above, it may be the case that it has simply learnt to map the second bit of the input pattern onto the output node. If one looks carefully at the input patterns in the training set, it is clear that if bit 2 (that is to say, the second binary digit) in the input pattern has the value 1, then—as it happens—the output should also be 1, whereas if bit 2 is 0, the output should be 0. Given this 100% correlation in the input/output pattern pairs in the training set, the network may simply form a strong connection between the second input unit and the output unit and ignore everything else. This explains why our testing pattern 1 0 0 1 1 0 0 1 could elicit a 0 response: the model has learnt a solution, but to a different problem! It is therefore imperative that the modeller expose the network to enough variety of input to maximise the probability of a correct solution or, more specifically, to maximise the probability of a solution to the correct problem.

2.5 A Brief Note about Rules

The word 'rule' seems to be used by many researchers in entirely different (and often unthinking) ways: as has already been mentioned, some consider that connectionist networks contain no rules; others that they involve *implicit* rules, or a mixture of the two (see Davies (1995) for a detailed discussion of the issue). Rumelhart & McClelland (1986), for instance, said:

> We would... suggest that parallel distributed processing models may provide a mechanism sufficient to capture lawful behaviour, without requiring the postulation of explicit but inaccessible rules.

Similarly, Rumelhart, Hinton & Williams (1986) argued:

> [W]e do not assume that the goal of learning is the formulation of explicit rules. Rather, we assume [that] it is the acquisition of connection strengths which allow a network of simple units to act *as though* it knew the rules.

And again, Norman (1986) wrote:

> [A]lthough the system develops neither rules of classification nor generalizations, it *acts* as if it had these rules... It is a system that exhibits intelligence and logic, yet that nowhere has explicit rules of intelligence or logic.

More recently, researchers such as Mark Seidenberg have moved away from the notion of explicit rule, by talking of "the explicit notion of rule" (Seidenberg, 1989), and Bates & Elman (1993) specifically argue that connectionist networks *do* contain rules.

Certain aspects of the connectionist networks discussed in this thesis are assumed to be explicitly rule-driven: they use predetermined learning algorithms; they also de-

pend on the activation function being used, and the multiplicative weighting function for the connections—all of which are to some degree explicit rules. However, these rules are *generic*, rather than specific to the problem being modelled: the models do not contain any explicit rule relating directly to the problem. Any impression of *linguistic rule*-driven behaviour may therefore be considered to be an epiphenomenon of processing.

2.6 Conclusion

In this chapter, the basic theory and terminology behind connectionism have been described, including node function, network structure and learning principles. With these basic tools, psychologists have been able to develop new, testable theories in many areas of cognition. In the next chapter we shall see how connectionism has helped to shape the development of psychological theory in one particular area: inflectional morphology.

Chapter 3

The Impact of Connectionism on Inflectional Morphology Research

3.0 Introduction

Neural network (connectionist) models are, in essence, statistical inference machines which can represent rules in a distributed form by extracting patterns and implicit structures from their input environment. They are, therefore, very good at modelling analogy-based systems such as the irregularisation process described in chapter 1 above, whereby the irregularisation of a verb is based on its similarity to known irregular verbs. As stated above, if regularisation could also be shown to be based on analogy with known *regular* verbs rather than on an explicit default rule, then the explicit rule route may be unnecessary, and a single-route system such as a connectionist model could suffice to accommodate the data on past tense morphology. This chapter rehearses the key events in the single-route/dual-route debate, beginning with the first connectionist paper on the English past tense.

3.1 Rumelhart & McClelland

David Rumelhart and James McClelland are widely regarded as the Godfathers of connectionism. Even though the theories and mechanisms behind Parallel Distributed Processing have existed for some 50 years, its applications were not widely appreciated for some time (due, at least in part, to Minsky & Papert's (1969) pessimistic paper 'Perceptrons'). Indeed, the field only really came alive with the publication of Rumelhart & McClelland's 'Bible' of PDP in 1986.

In that work, Rumelhart & McClelland (hereafter R&M) developed a PDP net-

work which was designed to attempt to model the acquisition by children of the past tenses of English verbs. R&M took the view that language acquisition need not involve the formulation, evaluation and modification of a set of rules which could then be consulted by production or comprehension systems. They argued that 'implicit knowledge' of language could be stored in the weight matrices of connectionist networks, and addressed themselves to the issue of whether a connectionist network could model the data of past tense acquisition as well as—or better than—a symbolic, explicit-rule based system. The particular phenomenon they attempted to replicate was U-shaped learning and, more specifically, the types of overregularisation errors made in stage two of the acquisition process (see chapter 1 above; also Kuczaj (1977) below).

R&M argued that one should not be misled by the description of past tense acquisition as a three-stage process, in that there is no clear demarcation between the stages of U-shaped learning. For example, it has been shown that the transition between stage two and stage three extends over several years, and that different children show great variety within stage two (Kuczaj, 1977). Nonetheless, R&M attempted to model U-shaped learning and, more specifically, to model Kuczaj's finding that early in stage two children tend to make 'stem + suffix' overregularisations, whereas later in stage two 'past tense form + suffix' errors (*e.g.* */sæŋd/) are more common.

3.1.1 The English Past Tense Model

R&M's simulation consisted of a fully connected two-layer[1] pattern associator with encoding/decoding networks at input and output, to convert the phonological form of the input to a featural representation and *vice versa*.

R&M used a system developed by Wickelgren (1969) to represent the 'features' of their phonological input to the network. First, the phonological representation was converted into 'Wickelphones'. These are tripartite structures consisting of each phoneme of the word, flanked by its predecessor and successor. Thus the verb 'sing' would be converted from /sɪŋ/ to the 3 Wickelphones: '#sɪ', 'sɪŋ' and 'ɪŋ#', where word boundaries are marked by a '#'. Wickelphones are therefore context-sensitive; they are not, however, position-specific. These resulting triples were then converted into 'Wickelfeatures', which are featural decompositions of the Wickelphones; that is, Wickelfeatures take features from the predecessor, central and successor phonemes of the Wickelphone, in a sense compressing the data into a more efficiently-encoded format (see Wickelgren (1969) for a detailed discussion).

The network was trained by using verb stem forms as input and the correct inflected form as target output. The stimuli used were 420 verbs from the Kučera & Francis (1967) word list, ordered by frequency. R&M used 10 high-frequency verbs (2 regular, 8 irregular) and 410 medium-frequency verbs (334 regular, 76 irregular). 86 low-frequency verbs (72 regular, 14 irregular) were reserved for testing the model after training was complete.

[1]That is, containing two layers of nodes, with one layer of connections.

Learning took place according to the 'perceptron convergence procedure' (Rosenblatt, 1959). With this algorithm, each output unit is analysed after a learning trial and compared with the target output activation. Where the outputs match (*i.e.* $o_a = o_t = 1$ or $o_a = o_t = 0$)[2] no weights are adjusted; where a false hit occurs (*i.e.* $o_t = 0$, $o_a = 1$), all weights to that output unit are reduced by a small, constant amount η, as is the bias weight; conversely a false miss ($o_t = 1$, $o_a = 0$) will result in all connections—and the bias weight—to that output unit being increased by η. In this way, the probability that the actual output will match the target output (*i.e.* that $o_t = o_a$) will be increased for any output unit which did not match the target on that trial.

The model was trained using an incremental training regime. This means that, instead of the entire training set being presented every epoch (mass training), only a subset of the final training set was used to begin with; the size of the training set was then increased during training, until all patterns were being presented every epoch.

In this case, the network was first trained on the 10 high-frequency stimuli for 10 epochs; at that stage it was found to perform in a manner corresponding to a child in stage one of the acquisition process, that is to say, inflecting regular and irregular verbs correctly. The 410 medium-frequency verbs were then added (all at once) to the training set, in order to reflect the so-called 'vocabulary spurt'[3] (Brown, 1973); the model was then trained for a further 190 epochs.

[2]where o_t is the target output and o_a the actual output.

[3]The sudden increase in productive vocabulary noticeable at around 12–18 months.

R&M suggested that during the early part of this learning phase (from epoch 11 to around epoch 80), the output data matched the second stage of the human acquisition data, whilst the later part matched stage three. They therefore argued that their model accurately reflected the U-shaped learning curve of the irregular past tense: addition of the mid-frequency verbs at epoch 11 had no effect on the rate of learning regular verbs, but *did* interfere with the irregular verbs, leading to a dip in performance which never regained par with regular performance throughout training. Furthermore, analysis of the errors made on individual irregulars after the 10th epoch showed that between epoch 11 and epoch 40 the overregularisation alternatives (stem+ed, past+ed) were more popular with the network than the correct response, and this gradually diminished until the correct response was much the best preferred. They suggested that not only did the model capture the three learning stages—including the gradual transition from stage two to stage three—it also mimicked the progression from rote-based learning to rule-based learning to a blend of rule/analogy/rote-based learning in, as they described it, "a single, homogeneous learning mechanism".

3.1.2 Comparison with Data from Bybee & Slobin

R&M also analysed their results in the context of Bybee & Slobin's (1982) distinction between classes of verbs.

Bybee & Slobin (1982) divided irregular verbs into 9 different classes, depending on the inflectional paradigm they follow:

I Zero-markers, *e.g.* /hɪt/↦/hɪt/;

II Final /d/ ↦ /t/, *e.g.* /bɪld/↦/bɪlt/;

III Vowel-change + final dental, *e.g.* /hɪə/↦/hɜːd/;

IV Vowel-change - final consonant + final dental, *e.g.* /θɪŋk/↦/θɔːt/;

V Vowel-change where the verb already has a final dental, *e.g.* /ʃuːt/↦/ʃɒt/;

VIa Vowel-change from /ɪ/ to /æ/, *e.g.* /rɪŋ/↦/ræŋ/;

VIb Vowel-change from /ɪ/ (or /æ/) to /ʌ/, *e.g.* /dɪg/↦/dʌg/ or /hæŋ/↦/hʌŋ/;

VII All other vowel-changes *e.g.* /rʌn/↦/ræn/;

VIII Vowel-change with diphthongal sequence, *e.g.* /θrəʊ/↦/θruː/.

Three categories of regular verb were also distinguished:

1. Vowel or voiced consonant ending ↦ +/-d/, *e.g.* /stɑː/↦/stɑːd/;

2. Voiceless consonant ending ↦ +/-t/, *e.g.* /stɑːʧ/↦/stɑːʧt/;

3. /t/ or /d/ ending ↦ +/-ɪd/, *e.g.* /stɑːt/↦/stɑːtɪd/.

Type I verbs, or zero-markers, all end in /t/ or /d/, thus already (superficially) being past tense marked. Bybee & Slobin suggested that children learn very early in acquisition that most past tense forms of verbs end in /t/ or /d/ and therefore react correctly to zero-markers early; furthermore, they incorrectly generalise this

process to other verbs. Using an elicitation task, Bybee & Slobin found that verbs not ending in an alveolar plosive consonant were usually regularised, and that those that did were predominantly zero-marked.

They also studied vowel-change verbs (*i.e.* types III-VIII) by analysing the spontaneous speech of preschoolers between the ages of 1 and 5 years. The corpus they used (from Ervin-Tripp & Miller; Slobin; and Greenberg) was studied and the percentages of regularisations recorded for each class of verb. Their results showed a marked difference in the percentage of regularisations dependent on the class of vowel-change verb: types III and IV showed very few regularisations, whilst type VIII verbs showed by far the highest rate of regularisation.

R&M compared Bybee & Slobin's findings with their own model, and found that the network showed very similar results. First, it tended not to add a suffix to verbs which ended in an alveolar plosive, which therefore showed a strong no-change response. In the very early stages of learning (epochs 1 to 15), suffixation responses were stronger than no-change responses for stem-final alveolar verbs, but this rapidly adjusted and fitted the Bybee & Slobin data thereafter.

R&M also found that one subset of regulars were no easier to learn than type I irregulars, namely monosyllabic CVC[4] regulars which end in an alveolar plosive. They therefore predicted that human data would show that more no-change errors are made on these verbs than on polysyllabic regulars ending in alveolar stops, and argued that although the number of syllables in each stimulus was not explicitly

[4]Where C represents a consonantal phoneme and V represents a vocalic phoneme.

encoded, monosyllables never contain VCV Wickelphones, and there are no zero-marking verbs in English which do contain such a Wickelphone[5]. In other words, they argued that the model can learn never to zero-mark a stimulus containing a VCV Wickelphone.

They also compared Bybee & Slobin's percentage regularisation by verb type data (see above) with the model's output and found some similarities—type VIII verbs were again the most susceptible to overregularisation and types III and IV much less so, although types V, VI and VII yielded inconclusive results.

3.1.3 Comparison with Data from Kuczaj

Kuczaj (1977) analysed samples of the spontaneous speech of 15 children for the appropriate use / inappropriate use / nonuse of past tense inflection, with four hypotheses:

1 that irregular forms are acquired earlier than regular past tenses;

2 that the regular rule is more likely to be overgeneralised to irregular forms such as hit, shut, put, etc.

3 that partial regularity (*i.e.* verb types III and IV) blocks overgeneralisation;

4 that the two types of overregularisation (*e.g.* 'goed' vs. 'wented') have acquisitional relevance.

[5]This is untrue, since several compound irregulars *do* contain them, such as 'recast', 'beset', 'reset', and (at least in Southern British standard pronunciation) 'overbid', 'undercut' and so on.

He tested 14 subjects in a cross-sectional sample (recording 1 hour of their speech per week for 6 weeks), and 1 subject (Abe) as a longitudinal study (recording 1 hour per week between the ages of 2.4 and 4.1, and half an hour per week between 4.1 and 5.1) and found no evidence for the first and second hypotheses. He did find, however, that there was partial support for the prediction (hypothesis 3) that types III and IV verbs (to use Bybee & Slobin's terminology) exhibited the lowest proportion of overregularisation errors. Furthermore, he provided the interesting finding (in support of hypothesis 4) that, early on in the acquisition process, 'stem+ed' errors are more common than 'past+ed' errors, but that later on the tendency is reversed; that is, the proportion of past+ed forms increases relatively to the number of stem+ed forms as the child gets older.

R&M compared Kuczaj's results with their own simulation, and argued that their results were consistent with his findings. They found that verbs of types III, IV and VIII had a tendency to stem+ed overregularisations, whilst types II, V, VI and VII tended to past+ed errors. These latter types exhibited behaviour very much in accordance with Kuczaj's results, whilst the former never showed past+ed forms as often as stem+ed ones.

R&M finally tested their model against the 86 low-frequency verbs: of the 14 irregulars, all of the above-threshold responses were either regular, no-change or correct vowel-change generalisations. Of the 72 regulars, 6 produced no response, 48 were correctly regularised with only one response above threshold, and 18 produced either incorrect forms or two above-threshold responses. Of these 18, 12 produced the

correct response (albeit not uniquely) above threshold. They concluded from these results that the model had been generally successful at modelling the acquisition of the English past tense.

3.2 Critiques of R&M's Model

Lachter & Bever (1988) and—more thoroughly—Pinker & Prince (1988) provided a comprehensive critique of R&M's paper, and went on to make a wider criticism of PDP models generally as an experimental tool. Pinker & Prince (hereafter P&P) made several specific criticisms of the R&M model, which will be discussed here.

First, P&P argued that R&M had misunderstood the nature of the U-shaped learning curve, stating that their explanation of the results produced by the model was "demonstrably incorrect". They suggested that R&M had forced the overregularisation effect by using a very small input with many irregulars and then suddenly switching to a much larger set with many regulars. It is interesting to note that MacWhinney & Leinbach (1991) pointed out that both P&P and R&M seemed to be mistaken as to the nature of the human data, arguing that "there *is* no observational or experimental evidence for an across-the-board U-shaped curve" (page 130). They argued that the onset of the overregularisation period is very different for different stems, and that a global U-shape is in question; this argument has experimental and observational support from several sources (*e.g.* Brown, 1973; Kuczaj, 1977; Marcus et al., 1992)[6].

[6]See Marcus et al. (1992) and Plunkett & Marchman (1991) below for a more detailed argument.

Secondly, it will be recalled that children sometimes produce errors in which the regular suffix is added to an irregular past tense form (*e.g.* 'ated'). R&M's model successfully captured this pattern of error, but P&P argued that it did so for the wrong reasons, namely by concatenating the regular suffix with an already altered stem, whereas a child would make the error—according to P&P—by feeding into the mechanism a form already in the past tense (*i.e.* using a past tense as 'present tense' input). Evidence to support this view comes from Kuczaj (1977) and MacWhinney (1975), who argued that production of errors such as 'wenting' suggest that the irregular past is not always *coded* as a past. Similarly, errors such as 'jumpeded' suggest that the same could be said for regulars.

Thirdly, irregulars ending in an alveolar stop consonant were overregularised less often than other irregulars. P&P accorded this to 'misperception', whereby a stem form is perceived to have a past tense marker already in place, leading to a blocking of the regular concatenation process; the R&M model, on the other hand, provides the same data by means of the Wickelphonological coding scheme of the training set: all irregular verbs ending in /-t/ or /-d/ will share the Wickelphone '-t#' or '-d#', thereby emphasising the cohesion of the irregular subgroup.

P&P also provided other arguments against the method of sound representation in the model (the use of Wickelphones). First, they pointed out that, on the one hand, some words could be represented in exactly the same way to the network if Wickelphones were used. The classic example is 'algal' vs. 'algalgal', two adjectives from the Australian native language Oykangand (Sommer, 1980): using Wickelphonology,

these would be represented identically as input to a network, both containing only the Wickelphones '#æl', 'lgæ', 'gæl' and 'æl#'.

Secondly (and conversely), the words 'slit' and 'silt' would be encoded very differently (since they share no Wickelphones) even though they are phonologically very similar.

Thirdly, P&P argued that a representational system such as the R&M model should "exclude the impossible" (p. 99). No language uses the inflection pattern 'stem $\mapsto$ mirror of stem' (*e.g.* 'bar' $\mapsto$ 'rab'); however, the R&M model could, they claimed, learn such a transformational paradigm. This seems to be a specious argument: just because no known language uses a 'mirror' inflectional paradigm does not necessarily mean that such a language could not exist; besides, humans are quite capable (with practice) of learning to speak or write backwards at natural speeds, thereby performing such transformations with ease.

Fourthly, P&P further argued that the R&M model failed to capture significant generalisations about English sound structure. For example, English native speaker intuition suggests that, if 'Bach' (/baːx/) were a verb, the past tense would be /baːxt/ as opposed to /baːxd/ or /baːxɪd/. The R&M model would find this difficult since it has no clear system for featural representation.

Their final criticism about Wickelphone representation was that the model does not have the capacity to tell the difference between phonologically-identical stems, such as 'ring' ($\mapsto$ 'rang'), 'ring' ($\mapsto$ 'ringed') and 'wring' ($\mapsto$ 'wrung').

But perhaps their greatest criticism of the model was simply that it was not very

good at what it was supposed to do. Of the 72 regulars on which it was tested, 20 were incorrectly processed, with 6 (of those 20) producing no output at all. P&P argued that the errors made by the network were frequently unlike human errors (*e.g.*/meɪl/↦*/mɛmbəld/). They attributed much of the problem to the fact that the model did not, and could not, *modify* the input to create the output.

In conclusion of their critique, P&P expressed their objection to R&M using a single mechanism for the inflection of regulars and irregulars, arguing that there should be a distinction between formal rules and memorised exceptions. In particular, they pointed out that the regular pattern operates less through family resemblance than the minor (irregular) patterns; indeed, it does not need to, since it applies by default.

Marcus et al. (1992) analysed the 11,521 irregular past tense utterances made in the spontaneous speech of 83 children. They made seven main findings, which are worth recording here in support of Pinker & Prince's hypothesised distinction between irregular and regular processing routes:

1. Quantitative support was found for the fact that there is an extended period of correct verb inflection before the first error.

2. Overregularisation errors were relatively rare, suggesting that there is "no qualitative defect in the children's grammar which must be unlearned".

3. Such errors occurred at a low, constant rate between 2 years and 5–6 years, decreasing very slowly and affecting most irregulars.

4. Overregularisations first appeared when children began to mark regulars reli-

ably.

5. The more often a parent used an irregular token, the less often the child overregularised it.

6. There was no correlation between overregularisation and the number (or proportion) of regular verbs in parental speech, or in the child's speech or vocabulary, and they therefore argued that a symbolic model could not be replaced by a PDP one in which changes in input statistics lead to rotelike or rulelike behaviour.

7. Finally, they asserted that irregular verbs are protected from overregularisation during acquisition by strong neighbourhood clusters of similar irregulars, but are not attracted to the regular pattern by clusters of similar-sounding regulars.

This final point (point 7) will be returned to and examined empirically in the experimental chapters below.

3.3 New Connectionist Models

Plunkett & Marchman (1991) developed a new connectionist model, designed to address P&P's critique of R&M's network simulation by improving upon the methodology of the original network, and attempted to model a U-shaped pattern of verb acquisition using a specially-designed artificial language. This language consisted solely of word strings 3 phonemes in length; these strings were all either CVC, VCC or CCV in form. Each input unit coded for one of 6 phonemic features for one of the three phonemes in the 'word'. Four types of inflectional paradigm were learnt: regular, vowel-change, zero-marking and total change. Using a corpus of 500 'verbs',

a single-layer perceptron was trained with the perceptron convergence procedure (as in the R&M model). Since this was unable to find a solution to the mapping problems, Plunkett & Marchman then developed a multi-layer network, trained it using backpropagation as the learning algorithm (Rumelhart et al., 1986), and found that this model successfully learnt all four types of past tense mapping. Plunkett & Marchman also found that the model exhibited micro-U-shaped learning on most of the irregular patterns (but not all—consistent with human data[7]) without manipulating the input corpus. Despite having used a very restricted methodology—an artificial language only consisting of stimuli 3 phonemes long—Plunkett & Marchman made a particularly interesting finding, by varying the type and token ratios in the input corpus. It was found that optimum learning occurred when the ratios were closest to real English structure. For instance, arbitrary mappings were learnt best when the type frequency was low and the token frequency high (like /biː/↦/wɒz/ or /gəʊ/↦/wɛnt/ in English). Similarly, identity mappings required low token frequencies, whilst vowel-change paradigms were best learnt with high token frequencies. Furthermore, inflectional paradigms were also learnt more easily when phonological subregularities were added to the artificial language's structure, leading Plunkett & Marchman to suggest that it is these subregularities which make irregular mappings learnable. It should be noted, however, that there are many irregular verbs in English (around 30) with unique paradigms of inflection (at least in terms of rime[8]), so there

[7] *Vide supra* Kuczaj (hypothesis 3).

[8] Throughout this thesis the word 'rime' is used specifically to mean the combined nucleus and coda of a word (in other words its final vowel/consonant cluster)—usually of its stem, unless specified

is presumably no suggestion that phonological subregularities are *necessary* for an irregular item to be learnt.

MacWhinney & Leinbach (1991) also attempted to resolve the criticisms of Rumelhart & McClelland's model by constructing a new connectionist implementation of verb learning, this time using English rather than an artificial language. They argued that many of P&P's criticisms, while reasonable, were due entirely to the Wickelfeature binding system used in the R&M model. The MacWhinney & Leinbach model, on the other hand, involved a totally different type of input representation, output system, network architecture, learning algorithm and training corpus.

The input representation involved a two-part syllabic template and a feature coding system. The template had a left-justified CCCVVCCCVVCCCVVCCC section, in which a whole word could be represented, and a right-justified VVCCC section, in which only the nucleus and coda of the stimulus could be represented. The use of a two-part template of this kind strengthened the influence of the rime of a stimulus in the past tense mapping process, a concept supported by human experimental data[9]. Stimuli were coded into this template by featural coding for each phoneme according to the UNIBET system (MacWhinney, 1990). For instance the verb stem 'bet' would be represented by bCCEVtCCVVCCCVVCCC in the left-justified segment, and by VECCt in the right-justified segment of the template. There were also five units—of which only one could be activated at any one time—representing the inflectional out-

otherwise.

[9]See, for example, Marchman (1997).

put demanded of the system, *i.e.* one each for present, past, past participle, present participle and 3rd person singular present. In this way, MacWhinney & Leinbach attempted to counter P&P's claim that separate modules would be necessary for the production of different suffixes, thereby denying the network the capability of learning cross-suffix regularities[10]. The output system of the network was equivalent to the left-justified segment of the input template. Between input and output were two hidden rows, each consisting of 200 units. The network was fully connected, and there was an extra set of one-to-one connections mapping the left-justified input directly onto output. MacWhinney & Leinbach used the more powerful backpropagation algorithm (Rumelhart et al., 1986) in place of the perceptron convergence procedure, and training consisted of presenting a total of about 1.3 million verb stems over 24,000 epochs, the frequency of presentation being weighted to match the actual relative frequencies of the stimuli (from Francis & Kučera, 1982). The least frequent 10% each of regulars and irregulars were reserved for testing the model.

MacWhinney & Leinbach found that the model was very successful at learning the irregular past tense, although it did not demonstrate the 'stage one' early correct production of irregular inflections. The stem form was typically produced during the first 40–80 epochs, whilst zero-marking irregulars tended not to be overregularised until around the 200th epoch. Between epoch 200 and epoch 3,000, the network explored various vowel-change paradigms and different types of regular suffixation, and was mostly stable by the 3,000th epoch. For most verbs, the final 10,000 epochs

[10]See also Plunkett & Juola (1999).

were error-free, 'bled', 'brought' and 'thought' being the last to be learnt. 11 low-frequency irregulars remained unlearnt by the end of training. The model was also very successful at learning the regular past tense: this was 99% correct by the first test point (epoch 4,000), and perfect by the third (epoch 12,000).

The model seemed to answer P&P's criticisms—on the whole—very successfully. It successfully modelled all aspects of U-shaped learning except early correct usage, and comfortably solved the 'ated' and 'algalgal' problems presented by P&P. The network produced 'ated' when 'ate' was presented at input, in the same manner as it was hypothesised that children produce such errors, but interestingly it also produced 'inflected form + ed' errors spontaneously, by the simultaneous application of vowel-change and suffixation paradigms. The model also successfully captured aspects of phonological regularities by producing, for example, */sʌŋd/ rather than /sæŋ/, or */eɪtən/ instead of /eɪt/.

MacWhinney & Leinbach ran a second simulation, identical to the first except that the 'correct' inflection of each input was its mirror-image (*i.e.* the 'brag $\mapsto$ grab' problem, as they incorrectly termed it). After 24,000 epochs the network was only correct on 15% of the stimuli, and this network would therefore appear to conform to P&P's insistence that a model should "exclude the impossible".

A final auxiliary simulation was also run to establish whether a network could resolve the 'ring'/'wring' homophony problem: it was easily solved with the addition of a few semantic feature coding units, thereby rendering each stimulus nonunique. This is hardly surprising: if you make two similar-sounding stimuli look different to

the model, it can tell that they are different.

MacWhinney & Leinbach concluded that Pinker & Prince's criticisms were answerable, and that they should therefore not be taken as a critique of connectionist models in general, but rather the R&M model in particular.

3.4 A New Dual-Route Theory, and More Debate

As a result of the success of connectionist models described above, Pinker (1991) presented a modified (hybrid) version of the traditional symbolic account of past tense inflectional morphology, with one change, namely that instead of a rote-learnt list being responsible for irregular verb learning, there is now an associative network to map irregular stem forms onto their past tenses. A word presented to the system is processed in parallel by the associative network and the rule component. If a match is found in the network, an inhibitory signal is sent to the rule component, and the word is irregularised. If, however, no match is found in the network (that is, if the threshold to 'fire' a response is not reached), the word is inflected by the default route: the regular suffixation process.

Daugherty & Seidenberg (1994) considered two questions which were thought to be of central importance in the symbolic–connectionist debate: first, Pinker (1991) concluded that any past tense theory will have to include a rule component—is this true, or could the phenomena taken as evidence for the rule be elicited from within a connectionist network? Secondly, would such a model simply be an *implementation* of symbolic theory? Daugherty & Seidenberg presented a simple connectionist model,

and argued that the behavioural patterns Pinker used as evidence for the rule could in fact be modelled by PDP networks, therefore being consistent with both theories; more importantly, they showed that their network behaved in a way which standard or hybrid symbolic theories do not predict, and moreover that this behaviour is consistent with human psycholinguistic data.

The model took a nucleus-justified CCCVVCCC template, in which each phoneme was coded at input by 15 binary articulatory feature units. Between input and output was a 200-unit hidden layer which was fully connected, and output followed the same structure as the input. Training—as in MacWhinney & Leinbach's model—consisted of comparing the network's actual output with the target output (*i.e.* the inflected form of the input stem) and altering the weights of the connections according to the backpropagation algorithm.

In the first of two simulations, the training corpus consisted of all monosyllabic present–past tense pairs with a Kučera-Francis frequency greater than 1 (N=309 regulars, 104 irregulars). 112 regular verbs with a Kučera-Francis frequency of 1 were kept in order to test the network on its ability to generalise to novel forms. The stimuli were presented as a function of the logarithm of their frequency: in this case, the most frequent stimuli were presented once per epoch, the least frequent once per 100 epochs. The network reached equilibrium after 400 epochs, having learnt 100% of the regular verbs and 84% of the irregular forms. When tested for capacity to generalise on the unseen regulars, the model produced 77% correct responses. Daugherty & Seidenberg found that the model performed worse on 'regular inconsistent' verbs (*i.e.*

regular verbs which have many irregulars in their rhyming neighbourhood[11]) than on completely regular patterns (that is, regulars with no irregulars in their rhyming neighbourhoods), although they did not say whether the difference was significant.

They also investigated whether there was an interaction between frequency and regularity, on the basis that an effect of word frequency on production of irregulars but not regulars could be taken as evidence for a dissociation between irregular and regular processing[12]. However, they found no such interaction.

They then retrained the model on a different corpus (n=309 regulars, 24 irregulars), attempting to reflect the proportions of irregular to regular verbs consistent with English as a whole. Again the network learnt 100% of the regulars, and this time learnt 92% of the irregulars (all bar 2), with 94% of the unseen verbs being (correctly) regularised. Analysis for consistency and frequency effects showed that the model clearly *did* perform worse on 'regular inconsistent' verbs, and that there was a clear interaction between frequency and regularity. These results suggested that the 'explicit regularisation rule' route was *not* a necessary component of an inflectional morphology model, and that the model was no mere implementation of symbolic theory since it contained no explicit linguistic rule.

3.5 Conclusion

This chapter has shown the back-and-forth nature of the inflectional morphology debate since the original R&M model in 1986, and provided a background to the

[11]For example, 'bake' (cf. 'make', 'take', 'wake', etc.).

[12]See Ellis & Schmidt (1998), and chapter 7 below, for discussion.

field up to the mid-1990s. During that period, symbolic theories have been refined, and connectionist models have become more powerful, more accurate and better implemented. The next chapter looks more closely at what Daugherty & Seidenberg called 'consistency effects', and provides an analysis and criticism of a 1993 paper by Prasada & Pinker, which leads empirically into the experimental components of the thesis.

Chapter 4

Neighbourhoods

In this chapter it will be shown how the findings of certain empirical studies of inflectional morphology are undermined by unsound methodology, specifically the poor classification of stimuli according to their rhyming neighbourhood. The chapter will focus on two papers: Prasada & Pinker (1993), and a riposte by Lee (1995). Prasada & Pinker used novel verbs as stimuli in 3 behavioural studies to find results suggesting that regular verbs are rule-governed; Lee argued that their method of stimulus classification was unconvincing and, using an improved procedure, produced very different results, suggesting that regular verbs are actually inflected by analogy rather than by symbolic rule. After the classificatory techniques of those two papers have been analysed and shown to be flawed in objective terms, a new quantitative classification is discussed and shown to be more methodologically valid.

4.0 Introduction

4.0.1 Prasada & Pinker (1993)

Theories of generative grammar, such as that of Chomsky & Halle (1968), argued that irregular patterns, like regulars, were encapsulated in mental rules. However, such rules fail to account for certain aspects of irregular generalisation (Bybee & Slobin, 1982) and prototypicality (Bybee & Moder, 1983), suggesting that irregular patterns are in fact partly productive but not necessarily rule-generated. Furthermore, behavioural and connectionist studies indicated that rules may not be necessary at all—even for regular patterns (Bybee & Moder, 1983; Rumelhart & McClelland, 1986). Prasada & Pinker (1993) felt that, because of radical alternative linguistic

theories such as those generated by connectionist modelling, the explanations of the regular/irregular distinction and the psychological status of rules had been rendered unclear, and that "subtler differences in the extent and kind of productivity must also be considered" (page 4). One such difference is the ability of theoretical models to generalise from known to novel stimuli, and Prasada & Pinker's experiments were motivated by a desire to generate human data on the generalisation of knowledge to novel items, and to compare such data with results generated by connectionist models. They therefore designed 3 experiments using novel verbs and 3 connectionist simulations in an assessment of the relative merits of rule-driven and associative systems for inflectional morphology.

In two of the three experiments, subjects were asked to rate the 'phonological goodness' of possible past tense forms of novel verbs; the third, a past tense production task, will be discussed here.

They chose four irregular and four regular inflectional clusters as a basis for their novel verb stimuli, defined according to their vowel sound. The irregular vowel-change paradigms used were /ɪ/↦/ʌ/, /iː/↦/ɛ/, /əʊ/↦/uː/ and /eə/↦/ɔː/ (*e.g.* respectively 'swing', 'bleed', 'blow' and 'wear'); the regular paradigms used were /ɪ/↦/ɪ/, /iː/↦/iː/, /əʊ/↦/əʊ/ and /eɪ/↦/eɪ/ (*e.g.* 'drip', 'beam', 'croak' and 'base'). Each of these paradigms relates to a group of verbs which are inflected in the same way. For instance, the irregular /ɪ/↦/ʌ/ paradigm contains the irregular verbs cling, fling, sling, slink, sting, string, swing and wring, while the /ɪ/↦/ɪ/ regular paradigm contains such verbs as drip, sip, flip, clip, zip, dip, grip, nip, rip, tip and whip (to pick

a single rime cluster), and also kill, skid, blink, and so on[1].

Novel verbs were then chosen which were classed as being 'prototypical', 'intermediate' or 'distant' pseudoirregulars (or pseudoregulars) in relation to each irregular (or regular) paradigm. Stimuli in the prototypical pseudoirregular class were chosen to rhyme with real verbs in each irregular paradigm; intermediate pseudoirregular items were created by changing *either* the initial *or* the final consonant cluster of each prototypical stimulus; and distant pseudoirregular items were derived by changing both the initial *and* final consonant clusters of the prototypical stimuli. For example, in the /ɪ/↦/ʌ/ irregular paradigm, 'spling' was a prototypical pseudoirregular (rhyming with cling, fling, sling, etc.); 'fring' was an intermediate pseudoirregular (changing the initial consonant cluster); and 'trisp' was a distant pseudoirregular (changing the initial and final consonant clusters).

Similarly, prototypical pseudoregular stimuli were selected to rhyme with real verbs in each regular cluster; intermediate pseudoregulars used onset-nucleus sequences which do not exist in English verbs *and* rimes which do not exist[2]; and distant pseudoregular verbs differed from intermediate ones in that their codas do not exist in any English verbs. To give an example from the /iː/↦/iː/ paradigm: 'greem' was used as a prototypical pseudoregular (rhyming with beam, deem, gleam,

[1]It should be noted that Prasada & Pinker chose the regular paradigms "to rhyme with many regular verbs, with the constraint that they contained the vowels used in the pseudoirregular items" (p.11). While this is true for six of the eight paradigms, the regular /eɪ/ paradigm (base, case, chase, face, etc.) clearly does not contain the same diphthong as the irregular /eə/↦/ɔː/ paradigm (wear, swear, tear, bear, etc.).

[2]Again one should note that this is not true for the /ɪ/ paradigm (as Prasada & Pinker state themselves) as it is "not possible to find initial consonant cluster–vowel sequences which do not exist for the items with the vowel [ɪ]; for these items, the final cluster is the only one that does not exist".

etc.); 'smeeb' was used as an intermediate pseudoregular (as there are no 'smee-' verbs in English, nor any '-eeb' verbs); and 'smeenth' was used as a distant pseudoregular (possessing the non-existent verb coda '-nth').

By structuring the stimuli in this way, the novel verbs were thought to fall along a 'gradient of similarity' to known items.

Using these stimuli, Prasada & Pinker performed a production task, in which 24 subjects were required to use novel verbs in a productive sentential context. Meanings were ascribed to the novel stimuli, and subjects were asked to fill in blanks in test sentences using the most appropriate form of each novel word. 60 novel verbs were used as stimuli, as well as 30 novel adjectives and 30 novel nouns.

Their two main findings were as follows:

1. They replicated Bybee & Moder's (1983) finding that the rate of irregularisation of a novel verb by an irregular paradigm is a function of the 'global similarity' between the stimulus and existing irregular verbs (*e.g.* 'spling'↦'splung' is more likely than 'nist'↦'nust');

2. They found that, conversely, the generalisation of a novel verb by the regular paradigm does *not* appear to depend on its similarity to known regular verbs.

They concluded that 'rule' theories cannot explain irregular generalisations and connectionist models cannot account for regular generalisations, and therefore that a hybrid, modular system provides the optimal theoretical solution. In this model, irregulars are phonologically 'tagged' when stored in memory, allowing for generali-

sation by analogy, while regulars are inflected by a default suffixation process which can operate on *any* verb regardless of its phonological composition.

However, it has been argued that various inconsistencies, first in the selection of the stimuli, and secondly (and more fundamentally) in the definition of the measure of distance from prototypical clusters, call the validity of their findings into question (Lee, 1995).

4.0.2 Lee (1995)

Lee (1995) argued that, by using rimes which violate phonological constraints on possible English verbs in their intermediate and distant pseudoregular categories (*e.g.* /pləʊmf/), Prasada & Pinker biased subjects' responses against analogy: analogy of an 'impossible' verb to known verbs is itself impossible. He further maintained that the use of "orthographically opaque" irregular verbs also made it "virtually impossible" for a subject to produce a non-suffixed form. For example, the use of 'cloe' (theoretically mapping to /klu:/) as a prototypical pseudoirregular stimulus (cf. 'blow–blew') is confounded by the orthographic differences between 'oe' and 'ow'. More importantly, the criteria for classification of novel verbs by their phonological 'distance' from known verbs were shown to be highly inconsistent: for instance, determination of pseudoirregular and pseudoregular verbs used different criteria (*e.g.* intermediate pseudoirregular verbs allow for perfect rhyme or no rhyme, while intermediate pseudoregular items do not allow for rhyme). Lee's criticisms were taken into account in his own replication of their production task. First, he ensured that

the nuclei used in the irregular and regular paradigms matched each other. Secondly, he applied the same classificational criteria to the pseudoirregular and pseudoregular verbs, namely:

1. that prototypical verbs would have perfect rhyme and perfect or near-perfect alliteration with existing verbs in the relevant paradigm[3];

2. intermediate verbs would have perfect rhyme and no alliteration; and

3. distant stimuli would have no rhyme and no alliteration.

Lee's results showed that Bybee & Moder's finding for irregular verbs was again replicated, suggesting that irregulars are inflected by a process of analogy. However, he also found the same to be true of regulars, since the process of regularisation appeared in his data to depend on the similarity of the stimuli to known regular verbs. On the basis of his data it was argued that—contrary to Prasada & Pinker's conclusions—regular verbs are indeed inflected by a process of analogy rather than rule.

However, it remains to be seen whether the similarity metrics used by Prasada & Pinker or Lee are in fact adequate. Is the generation of novel verbs according to their distance (in terms of rhyme and alliteration) from *one particular* inflectional paradigm a sufficient method for determining their 'irregularity' or 'regularity'? It will be seen from the next section that the answer to this question is No.

[3]At first glance one might assume that if a stimulus has the same rime and the same onset as words in an existing paradigm, it is the same word as an item in that paradigm. This is not the case, however. In the /ɪ/↦/ʌ/ paradigm, for instance, a stimulus might take the onset from 'string' and the rime from 'slink' to produce the novel verb 'strink'.

The prototypical–intermediate–distant (hereafter PID) systems of classification for pseudoregular and pseudoirregular verbs (as used by both Prasada & Pinker and Lee) can be shown by quantitative computational analysis to be flawed even in Lee's replication, in that the PID classifications only apply to the stimuli within their specific vowel-change paradigm, and not throughout phonological space. For example, in terms of the irregular paradigm /iː/↦/ɛ/ ('feed', 'read', etc.), Lee chooses 'cleap' as a distant pseudoirregular item, although 'cleap' in fact relates to a large cluster of real irregular verbs (/iːp/↦/ɛpt/, *e.g.* 'sleep', 'creep'). When examined in the wider context of *all* verbs, the stimuli are found to bear little relation to the neat compartmentalisation suggested by Prasada & Pinker's classification.

The stimuli used in Prasada & Pinker's experiments, and those used in Lee's, will be analysed in detail here, before a different verb classification system is proposed and applied in the experiments which follow in this thesis.

4.1 Analysis of Stimuli from Prasada & Pinker and Lee

It is possible to represent phonologically any phoneme in terms of a binary sequence; that is, a sequence of 0s or 1s, each of which reflects the presence or absence of a phonetic feature, such as whether a phoneme is voiced or unvoiced. A whole word can also be represented as a vector by concatenating several of these phoneme vectors into one long string. The vector of a word may be regarded as the co-ordinates for the location of that word in phonological space. By rendering all English verbs into binary in the same manner, one can build up an abstract hyperdimensional representation of English verbs in terms of their phonological distribution[4]. This allows for the mathematical definition of the 'distance' between any two verbs, by calculating the root Euclidean distance between the two points in phonological hyperspace. This measure is calculated as the square root of the number of differing phonemic features of the two verbs being compared. For instance, the verbs 'heed' and 'heat' are nearly phonologically identical in that only one phonemic feature distinguishes them: despite the fact that they both end in an alveolar plosive, the coda is voiced in /hiːd/ and unvoiced in /hiːt/. There would therefore be only one bit different between each of the two vector representations, and the root Euclidean distance between the two verbs is therefore $\sqrt{1} = 1$.

It is, of course, possible to represent a *novel* verb as a binary sequence in exactly

[4]In this hyperdimensional space, all English verbs—being sequenced into binary co-ordinates—are therefore represented in the corners of a hypercube.

the same way, enabling its position in phonological hyperspace to be analysed by calculating its 'nearest neighbour', looking for nearby groups of verbs which are all inflected according to the same paradigm or cluster, and so on. A 'nearest neighbour' (hereafter NN) analysis was performed on the stimuli used by Prasada & Pinker and Lee.

They had proposed that the novel verbs used in their experiments were separable along a similarity gradient, as described above; the purpose of these tests was to establish whether in fact the stimuli could be shown to be categorisable *objectively* into a PID classification.

4.1.1 Nearest Neighbour Analysis

In the first analysis, two measures were calculated for each stimulus: NN—*i.e.* the phonologically closest real verb to the stimulus; and the root Euclidean distance between the stimulus and its NN, calculated as described above. Calculations were made by reference to a binary-sequenced translation of the CELEX database using a program designed and written by Dr Ramin Nakisa of this department.

Results of this analysis are shown in Tables 4.1 and 4.2. For each stimulus (in column 3) the table shows the nearest neighbour (column 4), the root distance between the NN and the stimulus (column 5), and the mean distance per category (column 6).

Two points arise from these results which suggest that there are problems with some of these stimuli in terms of their nearest neighbours. First, three of Lee's stimuli

Class	Category	Stimulus	NN	Distance	Mean
Pseudo-Irregular	P	spling	spring	1.0000	1.32
		skring	spring	1.7321	
		sprink	prink	2.6458	
		cleed	clean	1.4142	
		preed	breed	1.0000	
		queed	queen	1.4142	
		cloe	glow	1.0000	
		froe	flow	1.0000	
		plare	blare	1.0000	
		quare	queer	1.0000	
	I	fring	fling	1.0000	1.24
		ning	ring	1.4142	
		frink	frank	1.4142	
		cleef	cleave	1.0000	
		preek	freak	1.4142	
		queef	quaff	1.7321	
		foa	sew	1.0000	
		voa	bow	1.4142	
		jare	jeer	1.0000	
		grare	glare	1.0000	
	D	trisp	drift	1.7321	1.44
		nist	nest	1.0000	
		plip	trip	1.4142	
		gleef	cleave/grieve	1.4142	
		keeb	keep	1.0000	
		meep	meet	1.0000	
		goav	goad	1.7321	
		joam	drone	2.2361	
		flape	freight	1.4142	
		blafe	place	1.4142	
Pseudo-Regular	P	plip	trip	1.4142	1.24
		glip	grip	1.0000	
		brip	drip	1.0000	
		greem	cream	1.0000	
		pleem	plume	1.4142	
		treem	dream	1.0000	
		gloke	cloak	1.0000	
		proke	croak	1.7321	
		slace	slave	1.4142	
		nace	race	1.4142	
	I	brilth	print	2.2361	1.63
		glinth	glint	1.0000	
		plimph	primp	1.7321	
		smeeb	speed	2.0000	
		smeeg	sneak	1.4142	
		smeej	smudge	1.7321	
		ploab	probe	1.0000	
		ploag	plague	1.7321	
		smaib	spade	2.0000	
		smaig	snake	1.4142	
	D	frilg	frisk	2.2361	2.30
		krilg	catalogue	2.4495	
		smeelth	smelt	2.0000	
		smeenth	cement	2.4495	
		smeerg	snake	2.2361	
		ploamph	flounce/trounce	2.4495	
		ploanth	flounce/trounce	2.4495	
		smairg	snake	2.0000	
		smairph	space	2.4495	

Table 4.1: NN analysis of Prasada & Pinker's stimuli

Class	Category	Stimulus	NN	Distance	Mean
Pseudo-Irregular	P	gling	cling	1.0000	1.03
		strink	drink	2.8284	
		veed	feed	1.0000	
		pread	breed	1.0000	
		krow	crow	0.0000	
		plow	blow	1.0000	
		zear	dare	1.4142	
		pear	pair	0.0000	
	I	bling	bring	1.0000	1.21
		pling	bring	1.4142	
		geed	guard	1.4142	
		cleed	clean	1.4142	
		frow	flow	1.0000	
		shrow	shallow	1.0000	
		jear	jeer	1.0000	
		mear	may	1.4142	
	D	blik	black	1.4142	1.62
		plig	prick	1.4142	
		geem	keen	1.4142	
		cleep	creep	1.0000	
		frowl	troll	1.7321	
		shrowth	shroud	2.0000	
		jearp	jeer	2.2361	
		mearp	mate	1.7321	
Pseudo-Regular	P	frip	flip	1.0000	1.05
		krip	grip	1.0000	
		keam	keen	1.0000	
		geem	keen	1.4142	
		boke	poke	1.0000	
		zoke	soak	1.0000	
		gare	care	1.0000	
		vare	fare	1.0000	
	I	yip	whip	1.4142	1.08
		jip	gyp	0.0000	
		feam	seem	1.0000	
		leam	lean	1.0000	
		doke	poke	1.4142	
		voke	soak	1.4142	
		yare	wear	1.4142	
		nare	near	1.0000	
	D	yid	wed	1.7321	1.43
		jiz	jazz	1.4142	
		feap	seep	1.0000	
		leat	lead	1.0000	
		dofe	dose	1.0000	
		vome	foam	1.0000	
		yairth	wait	2.2361	
		nairth	note	2.0000	

Table 4.2: NN analysis of Lee's stimuli

have a NN Euclidean distance value of 0.0000, in other words are phonologically identical to real verbs ('krow', 'pear' and 'jip', cf. 'crow', 'pair' and 'gyp')[5]. Note also that the first two of these are described as "prototypical pseudoirregular" stimuli, although the real verbs 'crow' and 'pair' are both regular!

Secondly, if one calculates the mean root Euclidean distance between stimuli and their NNs in each category, one sees that Prasada & Pinker's stimuli range in mean distance per category from 1.24 to 2.30; Lee's range between 1.03 and 1.62. 5 of Prasada & Pinker's categories lie between 1.24 and 1.63, the obvious outlier being the Distant Pseudoregular category (2.30). In other words, the stimuli in this category are not merely further away from regular verbs than the other categories but are further away from *all* verbs. One would therefore expect those verbs to be inflected by the regular paradigm rather than the irregular one, whether by analogy or by the application of a default regular rule; so it is difficult to appreciate the validity of defining these stimuli as *distant* pseudoregulars.

It should be pointed out that the NN classification program used to calculate the NNs and distances shown above does not use the same definition of proximity as a human, since it gives equal weight to the onset and the rime of a verb. Anecdotal evidence suggests that people put more emphasis on rime than on onset when comparing two words and would therefore think, for instance, that the novel word 'mear' (pronounced /mɛə*/ to rhyme with 'wear') is closer to the word 'wear' than to the

[5]Lee acknowledges this obvious deficiency but argues "However, subjects have been briefed to treat all the words posed in the task as non-existent". This is unsatisfactory.

word 'may'. The NN classifier would think the opposite.

4.1.2 Partword Analysis

The specific analyses carried out below (which examine *all* verbs with the same rime as each stimulus) provide strong evidence that the PID categories designed by Prasada & Pinker do not—objectively—distinguish between the stimuli as they maintained.

Tables 4.3 and 4.4 show the results of a second (partword) analysis whereby, for each stimulus, four measures were taken:

1. the number of English verbs sharing the onset of the stimulus;
2. the number of these (shared-onset) verbs which are irregular;
3. the number of English verbs sharing the rime of the stimulus; and
4. the number of these (shared-rime) verbs which are irregular.

The data for the onsets yields little information of value, as one might expect; however, the rime data are much more revealing.

Looking at the pseudoirregular stimuli in Table 4.3 in terms of rime it is clear that 7 of the prototypical stimuli are identical to 7 of the intermediate stimuli. On the basis of rime only, the stimuli are therefore inseparable into these two categories[6]. Within the intermediate category, the verbs 'cleef' and 'queef' rhyme with 5 regulars

[6]It should be remembered here that Prasada & Pinker defined the difference between these two categories in terms of onset rather than rime.

Class	Category	Stimulus	N(Onset)	(Irreg)	N(Rime)	(Irreg)
Pseudo-Irregular	P	spling	10	1	17	13
		skring	23	0	17	13
		sprink	10	2	20	9
		cleed	54	2	35	14
		preed	139	2	35	14
		queed	31	0	35	14
		cloe	54	2	68	12
		froe	38	1	68	12
		plare	37	0	35	6
		quare	31	0	35	6
	I	fring	38	1	17	13
		ning	75	3	17	13
		frink	38	1	20	9
		cleef	54	2	5	0
		preek	139	2	18	3
		queef	31	0	5	0
		foa	174	16	68	12
		voa	69	0	68	12
		jare	72	0	35	6
		grare	44	2	35	6
	D	trisp	77	1	2	0
		nist	75	3	22	0
		plip	37	0	33	0
		gleef	22	0	5	0
		keeb	384	7	0	-
		meep	250	14	16	8
		goav	80	4	2	0
		joam	72	0	6	0
		flape	60	5	11	0
		blafe	52	2	2	0
Pseudo-Regular	P	plip	37	0	33	0
		glip	22	0	33	0
		brip	48	6	33	0
		greem	44	2	17	1
		pleem	37	0	17	1
		treem	77	1	17	1
		gloke	22	0	17	0
		proke	139	2	17	0
		slace	48	6	27	0
		nace	75	3	27	0
	I	brilth	48	6	0	-
		glinth	22	0	0	-
		plimph	37	0	0	-
		smeeb	18	2	0	-
		smeeg	18	2	3	0
		smeej	18	2	1	0
		ploab	37	0	3	0
		ploag	37	0	1	0
		smaib	18	2	0	-
		smaig	18	2	2	0
	D	frilg	38	1	0	-
		krilg	63	2	0	-
		smeelth	18	2	0	-
		smeenth	18	2	0	-
		smeerg	18	2	0	-
		ploamph	37	0	0	-
		ploanth	37	0	0	-
		smairg	18	2	0	-
		smairph	18	2	0	-

Table 4.3: Partword analysis of Prasada & Pinker's stimuli

Class	Category	Stimulus	N(Onset)	(Irreg)	N(Rime)	(Irreg)
Pseudo-Irregular	P	gling	22	0	17	13
		strink	34	4	20	9
		veed	69	0	35	14
		pread	139	2	35	14
		krow	63	2	68	12
		plow	37	1	68	12
		zear	6	0	34	6
		pear	258	2	34	6
	I	bling	52	2	17	13
		pling	37	0	17	13
		geed	80	4	35	14
		cleed	54	2	35	14
		frow	38	1	68	12
		shrow	10	2	68	12
		jear	72	0	34	6
		mear	250	14	34	6
	D	blik	52	2	24	1
		plig	37	0	7	1
		geem	80	4	17	1
		cleep	54	2	16	8
		frowl	38	1	24	0
		shrowth	10	2	0	-
		jearp	72	0	0	-
		mearp	250	14	0	-
Pseudo-Regular	P	frip	38	1	33	0
		krip	63	2	33	0
		keam	384	7	17	1
		geem	80	4	17	1
		boke	262	27	17	0
		zoke	6	0	17	0
		gare	80	4	34	6
		vare	69	0	34	6
	I	yip	27	0	33	0
		jip	72	0	33	0
		feam	174	16	17	1
		leam	153	10	17	1
		doke	445	4	17	0
		voke	69	0	17	0
		yare	27	0	34	6
		nare	75	3	34	6
	D	yid	27	0	8	7
		jiz	72	0	4	0
		feap	174	16	16	8
		leat	153	10	31	4
		dofe	445	4	1	0
		vome	69	0	6	0
		yairth	27	0	0	-
		nairth	75	3	0	-

Table 4.4: Partword analysis of Lee's stimuli

and no irregulars; these stimuli would perhaps be better defined as distant pseudoirregulars. 'Smeeg' and 'ploab' are scarcely quantatively different from 'cleef' and 'queef', although they find themselves in the intermediate pseudoregular category.

In the distant pseudoirregular category one finds that all bar one of the stimuli rhyme with no irregular verbs, as one might hope; but the exception is 'meep'; as a stimulus with 8 of its 16 rhyming verbs being irregular, a far higher proportion than most verbs, one might expect it to be classified as a prototypical rather than distant pseudoirregular.

There is clearly little or no difference between distant pseudoirregulars and prototypical pseudoregulars, nor again between the intermediate and distant pseudoregular categories. The stimuli (with one or two exceptions) seem to divide into 3 categories rather than 6: of the six categories in the Prasada & Pinker partword analysis (Table 4.3 above), the top two contain stimuli which rhyme with many real verbs, of which some are irregular; the next two contain stimuli which rhyme with regular but not irregular verbs; and the stimuli in the last two categories rhyme with almost no regulars and absolutely no irregulars. As was suggested by the NN analysis, the inclusion of stimuli which rhyme with *no* verbs seems strange, since one would expect them to be regularised, even though they are classified as *distant* pseudoregulars.

Looking at Lee's stimuli (Table 4.4) we see a similar pairing of categories. In terms of rime alone, there is obviously little difference—by definition—between prototypical and intermediate stimuli, whether within the pseudoirregular or pseudoregular paradigm, in accordance with Lee's classification (see above). Furthermore, there

is no coherent similarity within the distant pseudoirregular category—this includes 3 stimuli which do not rhyme with any verb; 3 which rhyme with some (including an irregular); one which rhymes with some regulars but no irregulars; and finally the novel word 'cleep': of the 16 verbs sharing its phonological rime, 8 are irregular, far above the median value of 1 (see below), and rhyming with as many irregular verbs as regular ones, a highly-qualified contender for inclusion in the prototypical pseudoirregular category.

Similarly, for the pseudoregular stimuli, the prototypical and intermediate categories are identical in terms of rime. It is curious to note that 'gare' and 'vare' are prototypical pseudoregulars when 'zear' and 'pear' are prototypical pseudo*ir*regulars, and 'yare' and 'nare' are intermediate pseudoregulars. As the only phonological difference between any of these stimuli is their onset, one can only assume that onset is being treated as having greater importance in verbal inflection than the literature would suggest.

Furthermore, it is difficult to see why, for instance, 'vome' is classified as a distant pseudoregular when all 6 shared-rime verbs are regular (as—for that matter—are 69 out of 69 shared-onset verbs). This stimulus would seem to be more appropriately classified as a prototypical pseudoregular (or perhaps distant pseudoirregular) stimulus.

The principal problem with the selection of the stimuli used by Prasada & Pinker and Lee seems to be that their typicality was calculated with reference to one specific inflectional paradigm only. As a result of this, similar stimuli are placed in very

different categories. It would be better if a stimulus were categorised, not by virtue of being near to or far away from a particular group of verbs, but in terms of its relation to *all* verbs.

The remainder of this chapter describes precisely that: a new classificatory system is defined, in an attempt to classify stimuli into certain categories by comparison with all English verbs rather than by comparison with only a particular subset of verbs.

4.2 Glushko (1979) and Marchman (1997)

Glushko (1979) described a system of 'neighbourhood structure' whereby all verbs belong in 'neighbourhoods' defined according to the rime of their stem: all verbs which have a particular rime are in the same neighbourhood, and all verbs in a particular neighbourhood have the same rime[7]. Within the neighbourhood of any specific 'target' verb, all other members may be defined as 'friends' and 'enemies' according to their past tense: friends are verbs which are inflected according to the same paradigm as the target verb; enemies are verbs which are inflected by a different paradigm. For an irregular verb, therefore, enemies may be subdivided into irregular and regular groups. For example, in the case of the irregular verb 'blow': 'throw' would be a friend; 'go' would be an irregular enemy; and 'flow' would be a regular enemy. For a regular target verb, there are only two categories: *e.g.* for 'flow', other /əʊ/ ↦ /əʊd/ verbs are friends (all regular), and all verbs with a different inflectional pattern (that is, all irregulars) are enemies.

[7]Note that these are not the same thing.

Marchman (1997) performed an elicited production task on children, in order to use this concept of neighbourhood structure to provide data which would contribute to the dual-route/single-route debate in inflectional morphology research. She argued (on page 287):

> ...both single- and dual-mechanism models suppose that frequency impacts error rates and that phonological features and neighborhood factors influence the production of irregularization errors like zero-marking. However, dual-mechanism models predict that regularization errors should occur independently of neighborhood similarity, whereas a single-mechanism view proposes that similar mechanisms underlie the production of both regularization and irregularization errors. Crucially, irregular verbs that are similar to suffixed verbs should be more vulnerable to regularizations than those that are not.

The task involved eliciting past tense forms of specific verbs from subjects, using a picture-book task (modified from Bybee & Slobin, 1982). 49 stimuli were used (11 regular, 38 irregular), and were coded according to the following criteria:

- as high or low frequency in their past tense form;
- according to whether they contained a stem-final alveolar;
- in terms of their past tense neighbourhood structure, specifically:
  - as 'High' or 'Low' in terms of Total Friends;
  - as 'High' or 'Low' in terms of Total Enemies;
  - as 'High' or 'Low' in terms of Suffixed Enemies (irregulars only).

The terms 'High' and 'Low' were defined by comparison with medians for a lexicon of 1,191 monosyllabic English verbs.

Irregular verbs were also coded for vowel dominance in their stem form (N=18), past tense form (N=7) or neither (N=13).

Childrens' responses to the stimuli were coded as 'correct', 'error' or 'nonvalid', where nonvalid answers included all responses which could not unambiguously be classed as "legitimate attempts to produce verb forms in the simple past tense". Errors were classed as 'suffixation', 'zero-marking' or 'vowel change'.

Marchman found that errors accounted for approximately 20% of all valid responses, and that most resulted from adding suffixes to irregular verbs. Her findings for irregulars and regulars will be discussed in turn.

She found that irregulars with a stem-final alveolar plosive were equally likely to be overregularised (*i.e.* suffixed) as others; items with high-frequency past tense forms were less likely to be suffixed than those with low-frequency past tenses; and stimuli with stem-dominant vowels were more likely to be suffixed than those with past-dominant vowels (consistent with Stemberger, 1993).

In terms of neighbourhood structure, items with many suffixed neighbours were more likely to be regularised than those with few suffixed neighbours, although no difference was found between irregular verbs with many friends compared to those with few friends.

Marchman concluded that 'enemy effects' result from an increased opportunity for interference among phonologically similar items. She argued that three factors lead to overregularisation:

1. low past tense frequency;

2. a dominant vowel in the verb stem; and

3. a high number of suffixed enemies.

She maintained that the presence of two factors leads to a significantly higher probability of suffixation than if one or no factors are present, and that the presence of all three factors leads to an even higher probability of overregularisation.

With regard to the 11 regular stimuli, Marchman found that more errors were made on verbs with a low-frequency past tense form than those with a high-frequency one. It was also found that more errors were made on verbs with few friends than on those with an above-median number of friends. Furthermore, having a high number of friends was found to facilitate correct inflection of both high and low frequency stimuli. However, only 11 stimuli were used in this analysis, and no statistical validation of these trends was provided.

Finally, it was found that zero-marking errors were more likely for stem-final alveolar verbs than for others, independent of verb class. However, Marchman argued that children were found to be using the zero-marked form productively, and not simply because they were "not trying to convert the verb to a past tense form at all" (Xu & Pinker, 1995).

Marchman concluded that phonological information influences production errors independently of verb class, thus supporting her earlier findings (for which see Marchman, Wulfeck & Ellis Weismer, 1999), and that her results run counter to the dual-

route theory of inflectional morphology. She argued that the results were consistent with those found for adult English speakers (Marchman & Callan, 1995), with those found for children when a different stimulus set was used (Marchman et al., 1999), and with the behaviour of connectionist systems (*e.g.* Plunkett & Marchman, 1993). Finally, she argued that:

> ... past tense findings are consistent with the conclusion that the mechanisms underlying the productive use of regular and irregular patterns are actually more similar than they are different. (Page 300)

The next section describes a classification of neighbourhood structure based on the ideas of Glushko and Marchman[8] which attempts to quantify in detail the relations between irregular and regular verbs in terms of all phonological rime neighbourhoods.

4.3 A New Classification of Neighbourhood Density

By an analysis of the values of various subsets of neighbourhoods and the median values across all irregular and regular clusters, it is possible to develop a more stringent system of classification particularly for novel verbs in terms of the potential for inflectional mapping by analogy in a regular or irregular direction.

The terms 'neighbourhood' and 'cluster' are used in very specific ways in this classification: as stated above, a neighbourhood is defined according to the phonological rime of the stem; within a neighbourhood, clusters are defined in terms of the phonological rime of the inflected verb. For example, the verbs 'throw' and 'go' are

[8]And others: for example Jared, McRae & Seidenberg (1990).

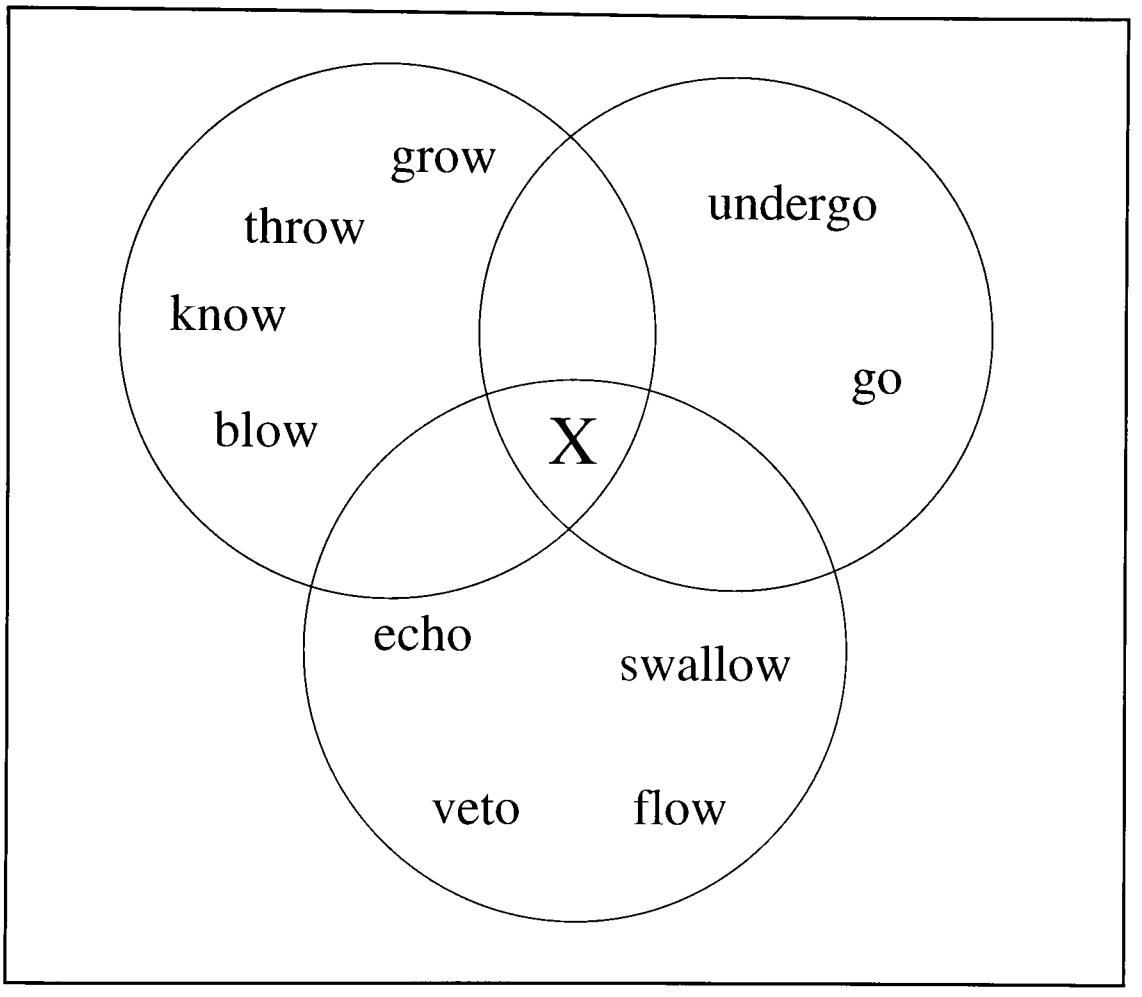


Figure 4.1: Venn diagram with samples of each cluster in the /əʊ/ neighbourhood

members of the same neighbourhood (*i.e.* the /əʊ/ neighbourhood) but within that, each of them belongs to a different cluster: 'throw' belongs to the /əʊ/↦/uː/ cluster, while 'go' is part of (and, with its compounds, comprises) the /əʊ/↦/ɛnt/ cluster. This is depicted visually in Figure 4.1.

In the Venn diagram shown here, the rectangular border represents the exclusive set of verbs which have the same rime, in this case the /əʊ/ neighbourhood. Within this environment, the circles represent clusters of these verbs which share inflectional

patterns. In this case, it is clear that there are three inflectional possibilities for an /əʊ/ verb: mapping to /uː/, /ɛnt/ or /əʊd/. A non-exhaustive set of exemplars is depicted in each circle in the figure. It is generally true that clusters will have no overlapping members, since most verbs can only be inflected in one way. However, in the /æŋ/ neighbourhood, for example, a verb such as 'hang' would rightly belong in the overlapping area between the /æŋ/↦/ʌŋ/ and /æŋ/↦/æŋd/ clusters, since it is both a regular and an irregular verb, depending on context. All the clusters are shown as overlapping here to allow for such possibilities.

All novel verbs in the neighbourhood might be considered to belong in one of two places: in the area of shared space created by the overlapping clusters (depicted in Figure 4.1 by an X) if it is inflected by a known pattern or, if inflected by a totally different pattern, in the space outside all of the clusters within the environment.

In the case of target novel verbs, it is of course inappropriate to use the terminology of friends and enemies, since there is no *correct* inflection of a novel verb. However, one can talk of—and, importantly, one can quantify—the irregular and regular 'neighbours' of a novel verb. So, for instance, the novel verb 'cleed' (and indeed all novel verbs with the rime /iːd/) has 21 regular neighbours (*e.g.* 'weed', 'succeed'). But how many irregular neighbours does it have? This is contentious as one must choose whether to include or exclude compounds from the calculation: for instance, 'breed' and 'read' are certainly irregular neighbours, but should one also include 'crossbreed', 'interbreed', 'proofread', 'misread', and so on? If one does, then the presence of a single irregular with many compounds could have a large effect on

the classification of a verb.

In the following classification of neighbourhood density, both figures are given: the number of simple irregulars in a neighbourhood, and the total (*i.e.* simple and compound) irregulars in the neighbourhood[9].

Two lists of British English verbs with irregular past tense forms were created. The list used by Pinker (1995) was taken as a starting point and altered in 3 ways: items irregular in American but not British English were removed (*e.g.* 'sneak'); items with irregular past participles but regular past tenses were removed (*e.g.* 'sow'); and items which had been omitted from the list were added (N=7, *e.g.* 'bet'). Furthermore, there was no consistency in whether compounds were included in the original list in that, for example, 15 irregular compounds such as 'foretell' and 'forswear' were included, but 103 (*e.g.* 'oversleep') were omitted. Two lists were created: one including simple irregulars only (list α), the other simple and compound irregulars (list β). These lists may be found in appendices A and B.

A table was then drawn up showing the cluster structure of all the verbs in both lists. In other words, all the verbs were tagged according to their phonological stem rime and phonological inflected rime (*e.g.* 'string' is an /ɪŋ/↦/ʌŋ/ verb). There are 60 irregular neighbourhoods in English, and several of these neighbourhoods contain more than one irregular cluster, yielding a total of 87 clusters. For each of these 87 stem-and-inflection clusters, 5 measurements were taken (from within the same rime

[9]It should be noted that Prasada & Pinker claim that all irregulars are either monosyllables or compounds of monosyllabic irregulars; this is not quite true, since irregular verbs such as 'beseech' and 'bereave' are neither monosyllabic nor compounds of monosyllabic irregulars.

neighbourhood):

$\mathsf{Is}\alpha$ Number of irregulars in the same cluster using list α;

$\mathsf{Id}\alpha$ Number of irregulars in different clusters using list α;

$\mathsf{Is}\beta$ Number of irregulars in the same cluster using list β;

$\mathsf{Id}\beta$ Number of irregulars in different clusters using list β;

R Number of regular enemies.

Note that I and R represent 'irregular' and 'regular' neighbours, and s and d represent clusters of 'same' and 'different' inflections.

These measurements were made by counting and cross-referencing all verbs in the CELEX database (Baayen, Piepenbrock & Gulikers, 1995). The complete data are shown in Table 4.5.

Medians were calculated across all British English irregulars for the 5 measurements described above, and stem-inflection clusters could then be classified as 'High' or 'Low' Is; 'High' or 'Low' Id and 'High' or 'Low' R in relation to median values across all irregular verb clusters. Clusters were defined as High Is if $\mathsf{Is}\alpha > \mu(\mathsf{Is}\alpha)$ and $\mathsf{Is}\beta \geq \mu(\mathsf{Is}\beta)$, otherwise as Low Is[10]. Similarly, High Id clusters were defined by $\mathsf{Id}\alpha > \mu(\mathsf{Id}\alpha)$ and $\mathsf{Id}\beta > \mu(\mathsf{Id}\beta)$, and High R clusters by $\mathsf{R} > \mu(\mathsf{R})$.

The median values across all irregular clusters were: $\mu(\mathsf{Is}\alpha) = 1$; $\mu(\mathsf{Id}\alpha) = 1$; $\mu(\mathsf{Is}\beta) = 3$; $\mu(\mathsf{Id}\beta) = 1$; and $\mu(\mathsf{R}) = 17$.

[10]where μ is the median across *all* English irregular clusters.

Cluster	lsα	ldα	lsβ	ldβ	R	Cluster	lsα	ldα	lsβ	ldβ	R
/ɑːst/↦/ɑːst/	1	0	8	0	8	/ɪd/↦/æd/	1	2	2	5	2
/ɔː/↦/uː/	1	0	3	0	17	/ɪg/↦/ʌg/	1	0	1	0	6
/ɔːl/↦/ɛl/	1	0	2	0	29	/ɪk/↦/ʌk/	1	0	1	0	23
/eɪ/↦/ɛd/	1	1	3	1	67	/ɪl/↦/ɪlt/	1	0	1	0	18
/eɪ/↦/uː/	1	1	1	3	67	/ɪld/↦/ɪlt/	1	0	2	0	1
/eɪk/↦/eɪd/	1	5	2	11	13	/ɪm/↦/æm/	1	0	1	0	9
/eɪk/↦/əʊk/	2	4	3	10	13	/ɪn/↦/ʌn/	2	1	2	1	26
/eɪk/↦/ʊk/	3	3	8	5	13	/ɪn/↦/æn/	1	2	1	2	26
/aɪ/↦/ɔːt/	1	2	1	5	122	/ɪt/↦/ɪt/	4	2	5	2	89
/aɪ/↦/eɪ/	1	2	3	3	122	/ɪt/↦/æt/	2	4	2	5	89
/aɪ/↦/uː/	1	2	2	4	122	/ɪv/↦/eɪv/	1	0	3	0	5
/aɪd/↦/əʊd/	3	3	6	4	17	/ɒst/↦/ɒst/	1	0	1	0	5
/aɪd/↦/ɪd/	3	3	4	6	17	/ʊt/↦/ʊt/	1	0	1	0	4
/aɪk/↦/ʌk/	1	0	1	0	8	/ʌm/↦/eɪm/	1	0	3	0	13
/aɪn/↦/ɒn/	1	0	2	0	42	/ʌn/↦/æn/	1	0	4	0	5
/aɪnd/↦/aʊnd/	4	0	8	0	4	/ʌst/↦/ʌst/	1	0	1	0	11
/aɪt/↦/ɔːt/	1	4	2	9	30	/ʌt/↦/ʌt/	2	0	3	0	10
/aɪt/↦/əʊt/	2	3	5	6	30	/iː/↦/ɔː/	1	2	4	2	19
/aɪt/↦/ɪt/	2	3	4	7	30	/iː/↦/ɛd/	1	2	1	5	19
/aɪv/↦/əʊv/	5	0	5	0	10	/iː/↦/ɒz/	1	2	1	5	19
/aɪz/↦/əʊz/	1	0	2	0	277	/iːʧ/↦/ɔːt/	2	0	2	0	11
/ɜːn/↦/ɜːnt/	2	0	3	0	13	/iːd/↦/ɛd/	6	0	14	0	21
/ɜːst/↦/ɜːst/	1	0	1	0	2	/iːk/↦/ɔːt/	1	1	1	2	15
/ɜːt/↦/ɜːt/	1	0	1	0	21	/iːk/↦/əʊk/	1	1	2	1	15
/ɔɪl/↦/ɔɪlt/	1	0	1	0	14	/iːl/↦/əʊl/	1	3	1	4	18
/əʊ/↦/ɛnt/	1	4	3	8	56	/iːl/↦/ɛlt/	3	1	4	1	18
/əʊ/↦/uː/	4	1	8	3	56	/iːm/↦/ɛmt/	1	0	1	0	16
/əʊld/↦/ɛld/	1	0	4	0	9	/iːn/↦/ɛnt/	2	0	2	0	20
/ɪə/↦/ɜːd/	1	0	4	0	40	/iːp/↦/ɛpt/	6	0	8	0	7
/ɛə/↦/ɔː/	4	0	6	0	27	/iːt/↦/eɪt/	1	2	1	3	27
/ɛd/↦/ɛd/	3	1	3	1	12	/iːt/↦/ɛt/	1	2	1	3	27
/ɛd/↦/ɒd/	1	3	1	3	12	/iːt/↦/iːt/	1	2	2	2	27
/ɛl/↦/əʊld/	2	3	6	5	17	/iːv/↦/əʊv/	3	2	4	2	20
/ɛl/↦/ɛlt/	3	2	5	6	17	/iːv/↦/ɛft/	2	3	2	4	20
/ɛnd/↦/ɛnt/	5	0	7	0	36	/iːz/↦/əʊz/	1	0	1	0	10
/ɛt/↦/ɛt/	4	1	12	3	18	/uː/↦/ɪd/	1	0	5	0	60
/ɛt/↦/ɒt/	1	4	3	12	18	/uːt/↦/ɒt/	1	0	3	0	41
/ɪŋ/↦/ɔːt/	1	10	1	12	4	/uːz/↦/əʊz/	1	1	1	1	27
/ɪŋ/↦/ʌŋ/	7	4	9	4	4	/uːz/↦/ɒst/	1	1	1	1	27
/ɪŋ/↦/æŋ/	3	8	3	10	4	/æʧ/↦/ɔːt/	1	0	1	0	12
/ɪŋk/↦/ɔːt/	1	5	3	6	11	/æŋ/↦/ʌŋ/	1	0	2	0	9
/ɪŋk/↦/ʌŋk/	1	5	1	8	11	/ænd/↦/ʊd/	1	0	4	0	10
/ɪŋk/↦/æŋk/	4	2	5	4	11	/æv/↦/æd/	1	0	1	0	0
/ɪd/↦/ɪd/	2	1	5	2	2						

Table 4.5: Neighbourhood density analysis of English irregular verbs

The pseudoirregular clusters can therefore be divided into 4 categories according to 2 variables. These variables are Is+Id, *i.e.* the total number of irregulars with the same rime as the stimulus; and R, *i.e.* the number of regular verbs with the same rime as the stimulus. Any verb or novel verb can therefore be categorised—by reference to the median values across all irregular clusters—as HighI or LowI and as HighR or LowR. An attempt has been made to stabilise Id at 0 as far as possible, so that analogy with an irregular verb can only lead to one inflection-type for each stimulus: this also minimises the possibility of competition effects between Is and Id within a cluster.

Thus all irregular verbs can be placed into one of the following four categories: HighILowR; HighIHighR; LowILowR; and LowIHighR.

The same process was carried out for regular verbs: a list was created of all British English regular verbs contained in the CELEX database. These verbs were used to produce a categorisation of the neighbourhood structure of all regulars. For each stem neighbourhood, 3 measurements were taken (from within that neighbourhood):

Iα Number of (irregular) enemies using list α;

Iβ Number of (irregular) enemies using list β;

R Number of (regular) friends.

Medians were calculated exactly as for irregular verbs, and regular rimes defined as 'High' or 'Low' Iα, 'High' or 'Low' Iβ and 'High' or 'Low' R in the same manner. It should be noted here that μ(Iα), μ(Iβ) and μ(R) are different from their values in the

irregular paradigm since the former are calculated across all neighbourhoods in which there is at least one regular verb (N=388), as opposed to across all clusters in which there is an irregular verb (N=87). The medians for all regular neighbourhoods were: $\mu(\mathsf{I}\alpha) = 0$; $\mu(\mathsf{I}\beta) = 0$; and $\mu(\mathsf{R}) = 6$. Regulars were defined as HighI if $\mathsf{I}\alpha > \mu(\mathsf{I}\alpha$ and $\mathsf{I}\beta > \mu(\mathsf{I}\beta$, otherwise as LowI; and as HighR if $\mathsf{R} > \mu(\mathsf{R})$. A complete list of the neighbourhood structure of all neighbourhoods containing at least one regular verb is shown in Tables 4.6 and 4.7.

4.3.1 Re-Analysis of Prasada & Pinker's and Lee's Stimuli

In the light of this neighbourhood classification, Prasada & Pinker's stimuli—and those of Lee—were analysed again, by counting the number of irregulars and regulars in the same (rhyming) neighbourhood as each stimulus. The results are given in Tables 4.8 and 4.9.

For each stimulus (in column 3) 3 figures are given: the number of simple irregular neighbours ($\mathsf{I}\alpha$); the number of compound irregular neighbours ($\mathsf{I}\beta$); and the number of regular neighbours (R).

From the data on Prasada & Pinker's stimuli, it is clear that some stimuli in the intermediate and distant pseudoirregular categories are misplaced when viewed in terms of the new classification. For example, 'cleef' and 'queef' (intermediate pseudoirregulars) look as though they belong in the distant category whilst, conversely, 'meep' (a distant pseudoirregular) seems to belong to the prototypical category. Again, many of the stimuli in the prototypical and intermediate pseudoirregular categories share

Rime	$l\alpha$	$l\beta$	R	Rime	$l\alpha$	$l\beta$	R	Rime	$l\alpha$	$l\beta$	R	Rime	$l\alpha$	$l\beta$	R
/aː/	0	0	2	/eɪʤ/	0	0	15	/ɔɪ/	0	0	12	/ɪəs/	0	0	1
/aːʧ/	0	0	8	/eɪd/	0	0	36	/ɔɪŋk/	0	0	1	/ɛ*/	4	6	27
/aː*/	0	0	11	/eɪf/	0	0	2	/ɔɪd/	0	0	2	/ʊə*/	0	0	18
/aːθ/	0	0	1	/eɪg/	0	0	2	/ɔɪl/	1	1	14	/ʊəl/	0	0	2
/aːʒ/	0	0	4	/eɪk/	6	13	13	/ɔɪn/	0	0	10	/ʊəm/	0	0	1
/aːʤ/	0	0	9	/eɪl/	0	0	36	/ɔɪnt/	0	0	9	/ʊəns/	0	0	1
/aːb/	0	0	1	/eɪm/	0	0	22	/ɔɪs/	0	0	3	/ə/	0	0	3
/aːd/	0	0	14	/eɪn/	0	0	52	/ɔɪst/	0	0	1	/ə*/	0	0	325
/aːf/	0	0	12	/eɪnʤ/	0	0	11	/ɔɪt/	0	0	1	/əd/	0	0	7
/aːft/	0	0	7	/eɪnt/	0	0	4	/ɔɪz/	0	0	2	/ək/	0	0	2
/aːk/	0	0	16	/eɪp/	0	0	11	/əʊ/	5	12	53	/əks/	0	0	1
/aːl/	0	0	2	/eɪps/	0	0	1	/əʊð/	0	0	3	/əl/	0	0	16
/aːm/	0	0	13	/eɪs/	0	0	27	/əʊʧ/	0	0	6	/əld/	0	0	1
/aːn/	0	0	2	/eɪst/	0	0	5	/əʊb/	0	0	3	/əm/	0	0	8
/aːnʧ/	0	0	3	/eɪt/	0	0	570	/əʊd/	0	0	18	/əmf/	0	0	1
/aːnd/	0	0	5	/eɪv/	0	0	15	/əʊf/	0	0	1	/ən/	0	0	40
/aːns/	0	0	10	/eɪz/	0	0	18	/əʊg/	0	0	1	/ənd/	0	0	3
/aːnt/	0	0	9	/aɪ/	3	6	121	/əʊk/	0	0	17	/ənz/	0	0	17
/aːp/	0	0	2	/aɪð/	0	0	2	/əʊks/	0	0	2	/ənt/	0	0	6
/aːs/	0	0	6	/aɪʤ/	0	0	2	/əʊl/	0	0	23	/ənz/	0	0	1
/aːsk/	0	0	5	/aɪb/	0	0	12	/əʊld/	1	4	9	/əp/	0	0	9
/aːsp/	0	0	4	/aɪd/	6	10	17	/əʊlt/	0	0	4	/əs/	0	0	16
/aːst/	1	8	8	/aɪf/	0	0	1	/əʊm/	0	0	6	/əst/	0	0	3
/aːt/	0	0	10	/aɪk/	1	1	8	/əʊn/	0	0	23	/ət/	0	0	11
/aːv/	0	0	3	/aɪl/	0	0	17	/əʊp/	0	0	13	/ɛʧ/	0	0	6
/aɪz/	0	0	1	/aɪm/	0	0	11	/əʊs/	0	0	3	/ɛʃ/	0	0	4
/ɔː/	1	3	17	/aɪn/	1	2	42	/əʊst/	0	0	8	/ɛʤ/	0	0	8
/ɔːʧ/	0	0	2	/aɪnd/	4	8	4	/əʊt/	0	0	21	/ɛb/	0	0	2
/ɔː*/	0	0	27	/aɪp/	0	0	8	/əʊv/	0	0	2	/ɛd/	4	4	12
/ɔːʤ/	0	0	3	/aɪs/	0	0	9	/əʊz/	0	0	31	/ɛf/	0	0	1
/ɔːb/	0	0	3	/aɪt/	5	11	28	/aʊ/	0	0	15	/ɛg/	0	0	5
/ɔːd/	0	0	16	/aɪv/	5	5	10	/aʊð/	0	0	1	/ɛk/	0	0	11
/ɔːf/	0	0	1	/aɪz/	1	2	276	/aʊʧ/	0	0	7	/ɛks/	0	0	9
/ɔːk/	0	0	16	/ɜːʧ/	0	0	8	/aʊʤ/	0	0	1	/ɛkt/	0	0	45
/ɔːl/	1	2	29	/ɜː*/	0	0	25	/aʊd/	0	0	6	/ɛl/	5	12	17
/ɔːld/	0	0	1	/ɜːθ/	0	0	3	/aʊl/	0	0	8	/ɛlʧ/	0	0	2
/ɔːls/	0	0	1	/ɜːʤ/	0	0	10	/aʊn/	0	0	7	/ɛlʃ/	0	0	1
/ɔːlt/	0	0	10	/ɜːb/	0	0	3	/aʊnʤ/	0	0	2	/ɛld/	0	0	3
/ɔːm/	0	0	13	/ɜːd/	0	0	4	/aʊnd/	0	0	20	/ɛlm/	0	0	1
/ɔːn/	0	0	15	/ɜːf/	0	0	2	/aʊns/	0	0	10	/ɛlp/	0	0	3
/ɔːnʧ/	0	0	2	/ɜːk/	0	0	8	/aʊnt/	0	0	10	/ɛlt/	0	0	4
/ɔːnt/	0	0	6	/ɜːl/	0	0	8	/aʊs/	0	0	6	/ɛlv/	0	0	2
/ɔːp/	0	0	2	/ɜːm/	0	0	8	/aʊst/	0	0	3	/ɛm/	0	0	4
/ɔːs/	0	0	8	/ɜːn/	2	3	13	/aʊt/	0	0	11	/ɛmpt/	0	0	3
/ɔːst/	0	0	1	/ɜːp/	0	0	4	/aʊz/	0	0	10	/ɛn/	0	0	4
/ɔːt/	0	0	24	/ɜːs/	0	0	14	/ɪə*/	1	4	39	/ɛnʧ/	0	0	9
/ɔːz/	0	0	2	/ɜːst/	1	1	2	/ɪəd/	0	0	1	/ɛnʤ/	0	0	3
/eɪ/	2	4	64	/ɜːt/	1	1	20	/ɪənz/	0	0	1	/ɛnd/	5	7	36
/eɪð/	0	0	3	/ɜːv/	0	0	12	/ɪənt/	0	0	3	/ɛns/	0	0	8

Table 4.6: Neighbourhood density analysis of English regular verbs - Part 1

Rime	$l\alpha$	$l\beta$	R	Rime	$l\alpha$	$l\beta$	R	Rime	$l\alpha$	$l\beta$	R	Rime	$l\alpha$	$l\beta$	R
/ɛnt/	0	0	38	/ɪt/	7	8	89	/ʌd/	0	0	8	/iːv/	5	6	20
/ɛnz/	0	0	1	/ɪts/	0	0	1	/ʌf/	0	0	15	/iːz/	1	1	11
/ɛp/	0	0	3	/ɪv/	1	3	5	/ʌg/	0	0	15	/uː/	1	5	60
/ɛpt/	0	0	3	/ɪz/	0	0	4	/ʌk/	0	0	11	/uːð/	0	0	2
/ɛs/	0	0	51	/l̩/	0	0	311	/ʌkt/	0	0	10	/uːʧ/	0	0	2
/ɛsk/	0	0	1	/ɒʧ/	0	0	5	/ʌl/	0	0	8	/uːʒ/	0	0	1
/ɛst/	0	0	26	/ɒŋ/	0	0	8	/ʌlʧ/	0	0	1	/uːʤ/	0	0	2
/ɛt/	5	15	17	/ɒŋk/	0	0	3	/ʌlʤ/	0	0	3	/uːb/	0	0	3
/ɛv/	0	0	1	/ɒʃ/	0	0	10	/ʌlf/	0	0	1	/uːd/	0	0	17
/n̩/	0	0	105	/ɒθ/	0	0	1	/ʌlk/	0	0	2	/uːf/	0	0	11
/ɪ/	0	0	80	/ɒʤ/	0	0	4	/ʌlkt/	0	0	1	/uːk/	0	0	3
/ɪəʊ/	0	0	2	/ɒb/	0	0	10	/ʌlp/	0	0	2	/uːl/	0	0	13
/ɪʧ/	0	0	12	/ɒd/	0	0	6	/ʌlpt/	0	0	1	/uːm/	0	0	19
/ɪŋ/	11	13	4	/ɒf/	0	0	4	/ʌls/	0	0	3	/uːn/	0	0	17
/ɪŋk/	6	9	11	/ɒft/	0	0	1	/ʌlt/	0	0	5	/uːp/	0	0	13
/ɪŋks/	0	0	1	/ɒg/	0	0	13	/ʌm/	1	3	13	/uːs/	0	0	19
/ɪʃ/	0	0	49	/ɒk/	0	0	26	/ʌmf/	0	0	1	/uːst/	0	0	2
/ɪʤ/	0	0	29	/ɒks/	0	0	6	/ʌmp/	0	0	14	/uːt/	1	3	39
/ɪb/	0	0	6	/ɒkt/	0	0	1	/ʌn/	1	4	5	/uːv/	0	0	9
/ɪd/	3	7	2	/ɒl/	0	0	1	/ʌnʧ/	0	0	7	/uːz/	2	2	27
/ɪf/	0	0	4	/ɒlf/	0	0	1	/ʌnʤ/	0	0	4	/æʧ/	1	1	12
/ɪft/	0	0	8	/ɒlv/	0	0	8	/ʌnd/	0	0	2	/æŋ/	1	2	9
/ɪg/	1	1	6	/ɒm/	0	0	1	/ʌnt/	0	0	10	/æŋk/	0	0	16
/ɪk/	1	1	23	/ɒmp/	0	0	4	/ʌp/	0	0	5	/æʃ/	0	0	20
/ɪks/	0	0	8	/ɒmpt/	0	0	1	/ʌpt/	0	0	5	/æʤ/	0	0	1
/ɪkt/	0	0	12	/ɒn/	0	0	4	/ʌs/	0	0	8	/æb/	0	0	9
/ɪl/	1	1	18	/ɒnd/	0	0	6	/ʌsk/	0	0	2	/æd/	0	0	2
/ɪlʧ/	0	0	1	/ɒns/	0	0	1	/ʌst/	1	1	12	/æf/	0	0	1
/ɪld/	1	2	1	/ɒnt/	0	0	1	/ʌt/	2	3	10	/æg/	0	0	20
/ɪlk/	0	0	2	/ɒnz/	0	0	1	/ʌv/	0	0	2	/æk/	0	0	33
/ɪlm/	0	0	2	/ɒp/	0	0	28	/ʌz/	0	0	1	/æks/	0	0	8
/ɪlt/	0	0	5	/ɒpt/	0	0	1	/iː/	3	6	19	/ækt/	0	0	25
/ɪm/	1	1	9	/ɒs/	0	0	6	/iːð/	0	0	6	/æl/	0	0	1
/ɪmp/	0	0	7	/ɒst/	1	1	5	/iːʧ/	2	2	11	/ælp/	0	0	1
/ɪmps/	0	0	1	/ɒt/	0	0	20	/iːʃ/	0	0	1	/ælt/	0	0	1
/ɪn/	3	3	26	/ʊʃ/	0	0	3	/iːʤ/	0	0	1	/ælv/	0	0	1
/ɪnʧ/	0	0	7	/ʊd/	0	0	1	/iːd/	6	14	18	/æm/	0	0	11
/ɪnʤ/	0	0	11	/ʊk/	0	0	10	/iːf/	0	0	5	/æmp/	0	0	13
/ɪnd/	0	0	2	/ʊl/	0	0	3	/iːg/	0	0	3	/æn/	0	0	16
/ɪns/	0	0	5	/ʊlf/	0	0	1	/iːk/	2	3	12	/ænd/	1	4	10
/ɪnt/	0	0	13	/ʊm/	0	0	1	/iːl/	4	5	16	/æns/	0	0	2
/ɪp/	0	0	32	/ʊt/	1	1	4	/iːld/	0	0	4	/ænt/	0	0	8
/ɪps/	0	0	1	/ʌʧ/	0	0	4	/iːm/	1	1	15	/æp/	0	0	26
/ɪpt/	0	0	1	/ʌŋ/	0	0	1	/iːn/	2	2	20	/æps/	0	0	5
/ɪs/	0	0	23	/ʌŋk/	0	0	9	/iːp/	6	8	7	/æpt/	0	0	1
/ɪsk/	0	0	4	/ʌʃ/	0	0	7	/iːs/	0	0	12	/æs/	0	0	4
/ɪsp/	0	0	2	/ʌʤ/	0	0	12	/iːst/	0	0	1	/æt/	0	0	12
/ɪst/	0	0	21	/ʌb/	0	0	9	/iːt/	3	4	27	/æz/	0	0	1
/eɪə*/	0	0	1	/aɪə*/	0	0	25	/əʊ*/	0	0	1	/aʊə*/	0	0	15

Table 4.7: Neighbourhood density analysis of English regular verbs - Part 2

Class	Category	Stimulus	$l\alpha$	$l\beta$	R
Pseudo-Irregular	P	spling	11	13	4
		skring	11	13	4
		sprink	6	9	11
		cleed	6	14	21
		preed	6	14	21
		queed	6	14	21
		cloe	5	12	56
		froe	5	12	56
		plare	4	8	27
		quare	4	8	27
	I	fring	11	13	4
		ning	11	13	4
		frink	6	9	11
		cleef	0	0	5
		preek	2	3	15
		queef	0	0	5
		foa	5	12	56
		voa	5	12	56
		jare	4	8	27
		grare	4	8	27
	D	trisp	0	0	2
		nist	0	0	22
		plip	0	0	33
		gleef	0	0	5
		keeb	0	0	0
		meep	6	8	8
		goav	0	0	2
		joam	0	0	6
		flape	0	0	11
		blafe	0	0	2
Pseudo-Regular	P	plip	0	0	33
		glip	0	0	33
		brip	0	0	33
		greem	1	1	16
		pleem	1	1	16
		treem	1	1	16
		gloke	0	0	17
		proke	0	0	17
		slace	0	0	27
		nace	0	0	27
	I	brilth	0	0	0
		glinth	0	0	0
		plimph	0	0	0
		smeeb	0	0	0
		smeeg	0	0	3
		smeej	0	0	1
		ploab	0	0	3
		ploag	0	0	1
		smaib	0	0	0
		smaig	0	0	2
	D	frilg	0	0	0
		krilg	0	0	0
		smeelth	0	0	0
		smeenth	0	0	0
		smeerg	0	0	0
		ploamph	0	0	0
		ploanth	0	0	0
		smairg	0	0	0
		smairph	0	0	0

Table 4.8: Neighbourhood analysis of Prasada & Pinker's stimuli

Class	Category	Stimulus	$I\alpha$	$I\beta$	R
Pseudo-Irregular	P	gling	11	13	4
		strink	6	9	11
		veed	6	14	21
		pread	6	14	21
		krow	5	11	56
		plow	5	11	56
		zear	4	6	27
		pear	4	6	27
	I	bling	11	13	4
		pling	11	13	4
		geed	6	14	21
		cleed	6	14	21
		frow	5	11	56
		shrow	5	11	56
		jear	4	6	27
		mear	4	6	27
	D	blik	1	0	23
		plig	1	1	6
		geem	1	0	16
		cleep	6	8	7
		frowl	0	0	23
		shrowth	0	0	0
		jearp	0	0	0
		mearp	0	0	0
Pseudo-Regular	P	frip	0	0	32
		krip	0	0	32
		keam	1	1	16
		geem	1	1	16
		boke	0	0	17
		zoke	0	0	17
		gare	4	6	27
		vare	4	6	27
	I	yip	0	0	32
		jip	0	0	32
		feam	1	1	16
		leam	1	1	16
		doke	0	0	17
		voke	0	0	17
		yare	4	6	27
		nare	4	6	27
	D	yid	3	7	2
		jiz	0	0	4
		feap	6	8	7
		leat	3	4	27
		dofe	0	0	1
		vome	0	0	6
		yairth	0	0	0
		nairth	0	0	0

Table 4.9: Neighbourhood analysis of Lee's stimuli

the same rime and are thus inseparable by these measures, perhaps lending too much weight to the importance of onset in verb inflection.

The pseudoregulars also show anomalies: clearly the prototypical items have more neighbours (particularly, more regular neighbours) than the others, but 14 of the 19 intermediate and distant pseudoregulars have no neighbours of any type at all. These 14 items clearly do not belong in different categories on the basis of rime, nor can onset be seen as a factor, since most of them share onsets but are in different categories.

All of these intermediate and distant pseudoregular stimuli have no irregular and very few regular neighbours, and therefore would be expected to be regularised from either theoretical perspective outlined in the introduction: in symbolic terms, by the application of a default suffixation rule; and in connectionist terms, given that there are no irregular analogies to draw on, because of the type bias towards regularisation. That is, given that the overwhelming majority of English verbs are regular, one might expect a stimulus to be regularised if there is nothing to suggest otherwise.

The same criticisms apply to Lee's stimuli; again there is almost no numerical difference between the prototypical and intermediate categories for pseudoirregular or pseudoregular verbs, and again there are strange outliers: compare, for instance, 'gare' and 'vare', being *prototypical* pseudoregulars with 6 shared-rime irregulars and 27 shared-rime regulars, when 'leat' is classed as a *distant* pseudoregular with 4 shared-rime irregulars and, again, 27 shared-rime regulars. And again, 'cleep' is construed as a distant pseudoirregular, when it shares its rime with a large number of real irregular verbs.

4.4 Conclusion

This chapter has assessed the systems of neighbourhood classification used by Prasada & Pinker (1993) and Lee (1995), and has shown that their category definitions are flawed: in particular, a stimulus may be classified as a 'Distant' pseudoirregular (and therefore highly likely to be regularised) because of its distance from *one* group of irregular verbs, even though in practice it may be situated in the middle of another irregular group (and therefore more likely to be irregularised).

It has also been shown that, by using a neighbourhood classification scheme which reflects—for a given real or novel verb—the relative proximity of *all* verbs, it is possible to separate stimuli along a 'gradient of similarity' which is more objective than Prasada & Pinker's original Prototypical–Intermediate–Distant categories.

The remainder of this thesis will show how stimuli classified in this way were used:

- in four behavioural studies, to test the validity of Prasada & Pinker's experimental findings (chapter 5);
- in two electrophysiological studies, to assess and compare neural activation during regularisation and irregularisation of real and novel verbs (chapter 6);
- to test a set of connectionist simulations trained on an English past tense inflection task, to compare network productivity on novel stimuli with human behavioural findings (chapter 7).

Chapter 5

Past Tense Inflection of Novel Verbs

5.0 Introduction

In the previous chapter, it was shown how Prasada & Pinker (1993) concluded from their experimental work that irregular verbs are inflected by analogy to known verbs, while regulars are inflected by rule. Lee (1995) performed a similar experiment using an improved neighbourhood classification and found that effects of similarity were apparent for *both* irregularisation and regularisation. It has already been shown how both sets of findings are called into question due to the use of unsound neighbourhood classificatory systems, and a new system has been proposed which—it has been argued—is more thorough and objective.

This chapter describes the application of the new classification to a set of novel verb stimuli; a set of four experiments is described, which have implications for the single-route/dual-route debate in inflectional morphology.

In the following experiments, subjects were presented with novel verb stems and required to inflect them to appropriate past tense forms. The experiments are therefore, in essence, replications of Prasada & Pinker's and Lee's work, but using the new neighbourhood classification. By classifying the stimuli as having an above-median or below-median number of regular neighbours (HighR or LowR, respectively) and an above- or below-median number of irregular neighbours (HighI or LowI, respectively), one can assess the separate and conjoint impact of the two factors on the probability of regularisation (or irregularisation) of novel verbs.

According to the hybrid dual-route theory propounded by Pinker (1991), irregular

neighbourhood density will have an effect on the regularisation rate of a novel verb, since similarity to known irregulars may lead to above-threshold 'partial matches' in the associative irregular component, causing the default regular rule to be blocked and the novel verb to be irregularised. If the novel verb is not similar to known irregulars, the model predicts that it will be inflected by the regular rule component. The dual-route theory would therefore predict that HighI novel verbs are more likely to be irregularised than LowI novel verbs. This prediction is also expected from a connectionist perspective: single-route connectionist networks process input patterns by analogy to previously-experienced patterns, and are thus also more likely to irregularise HighI rather than LowI novel verbs.

However, the two theories make conflicting predictions about the influence of *regular* neighbours on the regularisation rate of novel verbs. According to Pinker's hybrid dual-route model, regular past tense forms are not stored in the lexicon, but are generated afresh every time they are processed, according to a default, explicit suffixation rule. The theory therefore predicts that the number of regular neighbours possessed by a novel verb will not have an effect on the probability of it being regularised or irregularised. Connectionist theory, on the other hand, predicts that regular neighbourhood density *will* have an impact on regularisation rate, since patterns are inflected by analogy to *all* verbs, and not just to irregulars.

The two theoretical approaches therefore offer mutually-exclusive and empirically testable predictions about the influence of neighbourhood density on the inflection of novel verbs. These predictions are tested here.

5.1 Method

Four experiments were conducted, two using an auditory, and two using a visual, presentation modality: in other words, in two experiments stimuli were presented in spoken form, and in the other two they were presented on a computer monitor. Within each of these presentation types, one experiment involved presenting the stimuli in the context of a cloze passage, and one involved presenting them with no context at all. This was done in order to establish whether presentation modality or the presence of context have any effect on novel verb inflection: this is useful to find out, as experimental data are often compared and contrasted directly even though different methodologies have been employed in the presentation of stimuli: such comparisons may not be valid if presentation modality and/or the presence or absence of context can be shown to result in differential processing.

5.1.1 Subjects

20 adult native speakers of Southern British English with standard pronunciation were used in each of the four experiments (mean age = 21.4 years; range = 19.6–23.9). The use of specific accent and dialect should optimise the probability, first, that stimuli would be perceived in a similar manner; and secondly, that inflection of existing verbs would be consistent across subjects (cf. 'plead': inflects regularly to 'pleaded' in Southern British English and irregularly to 'pled' in Scots English).

5.1.2 Materials

Context

The context for the two 'With Context' experiments comprised a cloze passage of approximately 150 lines—the short story used in Lee's original paper. The passage contained 48 blanks in which past tense novel verb forms could be inserted by subjects, and originally (in Lee's work) also had 12 blanks for novel adjectives, to prevent the subjects from treating the experiment purely as a past tense exercise. In these experiments, these were omitted, as it was found in Lee's work that subjects gave similar answers, whether or not the adjectives were included (Lee, *pers. comm.*). The cloze passage (and a practice passage which preceded the experiment) may be found in appendices C and D.

Stimuli

48 stimuli were used in these experiments. 24 were classed as 'pseudoirregular', having been classified by quantitative comparison with all neighbourhoods containing an *irregular* verb; the other 24 were classed as 'pseudoregular', having been classified by comparison with all neighbourhoods containing at least one *regular* verb. Note that pseudoregulars are not *per se* more likely to be regularised than pseudoirregulars, or vice versa. Rather, the terms are being used—as in Prasada & Pinker's paper—to indicate the frame of reference within which the stimuli were classified.

First, quantitative analyses of all verbs contained in the CELEX database were conducted, as described in detail in chapter 4 above. For methodological balance, it

was desirable to select an equal number of representative novel stimuli in each of the following 8 categories:

- HighILowR, HighIHighR, LowILowR and LowIHighR pseudoirregulars;
- HighILowR, HighIHighR, LowILowR and LowIHighR pseudoregulars.

To this end, 6 clusters were selected for each of these 8 verb categories according to their suitability to represent that category. As far as possible, Id of pseudoirregular clusters (that is, the number of irregular verbs in the neighbourhood of a pseudo-irregular cluster which are inflected by a different irregular paradigm) was stabilised at zero, in order to minimise competition effects between different types of irregularisation within a single stimulus. This was impossible to achieve in all cases: HighIsLowId clusters[1] are very few in number; in four cases, therefore, Id could not be maintained at exactly zero.

One novel verb was then selected for each of the 48 pseudoirregular and pseudoregular clusters; these are shown in Table 5.1 below.

The stimuli were drawn from a pool of 4282 monosyllabic novel verbs created by permuting all mono-, bi- and tri-gram onsets and bi- and tri-gram rimes of English verbs[2] and filtering out all results which were identical to real words, either in sound or orthography, or which were phonotactically or orthographically illegal. Each of the 48 selected stimuli was then examined to ensure that neither a probable irregular

[1] That is, clusters containing an above-median number of irregulars *which are all inflected by the same paradigm.*

[2] From the CELEX database.

Class	ls,ld,R	Cluster	lsα	ldα	lsβ	ldβ	R	Stimulus
Pseudo-Irregular	H L H	/əʊ/↦/uː/	4	1	8	4	56	shrow
		/ɛə/↦/ɔː/	4	0	6	0	27	plair
		/ɛnd/↦/ɛnt/	5	0	7	0	36	strend
		/ɪt/↦/ɪt/	4	3	5	3	89	thrit
		/iːd/↦/ɛd/	6	0	14	0	21	veed
		/iːl/↦/ɛlt/	3	1	4	1	18	preel
	H L L	/iːʧ/↦/ɔːt/	2	0	3	0	11	steech
		/aɪv/↦/əʊv/	5	0	5	0	10	scrive
		/ɜːn/↦/ɜːnt/	2	0	3	0	13	purn
		/ɛd/↦/ɛd/	3	1	3	1	12	cled
		/ʌt/↦/ʌt/	2	0	3	0	10	plutt
		/iːp/↦/ɛpt/	6	0	8	0	7	cleep
	L L L	/aɪk/↦/ʌk/	1	0	1	0	8	drike
		/ɜːst/↦/ɜːst/	1	0	1	0	2	hurst
		/ʌst/↦/ʌst/	1	0	1	0	11	wust
		/ɪm/↦/æm/	1	0	1	0	9	thrim
		/ɒst/↦/ɒst/	1	0	1	0	5	jost
		/æv/↦/æd/	1	0	1	0	0	rav
	L L H	/ɔːl/↦/ɛl/	1	0	2	0	29	grawl
		/aɪn/↦/ɒn/	1	0	2	0	42	stine
		/aɪz/↦/əʊz/	1	0	2	0	277	plyse
		/ɜːt/↦/ɜːt/	1	0	1	0	21	lurt
		/ɪk/↦/ʌk/	1	0	1	0	23	thrick
		/ɪl/↦/ɪlt/	1	0	1	0	18	pril
Pseudo-Regular	- H H	/ɑːst/↦/ɑːstɪd/	-	1	-	8	8	thrast
		/əʊld/↦/əʊldɪd/	-	1	-	4	9	crold
		/ʌm/↦/ʌmd/	-	1	-	3	13	vum
		/iːm/↦/iːmd/	-	1	-	1	15	bleem
		/iːz/↦/iːzd/	-	1	-	1	11	creeze
		/æŋ/↦/æŋd/	-	1	-	2	9	splang
	- H L	/aɪnd/↦/aɪndɪd/	-	4	-	8	4	plynd
		/ɪg/↦/ɪgd/	-	1	-	1	6	splig
		/ɪld/↦/ɪldɪd/	-	1	-	2	1	sprild
		/ɪv/↦/ ɪvd/	-	1	-	3	5	miv
		/ɪd/↦/ ɪdɪd/	-	3	-	7	2	sprid
		/ʌn/↦/ʌnd/	-	1	-	4	5	chun
	- L L	/aɪf/↦/aɪft/	-	0	-	0	1	chife
		/ɜːf/↦/ɜːft/	-	0	-	0	2	purf
		/ɪlm/↦/ɪlmd/	-	0	-	0	2	vilm
		/ɪps/↦/ɪpst/	-	0	-	0	1	fipse
		/ɒnt/↦/ɒntɪd/	-	0	-	0	1	zont
		/ælp/↦/ælpt/	-	0	-	0	1	calp
	- L H	/ɔːd/↦/ɔːdɪd/	-	0	-	0	16	tord
		/aɪl/↦/aɪld/	-	0	-	0	17	scrile
		/aɪm/↦/aɪmd/	-	0	-	0	11	thime
		/ɜːs/↦/ɜːst/	-	0	-	0	14	lerse
		/ɪʧ/↦/ɪʧt/	-	0	-	0	29	flitch
		/ɪp/↦/ɪpt/	-	0	-	0	32	glip

Table 5.1: Neighbourhood analysis of novel verb stimuli

inflection, nor a regular past tense suffixation, would produce a real verb, either phonologically or orthographically.

For the pseudoregulars, it was hypothesised that two of the four sets of stimuli used (LowILowR and LowIHighR) would show high regularisation rates, because by definition (given that $\mu(\mathsf{I}) = 0$) they have no irregular verbs in their neighbourhoods. From either theoretical perspective, one would expect these stimuli (which rhyme with *no* irregular verbs) to be regularised: either by a default rule in a symbolic dual-route system, or by the general pull towards regularisation in a single-route associative system. More importantly, if a statistical difference could be shown between the other two categories which *do* have neighbouring irregulars (HighIHighR and HighILowR), then it would seem to be the case that the number of regulars in the neighbourhood of a stimulus does have an effect on regularisation probability. Similarly, for pseudoirregulars, both theoretical standpoints predict differential regularisation rates for HighI vs. LowI pseudoirregulars, but dual-route theory—unlike the connectionist perspective—predicts no effect of Rdensity (*i.e.* HighR vs. LowR) on regularisation rate. This, therefore, is the critical area of investigation in these experiments.

5.1.3 Procedure

In each of the four paradigms, there was a brief test period before each subject began the task to ensure that it was understood what he/she was to do.

Auditory/With Context Experiment Subjects were seated at a computer in a quiet room. The context passage was presented line by line on the screen at a pace determined by the subject: every time the subject pressed a key, the next line of the passage was presented. Whenever a line appeared which contained a blank, the computer automatically presented one of the recorded stimuli (selected randomly) via the computer's loudspeakers. On hearing the stimulus, the subject verbally produced the most appropriate-sounding past tense form in the context of the cloze passage. Responses were recorded onto tape.

Auditory/No Context Experiment In this experiment, subjects were seated at a computer in a quiet room. After instructions had been presented on the screen, the stimuli were presented (in random order) over loudspeakers as before, every time the subject pressed a key. Responses were recorded onto tape.

Visual/With Context Experiment This experiment was identical to the first 'With Context' experiment, except that when a line of context containing a blank appeared, a stimulus was presented visually rather than auditorily. When the subject pressed a key for the visual stimulus to be presented, a randomly-selected stimulus was presented in the centre of the screen in large, lower case letters and the subject verbally provided an appropriate past tense inflection, which was recorded. The stimulus remained on the screen until the subject pressed a key, after which the context passage reappeared on screen as before.

Visual/No Context Experiment Stimuli were presented in exactly the same way as for the 'Visual/With Context' experiment, except that no context passage was presented. The screen was blank until the subject pressed a key, when a stimulus appeared as before, remaining on screen until a key was pressed again, triggering the presentation of the next stimulus; this process was repeated until all 48 stimuli had been presented once.

5.2 Results

In each of the four experiments, 20 subjects were each presented with all 48 stimuli exactly once. The resulting 3840 responses were coded as regularised if the subject inflected the stimulus according to regular suffixation rules. That is, if the subject added /-ɪd/ to a stem-final alveolar plosive, /-t/ to a stem-final voiceless consonant other than /t/, or /-d/ to any other stimulus ending, the response was classified as a regularisation; otherwise, responses were coded as irregularised. There was a potential complication here, since a 'garbage' response would be coded as irregularised. However, it was necessary to classify subject responses in this manner in order to maintain optimal comparability with the results of the Prasada & Pinker and Lee studies, both of which used the same classification; in the event, very few stimuli were corrupted in this way by the subjects ($N < 10$ out of 3840 responses).

Results for pseudoregular and pseudoirregular items were analysed separately, as they are definitionally incomparable, since 'High' and 'Low' are based on different median values dependent on the context of regular or irregular neighbourhood density.

5.2.1 Results for Pseudoirregular Stimuli

Figure 5.1 shows the results of the four experiments for the pseudoirregular stimuli. In each experiment, there were a maximum 120 irregularisations to be made (y axis) for each of the four pseudoirregular stimulus cells (HighI LowR, HighI HighR, LowI LowR, and LowI HighR on the x axis), as 20 subjects were presented with six pseudoirregular stimuli per cell per experiment. A similarity gradient is clear from this graph: there was a decrease in irregularisation rate as a function of neighbourhood density. The differences between cells are clear; but are they significant? Furthermore, was there an effect of Presentation Modality or Context on irregularisation rate?

To answer these questions, a 4-way repeated-measures (2 within, 2 between) analysis of variance (ANOVA) was carried out on the pseudoirregular results (F1), as well as an items analysis (F2).

The two within-subjects variables were the number of irregulars (HighI vs. LowI) and regulars (HighR vs. LowR) in the neighbourhoods of the stimuli; and Presentation Modality (Auditory vs. Visual) and Context (vs. No Context) were the between-subjects variables. In the items analysis, Presentation and Context were the within-items variables, and irregular and regular neighbourhood density were the between-items variables.

A very strong overall effect of the number of irregulars neighbouring a stimulus was found ($F1(1,76) = 460.617; p < 0.0001$; $F2(1,20) = 724.638; p < 0.0001$), as was a strong effect of the number of neighbouring regulars ($F1(1,76) = 110.452; p < 0.0001$;

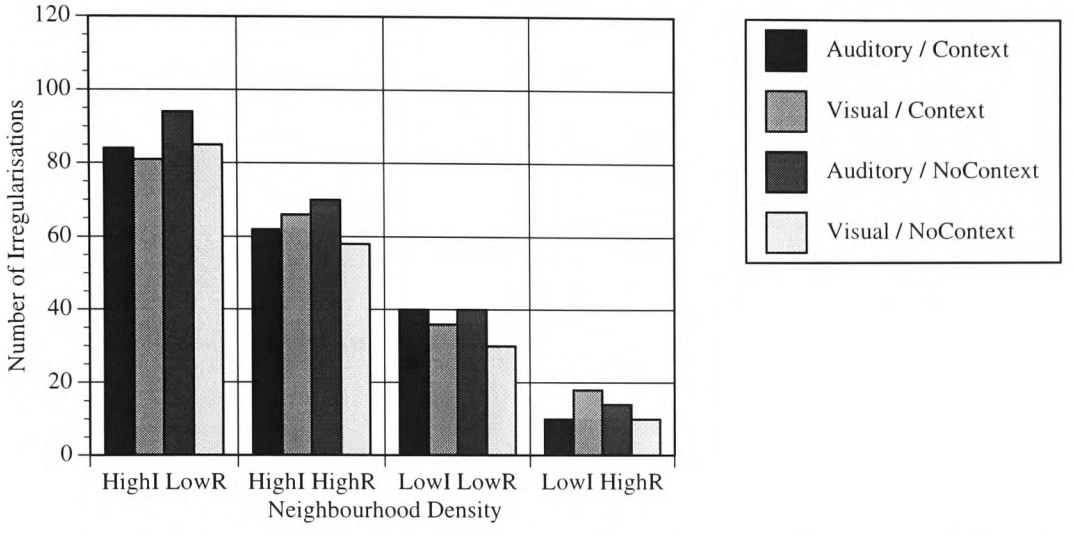


Figure 5.1: Overall results for pseudoirregular stimuli across all paradigms

$F2(1, 20) = 146.739; p < 0.0001$); there was no higher-order interaction between these two within-subject variables. This shows that the number of irregular verbs and the number of regular verbs in a stimulus' rime neighbourhood are both important factors in how that stimulus is inflected: HighI stimuli were more likely to be irregularised than LowI stimuli, and LowR stimuli showed higher irregularisation rates than HighR novel verbs. No other main effects were found.

One interaction was found, between the two between-subject variables Presentation Modality and Context ($F1(1, 76) = 6.957; p = 0.0101$; $F2(1, 20) = 5.879; p = 0.0249$), although no main effect of Context was found and the main effect of Presentation Modality was only nearly significant in the subject analysis ($F1(1, 76) = 3.913; p = 0.0515$), and insignificant in the items analysis ($F2(1, 20) = 3.062; p = 0.0954$). This Modality $\times$ Context interaction is shown in Figure 5.2 below.

A post-hoc Scheffé analysis of this interaction showed a statistically significant

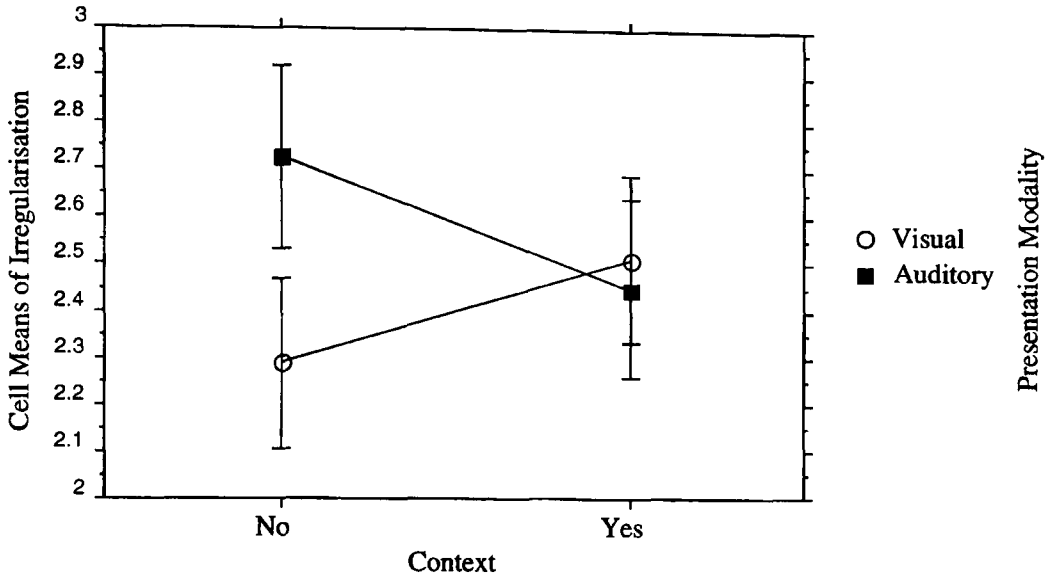


Figure 5.2: Modality × context interaction for pseudoirregular stimuli

difference between Auditory and Visual presentation in the No-Context paradigms ($F = 10.65; p < 0.001$), but not in the Context paradigms ($F = 0.22; p > 0.05$). Furthermore, the presence or absence of Context did make a difference in the Auditory paradigms ($F = 4.21; p < 0.05$), but not in the visual paradigms ($F = 2.82; p > 0.05$).

Figure 5.2 shows that subjects were more likely to irregularise pseudoirregular stimuli when they were presented aurally than when they were presented visually, but *only* when they were presented out of context. In other words, there seems to be a distinction between inflecting a stimulus one sees and a stimulus one hears, but this distinction is lost when the stimuli are presented embedded in real sentences.

5.2.2 Results for Pseudoregular Stimuli

The analysis of pseudoregular items was carried out in an identical manner to that for pseudoirregular items, except that the number of regularisations was counted,

rather than the number of irregularisations. Again, two 4-way repeated-measures ANOVAs were calculated: in the subjects analysis, the number of irregulars and regulars in the neighbourhoods of the stimuli were the two within-subjects variables, and Presentation Modality (Auditory vs. Visual) and Context (vs. No Context) were the between-subjects variables; in the items analysis, Presentation and Context were the within-items variables, and irregular and regular neighbourhood density were the between-items variables.

The overall results for the pseudoregular stimuli are plotted below in Figure 5.3. Again a similarity gradient may clearly be seen.

As with the pseudoirregular stimuli, a very strong effect of irregular neighbourhood density was found ($F1(1,76) = 414.226; p < 0.0001$; $F2(1,20) = 414.529; p < 0.0001$), as was a strong effect of the number of regular neighbours ($F1(1,76) = 6.841; p = 0.0107$; $F2(1,20) = 6.108; p = 0.0226$): again, LowI stimuli were more likely to be regularised than HighI stimuli, and HighR stimuli were more likely to be regularised than LowR novel verbs. Furthermore, there was an interaction between the two ($F1(1,76) = 62.570; p < 0.0001$; $F2(1,20) = 65.942; p < 0.0001$). This interaction is depicted in Figure 5.4.

A post-hoc Scheffé analysis on this irregular × regular neighbourhood density interaction showed that all parts of the interaction are significant, although some more so than others. For the LowR stimuli, the difference between those with and without irregulars in the neighbourhood[3] was highly significant ($F = 53.23; p < 0.001$). This

[3]It will be recalled that LowI pseudoregular stimuli contain *no* irregulars in their rhyming neigh-

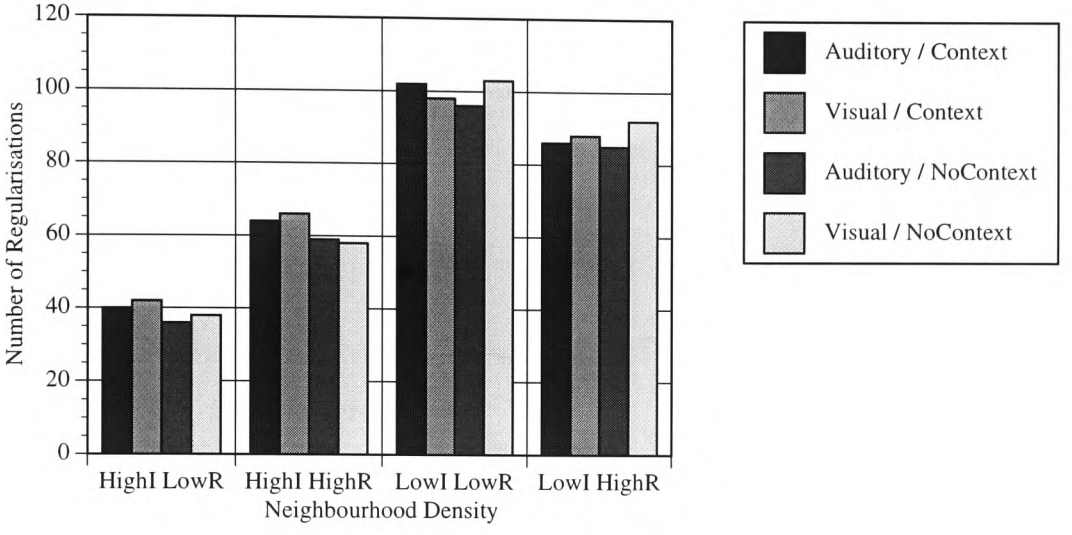


Figure 5.3: Overall results for pseudoregular stimuli across all paradigms

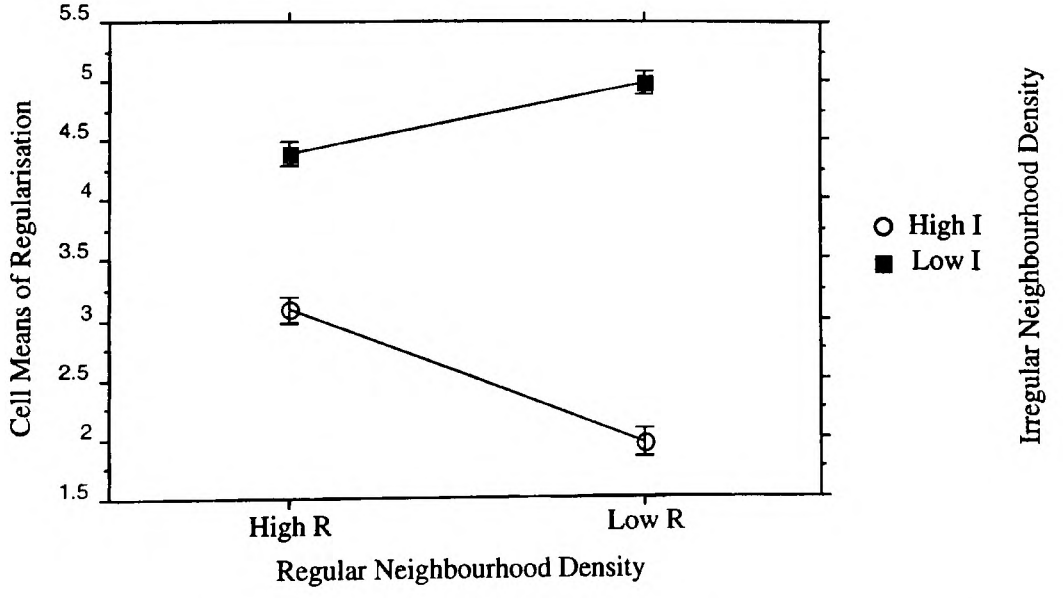


Figure 5.4: Irregular × regular interaction for pseudoregular stimuli

was also true of the HighR stimuli ($F = 6.70; p < 0.05$). Of the HighI stimuli, the number of regulars in the neighbourhood made a significant difference ($F = 14.63; p < 0.01$); and this difference was also clear for the LowI items ($F = 6.50; p < 0.05$). There were no other interactions.

5.3 Discussion

5.3.1 Neighbourhood Density Effects

Findings for Pseudoirregulars

From Figure 5.1 above, a similarity gradient may easily be seen between each cell, showing a decrease in the number of irregularisations made as a function of distance from irregular and regular clusters. The two middle cells of Figure 5.1 hint at the relative importance of the two factors, when placed in the context of the other two cells in the graph. The fact that the HighIHighR stimuli were irregularised more than the LowILowR novel verbs, orders the cells with the two HighI cells showing the highest irregularisation rates, and then the two LowI cells showing lower irregularisation rates, suggesting that the number of irregular verbs in the neighbourhood cluster was a more important factor than the number of regulars; however, *within* High and Low irregular clusters, it can be seen that the number of regulars was, in turn, influential.

The findings for pseudoirregulars therefore support Bybee & Moder's (1983) conclusion that novel verb stems are inflected by analogy to known irregular stems, and that the probability of a novel verb being irregularised is a function of its similarity

bourhoods.

to known irregulars. This also supports the 'neighbourhood structure' methodology used here as being an accurate way of classifying influences on target verbs, and suggests that rime alone is a very strong predictor of the analogous 'pull' on a verb, even without taking—for example—verb frequencies into account. More importantly, the results showed that *regular* verbs in the stimulus neighbourhoods also affected irregularisation rate. If regular verbs were inflected by a default suffixation rule, then the number of regulars in a target verb's neighbourhood would not be expected to make a difference to the probability of irregularisation, unless they were also stored in associative memory alongside irregular stem–past pairs, in which case the rule may not be necessary.

Findings for Pseudoregulars

There are two important points to note here: the first is the high probability of regularisation of those items which have no irregular neighbours and few regular ones (LowILowR stimuli).

This result is easily explicable from either an analogy-based or a symbolic perspective. The symbolic argument would hold that LowILowR items would be regularised as there are no irregulars to compete with the default suffixation rule, and the analogy theory would maintain that they would be regularised simply because of the presence of regular neighbours (and absence of irregular ones), and also the general 'pull' of the network to regularisation[4].

[4]Owing to the high type frequency of regular verbs compared to irregulars.

The second point is that there is a significant difference between HighILowR and HighIHighR stimuli: in other words, given a high number of irregulars in a stimulus' neighbourhood, the number of regulars nonetheless had an impact. This was less the case for the two cells of zero-irregular stimuli. Symbolic theory would suppose that the number of regulars in the high-irregular neighbourhoods should make no difference to regularisation rate: this supposition is not supported by these results. On the contrary, these data support single-route theories, which predict that regular neighbourhood density *should* affect regularisation rate.

5.3.2 Effects of Presentation Modality and Context

The only significant effects of Presentation or Context were found in a Presentation $\times$ Context interaction in the pseudoirregular paradigm. The interaction indicates that when the stimuli appear in the context of a sentence, it makes no difference whether they are presented visually or auditorily. However, if they are presented without any context, as free-standing items, they are more likely to be irregularised if heard than if seen.

It is interesting to note that this effect was not replicated for the pseudoregular stimuli.

5.4 Conclusion

This chapter has shown how novel stimuli were regularised or irregularised by adults at different rates, depending on their neighbourhood structure. Specifically, effects

of irregular and regular neighbourhood density were found for both pseudoregular and pseudoirregular novel verbs: in both cases, novel verbs with an above-median number of irregulars in their neighbourhoods (*i.e.* Highl stimuli) were more likely to be inflected by an irregular paradigm than stimuli sharing their rime with few irregulars (or none at all). Furthermore, and critically, novel verbs with an above-median number of regulars in their neighbourhoods were more likely to be regularised than those with few regular neighbours. It has been shown that this finding conflicts with the dual-route theory of inflectional morphology in which, since regular inflections are generated 'on-the-fly' (rather than being stored—like irregulars—in the lexicon), regular neighbourhood density should have no effect on irregularisation rate. This set of experiments therefore argues strongly against the hybrid dual-route theory and supports the alternative connectionist account, which predicts that irregular and regular neighbourhood density should *both* affect the rate of regularisation of novel verbs. Moreover, the direct conflict between the results of these experiments and the findings reported by Prasada & Pinker (1993) suggests that their data may—as suggested in chapter 4—have resulted from methodological flaws in the design of their stimulus categories.

Chapter 6

Electrophysiological Studies of Past Tense Inflection

6.0 Introduction

Inflectional morphology is a useful field for the investigation of the psycholinguistic processes underlying grammar. As has already been described, it has been argued that these processes are genetically determined, modular, rule-implementing systems (*e.g.* Pinker, 1994). Symbolic theories argue that the production of regular and irregular inflections involves two distinct routes (Marcus, Brinkman, Clahsen, Wiese & Pinker, 1995; Pinker & Prince, 1988; Prasada & Pinker, 1993); some connectionists, on the other hand, suggest that a single associative memory system could—and can—cope with both regular and irregular morphology (*e.g.* Rumelhart & McClelland, 1986; MacWhinney & Leinbach, 1991)[1].

Neurophysiological studies have suggested support for the dual-route theory with evidence from Alzheimer's patients, anomics and agrammatics (Ullman et al., 1993; Ullman, Corkin, Coppola & Hickok, 1997). Some Alzheimer's and anomic patients are impaired on production of irregular but not regular verbs, while agrammatics may show the reverse deficit, and it has been suggested that a semantic system may be involved in the production of inflected irregular (but not regular) verbs.

Ullman et al. (1994) make a specific claim about localisation of inflectional function in the brain, based on observation of brain-damaged patients. They predict that "dominant hemisphere anterior neocortical areas (*e.g.* Broca's area) subserve the computation of regulars, [as do the basal ganglia,] while dominant hemisphere

[1]See chapter 3 above for a detailed discussion.

posterior neocortical areas (*e.g.* areas 39, 40) subserve the computation of irregulars".

Recent psychological research into the English past tense system has been both prolific and useful, but is almost entirely based on behavioural experimentation. Very few neuroimaging studies have addressed the issue of inflectional morphology: of these, recent PET (Positron Emission Tomography) and ERP (Event-Related Potential) experiments will be discussed here.

6.0.1 Jaeger et al. (1996)

Jaeger et al. (1996) conducted a Positron Emission Tomographic study to investigate the locus or loci of verbal inflectional processes, by measuring neural activation in subjects inflecting regular, irregular and nonce verb stems to the past tense. They asserted that single-route theories of inflectional morphology would predict activation in the posterior (temporal and parietal) association cortices, in other words those areas thought to be involved in lexical and associative memory. Dual-system theories, on the other hand, would hypothesise that—since functionally distinct systems are involved in the production of regular and irregular inflections—physically distinct brain areas will be activated, although some overlap would be expected since both routes of the system involve the application of the grammatical category 'past tense' to the input. Jaeger et al. suggested that the explicit regular rule would be reflected in activation in the left frontal lobe (which is associated with grammatical processing), whereas irregular lexical retrieval would be likely to be reflected in temporal or parietal activity.

They tested these predictions by giving subjects five lists, each of 46 words, as follows:

Read verb: English verb stems to be read aloud;

Read nonce: novel verb stems to be read aloud;

Regular past: regular verb stems to be inflected to the past tense;

Irregular past: irregular verb stems to be inflected to the past tense; and

Nonce past: novel verb stems to be inflected to the past tense.

It should be noted that the novel verb forms to be inflected (*i.e.* list 5) were constructed so as not to be similar to any irregular subregularity; there were, for instance, no novel stimuli ending in /ip/, /ɪŋ/, /aɪ/ or /ɛnd/. This was done to optimise the probability that subjects would inflect the forms by the regular paradigm.

Each word in each list was presented individually on a computer monitor, 1 stimulus every 2 seconds, remaining on screen for 1250 msec., and subjects were required to state aloud the appropriate response as quickly as possible. Brain activity was measured using a positron emission tomograph (PET). The 5 different verb types were presented in separate lists since PET scanners have accurate spatial resolution but poor temporal resolution: subjects were therefore required to perform a single task (*e.g.* regular inflection) for at least 90 seconds in order for at least 60 seconds of accurate data to be recorded (the minimum amount required for statistically significant analysis).

Brain Area	Regular Past	Irregular Past	Nonce Past
I. Frontal Lobe			
L. motor strip	***	***	***
L. anterior insula	*	*	
L. inf. frontal g.	*	**	***
L. ant. cingulate g.		*	
R. motor strip	*		
II. Temporal, Parietal			
L. sup. temporal g.			***
L. mid. temporal g.			***
III. Occipital Lobe			
L. calcarine cortex	***	***	
L. cuneus	*	***	***
L. lingual g.	*		
L. inf. occipital g.	*	***	***
L. fusiform g.	*	***	***
L. mid. occipital g.		***	***
R. calcarine cortex	***	***	***
R. cuneus	*	***	***
R. mid. occipital g.			***
IV. Cerebellum			
R. hemisphere	***	***	***
L. hemisphere	***	***	***
midline	***	***	***
V. Other Regions			
midbrain		***	***
colliculi		***	***
L. pulvinar		***	***
R. hippocampus		*	*

Table 6.1: Brain regions activated in Jaeger et al. study

Their results, showing activation levels for the three past tense conditions as compared to the resting state, are summarised in Table 6.1.

In Table 6.1 above, column 1 shows the brain regions involved in the inflectional tasks, and columns 2–4 show the significance of the activation generated by 3 tasks when compared with the resting state[2]. For ease of reference, Table 6.1 is divided into

[2]Asterisks indicate the level of significance of activity: *** = significant at Bonferroni 0.05 threshold; ** = significant at omnibus 0.001 threshold; * = significant at omnibus 0.01 threshold.

five horizontal sections, to separate the cytoarchitectonic areas under consideration. Note that in the Table, 'L' and 'R' stand for left and right, 'g' stands for gyrus, and 'sup.', 'mid.', 'inf.' and 'ant.' are short for superior, middle, inferior and anterior respectively.

Jaeger et al. (1996) found that all tasks involved widespread brain activity compared to the resting state. All parts of the cerebellum were active across all inflection conditions, as were left motor strip and left inferior frontal gyrus in frontal lobes; and left and right cuneus, left inferior occipital gyrus, left fusiform gyrus and right calcarine cortex (in occipital lobes). This is to be expected, since the cerebellum is considered to play a significant function in linguistic (and other) tasks (see Schmahmann (1991) and Leiner, Leiner & Dow (1993) for review); left motor strip is active as all tasks involve speaking words as a response; and, as all tasks also involved reading words, it follows that calcarine (primary visual) cortex and the other occipital areas (secondary and associative visual cortex) were active across conditions.

For regular inflection, uniquely activated areas included right motor strip and left lingual gyrus. The only uniquely activated area in the irregular inflection task was left anterior cingulate gyrus, while the nonce task elicited unique activity in left superior and middle temporal gyrus and right middle occipital gyrus. The two 'real verb' tasks (*i.e.* columns 2 and 3) shared activity in left anterior insula and left calcarine cortex; this activity was not present in the nonce verb condition. Finally, left middle occipital gyrus, midbrain, colliculi, left pulvinar and right hippocampus showed similar patterns of activity for the irregular verb and nonce verb conditions,

but were inactive in the regular paradigm. Note that, conversely, no areas shared activity for the regular and nonce conditions only.

No other brain areas showed levels of activation compared to the resting state in these three tasks.

Jaeger et al. showed that the regular inflection task involved the smallest amount of brain activation, while the irregular verb and nonce verb conditions showed much more widespread activity. They also maintained that the activation pattern in the nonce condition was more similar to the activation for irregular verbs than for regulars, while also showing significant activation in other areas (*viz.* left superior and middle temporal gyrus, and right middle occipital gyrus).

They also compared the data from the three inflectional tasks with the data from the relevant 'read aloud' task ('read verb' or 'read nonce'), by subtracting the activity generated by reading/speaking from the total task activation, leaving the brain activity generated by the process of past tense generation alone. These results are summarised in Table 6.2.

They found that left middle and inferior frontal gyri (including Broca's area) were active in all three inflection conditions. Left dorsolateral prefrontal cortex and left anterior cingulate gyrus were active in the 'regular past minus read verb' and 'nonce past minus read nonce' analyses, but not in the 'irregular past minus read verb' analysis, while left superior frontal gyrus, left superior parietal lobule, left inferior occipital gyrus, the cerebellum and the colliculi were active in the 'irregular minus read' and 'nonce minus read' analyses, but not in the 'regular minus read' one.

Brain Area	Regular Past minus Read Verb	Irregular Past minus Read Verb	Nonce Past minus Read Nonce
I. Frontal Lobe			
L. motor strip			*
L. mid. frontal g.	*	**	*
L. inf. frontal g.	*	***	**
L. dorsolateral prefrontal	*		*
L. superior frontal g.		***	*
L. ant. cingulate g.	*		*
II. Temporal, Parietal			
L. mid. temporal g.		*	
L. inf. parietal lob.	*	**	
L. sup. parietal lob.		**	*
L. precuneus (mesial 7)	*	***	
R. precuneus (mesial 7)	*	*	
III. Occipital Lobe			
L. calcarine cortex		*	
L. cuneus			*
L. inf. occipital g.		*	*
L. fusiform g.		*	
L. sup. occipital g.			*
IV. Cerebellum			
R. hemisphere		**	*
L. hemisphere		**	*
midline		**	*
V. Other Regions			
colliculi		*	*

Table 6.2: Past tense tasks compared to reading tasks

Bilateral precuneus and left inferior parietal lobule were involved in both real verb conditions, but not the nonce one, and there was some unique activity: left calcarine cortex, left middle temporal gyrus and left fusiform gyrus were only active in the irregular paradigm; left motor strip, left cuneus and left superior occipital gyrus were only active in the nonce condition; and no unique brain activity was found for the regular paradigm.

Jaeger et al. concluded that regular, irregular and nonce past tense inflection produced different (but overlapping) patterns of neural activation, with irregular and nonce verbs eliciting more widespread activity than regular verbs.

So, what are the functions of those parts of the brain which were shown to be active in the three inflectional paradigms?

As has already been pointed out, Broca's area (and neighbouring cortex) was active across all three past tense conditions. It is thought that this part of cortex is involved in access of lexical and morphosyntactic information for grammatical processing to occur (*e.g.* Wise et al., 1991; Mazoyer et al., 1993). Patients with Broca's aphasia have been found to have a deficit in access to morphosyntactic information (Schriefers, Friederici & Van Kampen, 1994) and in production of regular past tense forms (Ullman et al., 1994).

Left hemisphere superior frontal gyrus was found to be active in the irregular and nonce conditions when the data were subtracted from the relevant reading tasks (see Table 6.2). This area seems to inhibit default, high-frequency responses to stimuli, and to be active in novel verb linguistic tasks (Raichle et al., 1994). The specific

meaning and implications of this will be dealt with in detail below.

Jaeger et al. suggested that left dorsolateral prefrontal cortex (active in the 'regular minus read verb' and 'nonce minus read nonce' conditions) may be involved in the computation of the regular suffix, hence its activation only in these two tasks.

Left middle temporal gyrus was active only in the 'irregular past minus read verb' condition; this area is thought to be associated with long-term semantic memory (Squire & Knowlton, 1995).

Inferior parietal lobule was involved in the two 'real word' conditions, and only in the inflection tasks rather than the simple reading tasks: the authors maintained that this area may be holding the phonological form of the stem in memory during the computation of its past tense, but only for existing lexical items.

Finally, occipital areas were substantially active across the three basic conditions, but much less so after subtraction from the reading tasks, as one would expect.

Jaeger et al. interpreted their results as being compatible with dual-route (but not single-route) theories, since "there are major differences in both the location and the amount of brain activation in the regular vs. irregular tasks" (p. 484).

6.0.2 Criticisms of Jaeger's Paper

Seidenberg & Hoeffner (1998) performed a critical analysis of Jaeger et al.'s work and drew three main conclusions:

1. The data did *not* provide evidence that regulars and irregulars are processed in separate brain regions;

2. The experiment was 'contaminated' by methodological problems; and

3. The observed dissociations are in fact not incompatible with connectionist theory, contrary to the conclusions of Jaeger et al.

These criticisms will be treated in turn.

Lack of Evidence for Separate Loci of Processing

As described in section 6.0.1 above, Jaeger et al. interpreted their data as providing crucial support for the dual-mechanism theory. In particular, part of left dorsolateral prefrontal cortex was active for regular and nonce stimuli, but not irregulars, and was said to be involved in generation by rule. Left middle temporal gyrus, conversely, was active in the irregular past condition but not in the regular past or nonce past conditions, and was thought to be associated with the retrieval of known irregular forms. Finally, left lateral orbitofrontal cortex was active for irregular and nonce verbs but not for regulars, and was said to inhibit default high-frequency responses.

Dual-route theory suggests that regular verbs and nonce verbs should show patterns of activity similar to each other and different from irregular verbs, and Jaeger et al. argued that this was found in their data. However, Seidenberg & Hoeffner pointed out that, by the same logic, one would not expect areas to be activated for both regular and irregular verbs, or for both irregular and nonce verbs, whereas the data also reflect *these* findings. In fact, all combinations of conditions activated unique brain regions, as did each type of stimulus alone. Seidenberg & Hoeffner argued that:

> The fact that some areas were activated for both irregulars and nonce items or for regulars and irregulars could just as well be taken as decisive evidence against the [dual-route] theory. In fact there are more areas activated in common by the irregulars and nonce words, which are supposed to be processed by separate mechanisms, than the regulars and nonce words, which are supposed to be generated by the rule.

They further argued that Jaeger et al.'s association of certain functions with certain brain regions was 'highly questionable' and not 'internally consistent'. To select one example: Jaeger et al. found that left superior frontal gyrus was active in the irregular and nonce conditions, and claimed that the area inhibits "default, high-frequency responses, and is inactive during overlearned, practised behaviors" (page 482). They assumed that regularisation was an 'overlearned, automatic response' (page 487) which was suppressed during irregularisation. However, this argument contradicts their explanation of why left dorsolateral prefrontal cortex was active for regular and nonce verbs, namely that it is involved in "on-line formulation of intentional novel behaviors" (page 488). Regularisation may be overlearnt and automatic, or intentional and novel, but not both. Furthermore, if left superior frontal gyrus *does* inhibit regularisation, as Jaeger et al. supposed, it should not have been active in the nonce condition at all, where the task was to apply *regular* inflections to novel stimuli. Jaeger et al. argued that its activation in the nonce condition resulted from inhibition, not of the regularisation process, but of "any number of possible irrelevant responses" to the novel stimuli (page 487). As Seidenberg & Hoeffner pointed out, this area is therefore assumed to serve different functions in the processing of irregular and nonce stimuli.

In summary, when regular and nonce verbs activated a single region, this was used as evidence for a single underlying mechanism; however, when irregular and nonce verbs activated one region, the data were interpreted in terms of *two* underlying mechanisms.

Methodological Problems

Seidenberg & Hoeffner also provided several criticisms of Jaeger et al.'s experimental design, arguing that methodological flaws had introduced factors which invalidated the data.

Their primary objection was the presentation of the stimuli in homogeneous blocks; that is, all regular stimuli were presented in one block, all irregulars in another and all nonce stimuli in a third. This was necessary owing to the poor temporal resolution of positron emission tomography, but in practice it meant that subjects knew what sort of stimulus was about to be presented. The stimuli were therefore susceptible to priming effects: similar stimuli facilitate processing, while dissimilar stimuli create interference. Seidenberg & Hoeffner pointed out that the high level of similarity between stimuli within the regular and nonce conditions would have facilitated performance, whereas the lack of similarity from trial to trial in the irregular condition would have created interference. Since PET studies are known to reflect priming effects in brain activity (*e.g.* Gabrieli et al., 1996; Demb et al., 1995), Jaeger et al.'s data may result from priming rather than from differences between the conditions *per se*.

Furthermore, in the regular and nonce tasks, subjects knew (before stimulus onset) that they merely had to concatenate the stem form with the '-ed' suffix, whereas in the irregular task stimuli would have to be inflected differently on each trial, requiring a greater depth of processing (Craik & Lockhart, 1972) than the other conditions. It has been shown that depth of processing affects brain activity (*e.g.* Gabrieli et al., 1996), and could therefore be responsible for some of the differences found in the Jaeger et al. study. While this confound is eliminated when stimuli are presented in heterogeneous blocks (as the subject does not then know what to expect), this is not possible within a PET methodology. Seidenberg & Hoeffner therefore argued that PET studies are "an inappropriate method for investigating the issues at hand".

They provided evidence for the problems introduced by using homogeneous blocks, by comparing Jaeger et al.'s latency data with data from an experiment conducted by Seidenberg & Bruck (1990). This experiment differed in design, in that regular and irregular stimuli were intermixed rather than presented in blocks of similar verbs. Seidenberg & Bruck's data showed mean net generation latencies[3] of 350 msec. for regulars and 455 msec. for irregulars. In Jaeger et al.'s data, equivalent figures were 38 msec. and 238 msec. respectively. In other words, Jaeger et al.'s subjects produced both types of verb quickly, but the regulars were generated nearly an order of magnitude faster than those in the Seidenberg & Bruck study. Subjects therefore seemed to be taking advantage of the fact that they knew in advance that the stimuli were to be regularised. There seems to be no other explanation, given that it has been shown

[3] *I.e.* the time taken to generate a past tense, minus the time taken to read a verb aloud.

elsewhere that subjects take approximately 100 msec. simply to identify a word, and that fixation points between saccades in reading are usually around 200–250 msec. (Rayner & Pollatsek, 1994).

Seidenberg & Hoeffner provided an alternative interpretation of Jaeger et al.'s data, since each of the three conditions (regular, irregular and nonce) differed in *difficulty*, arguing that the regular task was the easiest, then the nonce task, and that the irregular task was the hardest. They argued that these differences may have been reflected in the levels of brain activity found (as suggested by Just, Carpenter, Keller, Eddy & Thulborn, 1996) and in the latency data. They concluded that:

> The data are therefore consistent with the single-mechanism theory in which performance is determined by the degree of difficulty in the mapping between present and past tense, not type of verb. (Page 116)

This 'differential difficulty' interpretation does not, however, account for the fact that left dorsolateral prefrontal cortex was active in the regular and nonce (but not irregular) conditions. Seidenberg & Hoeffner suggested that this pattern of activation resulted from the facts that:

1. Subjects knew in advance that the regular suffix was required;
2. This information must have been maintained in working memory;
3. Left dorsolateral prefrontal cortex is implicated in working memory (Goldman-Rakic, 1988).

They therefore attributed the activity of left dorsolateral prefrontal cortex in the

regular and nonce conditions to this working memory aspect of the tasks, which was—they argued—not required in the irregular condition.

Do Jaeger et al.'s Data Implicate Separate Subsystems?

The third and final thrust of Seidenberg & Hoeffner's main argument addressed the issue of whether Jaeger et al.'s findings—even if they *were* valid—would in fact be compatible with a dual-route approach to verb inflection, but not with single-route theories, as the original authors concluded. They pointed out that it is easy to explain differential processing of regular and irregular stimuli in terms of two processing routes, "simply by investing each mechanism with whatever characteristics the data demand" (page 119). However, they argued that the converse is *not* true, in that connectionist systems are also capable of showing differential processing of different types of verb. In particular, double-dissociation data do not necessarily implicate separate mechanisms (Plaut, McClelland, Seidenberg & Patterson, 1996; Plaut, 1997), and overall propensity for regularisation also takes place in connectionist models, due to "the aggregate effects of exposure to all of the patterns in the corpus" (Daugherty & Seidenberg, 1992).

6.0.3 Other Recent Studies

As has already been mentioned, PET methods have the advantage of high-level spatial resolution, but the weakness of poor temporal resolution. By using 'event-related potential' (ERP) methodology, the reverse is true.

By studying electroencephalogram (EEG) recordings of brain activity, it is possi-

ble to isolate small changes in voltage which are temporally related to events, such as the presentation of a visual or auditory stimulus, or the initiation of a motor action. These time-locked voltage changes are known as event-related potentials. Spatial resolution is poor for several reasons, including the fact that each electrode attached to the skull can only detect voltage for a relatively large area (between 56 and 64 electrodes are usually used to cover the whole scalp); furthermore, the electroresistance of the skull is very high, and the voltages being detected are very small. In spite of this spatial inaccuracy, however, the method has the great advantage of being extremely temporally accurate, being a measure of 'real-time' neural activation, and is thus not prone to the methodological problems faced by positron emission tomography. By converting the analogue EEG to digital data for analysis at 250 Hz (*i.e.* by recording 250 data points per second from the EEG), neural activity can be analysed to the nearest 4 msec. (being $\frac{1}{250}$ of a second).

Early ERP experiments sought to discover differences in activation between semantic and syntactic processing. Kutas & Hillyard (1983) found that semantic violations led to a pattern of activation called the N400: a negative deflection with a peak amplitude occurring c. 400msec. after stimulus onset with typically wide distribution, maximal over right hemisphere and posterior sites; furthermore, this pattern did not appear to be correlated with other types of linguistic violation (*e.g.* agreement violations). Since then, the N400 has also been found to be elicited in response to nonsemantic variables such as frequency and sentence position (Kutas & Van Petten, 1994; Kutas & Kluender, 1994).

Similarly, some types of ERP effect appear to be elicited in response to syntactic, rather than semantic, stimuli. Work in this area has yielded data from which it is difficult to distil a coherent pattern of elicitation. Particular types of syntactic material have been shown to elicit two response patterns, generally referred to as 'left anterior negativity' (or LAN) and late positivity (also known as the 'syntactic positive shift' (SPS) or P600). Two points should be noted here, however. First, the experiments which have shown such patterns of activation have tended to use syntactic *violations* rather than syntactically correct stimuli, although the general opinion in the literature seems to be that correct syntactic stimuli may elicit similar—if smaller—effects. Secondly, LAN and SPS/P600 patterns may not be syntax-specific. In a recent review paper, Wicha (1998) pointed out that such effects have been elicited by semantic information (Shao, 1995), orthographic violations (Münte, Heinze, Matzke, Wieringa & Johannes, 1998) and even nonlinguistic stimuli (Patel, Gibson, Ratner, Besson & Holcomb, 1996). Conflicting schools of investigation have found that the LAN and SPS/P600 are specific to language (*e.g.* Friederici, Hahne & Von Cramon, 1998) or are domain-nonspecific activity patterns (*e.g.* Coulson, King & Kutas, 1998).

Wicha (1998) provided a comprehensive review of the literature in this area, and summarised the findings as follows:

- LAN found in response to:
  - phrase structure violations: syntactic word category (Friederici, Pfeifer & Hahne, 1993);

- specificity constraint violations (Neville, Nicol, Barss, Forster & Garrett, 1991);
- subcategorisation violations (Rösler, Pütz, Friederici & Hahne, 1993);
- processing grammatical object filler-gap constructions (Kluender & Kutas, 1993; King & Kutas, 1995; Müller, King & Kutas, 1997);
- grammatical violations in word-pair associations: pronouns (Münte et al., 1998);

- SPS/P600 found in response to:
  - subject-verb agreement violations (*e.g.* Kutas & Hillyard, 1983; Hagoort, Brown & Groothusen, 1993);
  - phrase structure violations: word order (Hagoort et al., 1993);
  - subcategorisation violations (*e.g.* Osterhout & Holcomb, 1992; Ainsworth-Darnell, Shulman & Boland, 1998);
  - movement violations: subjacency constraint and empty category (Neville et al., 1991; McKinnon & Osterhout, 1997);

- Biphasic LAN/P600 found in response to:
  - phrase structure violations such as word order and complement structure violations (Neville et al., 1991; Osterhout & Holcomb, 1992);
  - semantic violations such as violations of hyponomy and negative polarity (Shao, 1995);

- Biphasic N400/P600 found in response to:
  - verb tense violations (Friederici et al., 1993);
  - compound syntactic-semantic violations (Ainsworth-Darnell et al., 1998);
  - semantic improbability violations (Shao, 1995).

(List taken from Wicha (1998)).

She concluded from these experiments that the lack of clarity in these findings suggests a more general basis for the elicitation of the LAN and SPS/P600, probably in terms of 'processing costs or general cognitive demands'; in particular, she argued that "the LAN may be an indicator of working memory, while the P600 may be a measure of oddball effects similar to the P300" (p.35). Clearly, much more experimentation is needed before it will be possible to draw firm conclusions about the nature and function of the LAN and SPS/P600.

Most of the ERP literature in this area consists of experiments into violations of some sort; across the literature, very few papers involved morphological studies, and almost all of these required the subject merely to read stimulus sentences. Breyer (1997), on the other hand, performed one of the first investigations into ERP features of past tense *production*, and concluded "that production of regular and irregular past tense forms in English activates brain regions differentially". She presented 12 subjects with 118 stimuli, each consisting of the infinitive form of an English verb, which they were required to inflect to the past tense. Half of the stimuli had regular past tense inflections, and half irregular. Stimuli were controlled by matching for

length, visual distinctiveness, orthographic neighbourhood, spelling regularity and frequency. The stimuli were presented individually and foveally (in random order) on a computer monitor for 300 msec., and subjects were asked to read and inflect the verb silently. After a 1500 msec. pause, a rectangular bar was displayed on the screen, cuing subjects to state aloud the inflected form of the stimulus. The interval of 1500 msec. allowed for the collection of EEG data free from articulation-related artefact.

Breyer found significant differences in terms of representative waveform and voltage topography between irregular and regular verbs. The waveforms produced by different verb types diverged c. 200 msec. after stimulus onset: these differences were significant between 280-420 msec. over central midline and lateral electrode sites. At peak latency (around 350 msec.), regular verbs showed a prominent central midline negativity in comparison with irregulars. Furthermore, although waveform distribution was similar for both verb types, there were significant differences in amplitude between 500 and 600 msec. Irregular verbs produced central midline, anterior and right-lateralised positivity, whilst regulars showed less positivity in general, and bilateral occipital negativity in particular.

Breyer remarked that "[t]he findings suggest differences in the contribution of areas and/or mechanisms engaged by regular versus irregular verb past-tense production"; however, she also pointed out that "[d]ifferences may reflect the same area participating in different ways, or participation of separate brain-regions for past-tense generation".

Although ERPs are considerably less spatially accurate than PET scans, this study nonetheless produced extremely interesting results, from which Breyer concluded:

> While this finding would support an hypothesised distinction between systems for processing regular and irregular past-tense forms, and/or implicate differential weighting of brain regions within a common system, it would not comment on the nature of the mechanism (*e.g.* symbolic in nature rather than associative). (Page 84)

If one compares two recent papers (Jaeger et al., 1996; Ullman et al., 1997), one finds that both claimed to have found support for the dual-route theory, but the findings conflicted with each other with respect to localisation of function. For instance, Ullman et al. concluded that frontal lobes and basal ganglia are responsible for regularisation, and that temporal parietal cortex is responsible for irregular inflection. However, Jaeger et al. found frontal lobe activity in both types of processing, and no activity in basal ganglia for either type.

It may be the case, even if one could find areas of brain activity which were robustly dissociable for regular and irregular verbs, that such dissociations can be modelled using single-system connectionist techniques. Furthermore, the double-dissociations found for regulars and irregulars in impaired patients (Marslen-Wilson & Tyler, 1997, 1998) may be modellable by different types of selective damage to a single-route network.

6.1 Experiment One — Real Verbs

Two ERP experiments (one using real verbs, and one using novel verbs) will be described in the remainder of this chapter. These studies sought:

1. to replicate Breyer's results for regular and irregular verbs, using stimuli chosen by different criteria, and with an improved presentation procedure;

2. to discover whether neighbourhood density has an impact on brain activity;

3. to discover the patterns of brain activation produced when novel verbs are inflected;

4. to attempt to establish whether regularised and irregularised novel verb forms are similar or different in terms of neural activation, and to what extent their ERPs match those of regularised and irregularised real verb forms.

The reason for replicating Breyer's experiments using different stimuli was that the new stimuli were balanced for the number of regulars and irregulars in their rime neighbourhoods. It was therefore possible to perform data analyses of subsets of the stimulus set, to determine the impact of neighbourhood density (in terms of both the number of regulars and the number of irregulars in the stimulus neighbourhoods) on the levels of neural activation in past tense inflectional processing. Furthermore, more stimuli were used in each of these experiments than in Breyer's study and, most importantly of all, *novel* stimuli as well as real verbs were investigated.

6.1.1 Method

Subjects

20 paid subjects were used in the experiment, 10 male and 10 female. All were native English-speakers, with normal or corrected-to-normal vision, and all were right-handed[4]. The mean age of the subjects was 24.6 years (range = 16.1–45.9).

Stimuli

A total of 160 monosyllabic verb stem forms were used as stimuli in this experiment, 80 irregular and 80 regular.

The verbs were categorised according to the number of irregular and regular verbs which share the same rime as each stimulus, that is to say, the number of regular and irregular *neighbours* of a stimulus. For example, a stimulus which shares its rime with an above-median number of regular verbs would be said to be a 'High Regular' (or HighR) verb, a stimulus sharing its rime with a below-median number of irregular verbs would be classed as 'Low Irregular' (LowI), and so on (as described in detail in chapter 4). Irregular medians were determined on the basis of all the phonological clusters which contain at least one irregular verb (N=87), whilst regular medians were calculated in terms of all neighbourhoods which contain at least one regular verb (N=388). Any verb can therefore be classified as:

- Irregular or regular; and within that classification, as either:
  - HighI or LowI (Irregular neighbourhood density); and as either

[4]Verified using the modified Edinburgh handedness inventory (Oldfield, 1971).

Density Type		Irregular Stimuli		Regular Stimuli	
Friends	Enemies	Density	Freq.	Density	Freq.
High	Low	4.35-1.25/6.40-1.70-7.95	97.90	0.00-0.00-17.70	96.65
High	High	3.45-1.55/5.75-2.55-43.70	98.15	2.35-4.80-39.35	98.70
Low	Low	1.00-0.25/1.80-0.30-8.75	134.70	0.00-0.00-2.35	107.85
Low	High	1.00-1.20/1.95-2.25-57.55	176.85	3.00-5.10-3.85	43.10

Table 6.3: Means of stimulus types used in experiment 1

– HighR or LowR (Regular neighbourhood density).

Within each of the two paradigms (*i.e.* regular and irregular verbs), the four possible categories are thus HighIHighR, HighILowR, LowILowR and LowIHighR.

20 verbs were selected per category in each paradigm, making a total of 160 stimuli. These are shown in Tables E.1–E.4 in appendix E; a summary table of means is provided here (Table 6.3). In these Tables, each stimulus is accompanied by a numerical evaluation of its neighbourhood structure (in the 'Density' column), and a measure of Kučera-Francis (KF) token frequency (in the 'frequency' column).

The neighbourhood structure of irregular verbs is shown using 5 numbers, while regular neighbourhood structure is described using 3 numbers, as explained in detail in chapter 4. For each irregular verb, the 5 numbers are a count of:

- the number of simple irregulars with the same rime as the stimulus, and which are inflected by the same irregular subparadigm (Isα);
- the number of shared-rime simple irregulars which are inflected by a different irregular subparadigm (Idα);
- the number of simple and compound irregulars with the same rime as the stim-

ulus, and which are inflected by the same subparadigm (Isβ);

- the number of shared-rime simple and compound irregulars which are inflected by a different irregular subparadigm (Idβ); and

- the number of regulars with the same rime as the stimulus (R).

Thus, taking the verb 'to swing' by way of example, Table 3 shows that its neighbourhood cluster structure is 7-4/9-4-4. So there are 7 simple irregular verbs which follow the /ɪŋ/↦/ʌŋ/ subparadigm (*viz.* cling, fling, sling, sting, string, swing, wring); 4 simple irregulars which are inflected differently (bring, ring, sing, spring); 9 simple *and compound* irregulars which share the inflectional paradigm (*i.e.* the 7 above-named verbs plus hamstring and unstring); 4 simple *and compound* irregulars which are inflected differently (*i.e.* the 4 monosyllabic verbs named above plus no compounds); and finally 4 regular shared-rime verbs (ping, ring, ting, wing). Since the verb 'swing' shares its rime with more simple and compound irregular neighbours than an hypothetical 'average' (more specifically, 'median') irregular verb, and with fewer regular neighbours, it is said to be a HighILowR stimulus[5].

For the regular stimuli, the classification is simpler: three numbers are used, which refer to:

- the number of simple shared-rime irregulars (Iα);

- the number of simple and compound shared-rime irregulars (Iβ);

[5]Note that, being an irregular verb, it could also be referred to as a 'High Friend, Low Enemy' stimulus.

- the number of shared-rime regulars (R).

So, for instance, the regular verb 'to blind' has a neighbourhood structure of 4-8-4, meaning that there are four simple irregulars ending in the rime /aɪnd/ (*viz.* bind, find, grind, wind), or 8 if compounds are included (the above 4 plus rebind, spellbind, unbind, unwind), and 4 regular /aɪnd/ verbs (blind, mastermind, mind, remind)[6].

The median neighbourhood structure is, for irregular verbs, 1-1/3-1-17, and, for regulars, 0-0-6. Any verb can therefore be classified as HighI or LowI, and HighR or LowR, by a comparison of its neighbourhood structure with the medians for its class.

The stimuli are paired up for the sake of constraining external factors (such as frequency) which may confound the results of the experiment. Each stimulus is matched—as far as is possible within the constraints of English phonological neighbourhood densities—with one from the opposite class in terms of Kučera-Francis frequency, and these are thus shown side by side in the Tables. In this case, HighILowR irregulars will be matched with LowIHighR regulars, HighIHighR irregulars with HighIHighR regulars, etc. In other words, High 'Friend' Low 'Enemy' stimuli are matched with each other, High 'Friend' High 'Enemy' stimuli are matched, and so on. It has been possible to do this consistently in all cases except when matching LowIHighR irregulars with HighILowR regular stimuli (Table E.4), since there are only about 20 regular verbs in English which fit into this category: there is thus no

[6]It should be noted that for an irregular stimulus another irregular in its neighbourhood may be a 'friend' or an 'enemy' in terms of its inflectional paradigm, whereas for a regular stimulus, all irregulars are *de facto* enemies, hence there being only 3 numbers for regulars as opposed to 5 for irregulars.

opportunity for precise matching of KF frequencies.

Furthermore, means of the figures for familiarity and Nvalue—as defined in the Oxford Psycholinguistic Database (Quinlan, 1992)—were calculated for the 80 regulars and 80 irregulars, and these overall means were compared using two 2-way ANOVAs (with regularity/irregularity and either familiarity or Nvalue as between-items factors) to ensure that none of the categories could skew the data by one or more of these factors, and no effects of familiarity or Nvalue were found (Familiarity: $F(1,104) = 2.296, p = 0.133$; Nvalue: $F(1,158) = 0.562, p = 0.454$). The full data are shown in appendices G and H.

Procedure

Subjects were seated in an armchair in a semi-darkened soundproof booth, one metre away from a 20-inch computer monitor, which was positioned in such a way that the centre of the screen was approximately at the subject's eye-level. A keyboard on the subject's lap enabled him/her to regulate the rate of the presentation of stimuli.

Stimuli were presented using a simple TurboPascal script (see appendix I).

When the subject pressed a key, a centrally-oriented fixation point (a small +) appeared on the monitor for a randomly-determined period of time between 1200 and 1700 msec. The element of randomness was incorporated so as to reduce ERP artefact resulting from timelocked expectancy-related potentials.

The fixation point was then replaced by a randomly-selected experimental item (preceded by the infinitive marker 'to') in the centre of the screen. Since subjects

had been asked to foveate on the + sign, the presentation of a stimulus would not result in eye movement, which would distort the ERP recording. At this point the subject mentally inflected the stem to its past tense form. After 2000 msec. the stimulus disappeared, which was the cue for the subject to say aloud the inflected form. In this way, there was a period where the subject was inflecting the verb, but not articulating it, to minimise—and preferably eliminate—articulation-related artefact[7]. The entire presentation of a single stimulus therefore took between 3200 and 3700 msec., which was short enough for the subject to have little difficulty not blinking. As each presentation was self-initiated, the subject was free to blink between trials. All visual presentation was in white on an otherwise black screen, and stimuli subtended between 1.2° and 2.4° of horizontal visual angle.

Stimuli were presented in two blocks of 80, with a short break between the blocks. Each block was preceded by two practice presentations to give the subject an opportunity to get used to the task and to adapt to regulating his/her blinking; this also allowed time for the potentials to settle as the subject became accustomed to the experimental environment.

Subject responses were conveyed by intercom to the experimenter outside the booth, and were recorded as being correctly or incorrectly inflected.

[7]Breyer's presentation system involved showing the stimulus, removing it, and then showing a blue bar as a cue for response; however, the disappearance of the stimulus tended to create a visual-evoked potential which conflicted with the potentials evoked by the stimulus.

ERP Recording

Subjects were fitted with a 56-channel electrode cap, from which continuous electroencephalogram (EEG) was recorded (see Figure 6.1 for the montage). At the point of contact for each electrode, the scalp was prepared using a gently abrasive cleaning liquid (to ensure as good a contact as possible), and each electrode was filled with electroconductive gel. Electrodes were also attached (using tape) to the outer canthi of each eye, and to the upper and lower canthi of the right eye, so that horizontal and vertical eye movements (including blinks) could be monitored. Reference electrodes were also attached to the left and right mastoid[8]. Ground was situated on the midline, midway between sites FPZ and FZ.

Electrode impedance was reduced until it was less than 5KΩ at each site, and the EEG signal was amplified with a gain of 20,000 and filtered with a bandpass of 0.1–100 Hz.

EEG was recorded using NeuroScan and converted from analogue to digital data at 250 Hz (*i.e.* the acquisition rate is therefore 1 data point per site per 4 msec.). Whenever an item was presented, a code marking the event was superimposed on the computer file saving the continuous EEG; this code indicated both the verb type (*i.e.* regular or irregular) and the neighbourhood density of the stimulus.

[8]The active reference electrode was on the right mastoid.

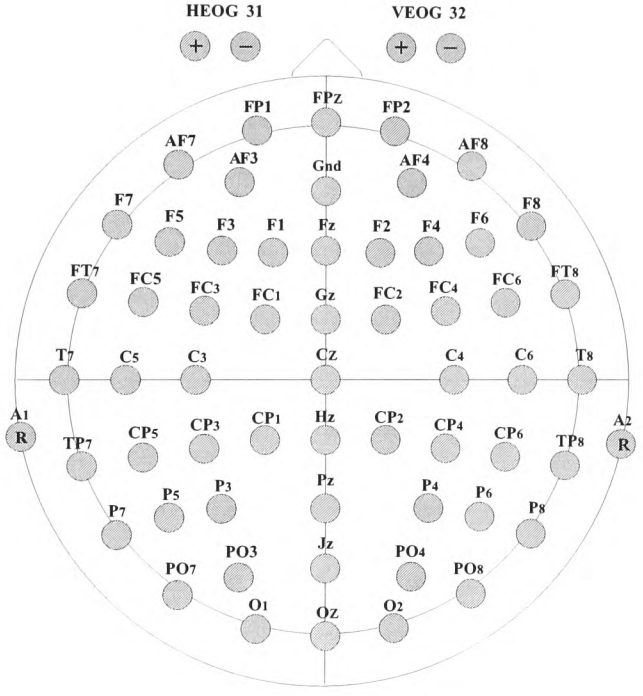


Figure 6.1: 56-channel electrode montage used in ERP experiments

6.1.2 Results

3200 trials were recorded, 160 for each of 20 subjects. This generated a substantial amount of data: in total, 91,750,400 discrete measurements were collected and saved in this experiment[9].

310 trials were rejected as a result of eye movement (9.69% of total trials; mean per subject=15.5, range: 0–50). 165 trials were rejected through incorrect responses across all subjects (5.16% of total trials; mean per subject=8.25, range: 1–31). The total number of trials rejected (through artefact and subject errors) across all subjects was therefore 475 (mean per subject=23.55; range: 2–55), or 14.84%.

The higher-than-expected number of incorrect responses seemed to be the product of the repetitive nature of the experiment, and the fact that the stimuli were presented out of sentential context: people are seldom aware of the linguistic functions they may be enacting when producing spoken verb inflections, and it is therefore understandable that subjects sometimes reported becoming confused by the stimuli simply because they were thinking in an unusual and nonnatural way.

Preparation of Data for Analysis

Once all the data had been collected for a subject, the EEG was prepared for analysis. The continuous EEG for each individual subject was re-referenced to a baseline consisting of the mean activity of the two mastoid electrodes, and then divided into epochs, one per stimulus presentation. The recording window for each epoch was

[9] 512 data points for each of 56 sites per trial per subject.

from 200 msec. before stimulus onset to 1848 msec. afterwards (total epoch length was thus 2048 msec.). Using the first 200 msec. of the epoch as a baseline, voltages recorded for each presentation at each electrode site were then normalised so that the mean pre-stimulus activation at each site was marked as 0 microvolts (μV)[10].

Epochs damaged by large eye movements (defined as those which showed a potential of >50μV or <-50μV on the horizontal or vertical eye-channels) were then eliminated, as were all trials on which the subject had produced an incorrect inflection of the stimulus, or had responded too early (or not at all).

Eight computer files were then created, containing mean data for all stimuli within each of the four stimulus classes (HighILowR, HighIHighR, LowILowR and LowIHighR) within each verb type (regular and irregular verbs).

This process was repeated for all 20 subjects, and the resulting 160 data files (each containing up to 20 epochs) were then concatenated in various ways to allow for detailed analysis of the influence of verb type and neighbourhood density on past tense inflection.

A grand-averaged file comparing all irregular with all regular stimuli across all subjects was compiled and, from this, waveform graphs (of microvoltage change against time) were created showing the averaged activity of each individual site for both regular and irregular stimuli. Topographical maps were also created, showing microvoltage change (μV) across cortex in 40 msec. slices for each condition (regulars and

[10]Note, therefore, that all microvoltages referred to hereafter are actually measurements of microvoltage *change* relative to the (baseline) activation before stimulus onset.

irregulars), with a constant μV range between +12μV and -2.0μV. Sample maps and waveforms are shown in Plates 1 and 2 below.

A summed-area analysis was then conducted on the grand-averaged file. This involved grouping the electrode sites into five sets (of 5 sites per hemisphere) by area, as follows:

Frontal: Sites FP1, AF3, AF7, F5, F7 in left hemisphere, FP2, AF4, AF8, F6, F8 in right hemisphere;

Midline: FPZ, FZ, CZ, PZ, OZ;

Lateral: FC5, FT7, T7, CP5, TP7 in left hemisphere, FC6, FT8, T8, CP6, TP8 in right hemisphere;

Midlateral: FC1, FC3, C3, CP1, CP3 in left hemisphere, FC2, FC4, C4, CP2, CP4 in right hemisphere; and

Posterior: P5, P7, PO3, PO7, O1 in left hemisphere, P6, P8, PO4, PO8, O2 in right hemisphere.

Note that, for mathematical simplicity and statistical balance, only 45 of the 56 electrode channels were used in the summed-area analysis.

The data from 0–1000msec. were also separated into 25 timebins, each containing 40msec. of data. 5 2- or 3-factor repeated-measures ANOVAs (1 per area) were then calculated for each of the 25 timebins, in order to ascertain whether differences between potentials for regular and irregular stimuli were statistically significant

across the within-subjects factors of verb type (regular or irregular), site and—where applicable—hemisphere. Where interactions were found in which site was a factor, a standard statistical technique (epsilon (ϵ) adjustment) was applied to the ANOVA—this acted to reduce the primary and secondary degrees of freedom of the interaction, to compensate for the fact that neighbouring sites are not strictly independent variables.

Of the 125 ANOVAs conducted in the grand-averaged analysis, 56 yielded statistically significant verb type-related differences; these will be reported in detail in the next section. It should be noted that, given the large number of ANOVAs conducted, and given a statistical level of significance of 5% (0.05), one would expect around $\frac{1}{20}$ of the significant findings to be anomalous. Therefore, significant effects are generally only reported here when they are supported in one of two ways: either by being reflected in another area during the same timebin; or by being reflected in the same area during the timebin immediately preceding or following the effect.

For example, a 'significant' effect between, say, 200–240msec. in frontal areas will only be reported:

1. if it immediately follows or precedes a similar effect in the same area (respectively an effect from 160–200 msec. or from 240–280 msec., also in frontal areas); or

2. if it parallels a contemporaneous effect in another area (an effect from 200–240 msec. in, say, midlateral areas).

Grand-Averaged Results

In an experiment such as this, the event-related potentials fall into two main types: visual ERPs and language-related ERPs. Visual ERPs reflect the initial registering of the (visually presented) stimulus in occipital cortex and tend to consist of an early positive peak followed by a negative deflection, both taking place around 80–120 msec. after stimulus onset. Language-related ERPs are found after this point (once the stimulus has been visually processed) and may be widespread. Findings corresponding to these two types of potential are described separately.

Visual ERPs Early visual-evoked potentials (VEPs) were somewhat similar across verb types, comprising a positive peak (P_1) maximal at approximately 90 msec. across occipital sites OZ, O1 and O2, followed by a negative deflection (N_1) maximal at approximately 120 msec. at sites PO7 and PO8. A brief laterality effect (*i.e.* effect of hemisphere) was noted in the latency of the P_1 in one ANOVA (timebin 40–80 msec.; $F(1, 19) = 5.908; p = 0.025$), where left hemisphere[11] showed greater positivity than right hemisphere during this timebin ($\bar{x} = 0.258\mu$V in LH; $\bar{x} = 0.014\mu$V in RH)[12]; no other significant differences were found within the first 240 msec. from stimulus onset.

Language-Related ERPs Plates 1 and 2 overleaf show microvoltage topographies and grand-averaged waveforms of ERPs elicited by regular and irregular verbs. These

[11] Hereafter referred to as LH; similarly RH will be used as shorthand for right hemisphere.
[12] Where $\bar{x}$ signifies the mean.

Irregular

Regular

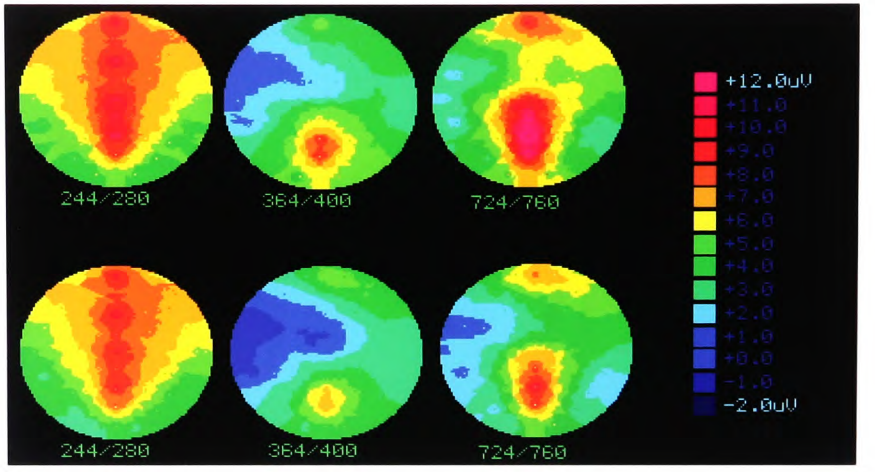


Plate 1: Grand-averaged ERP topographies for real verbs

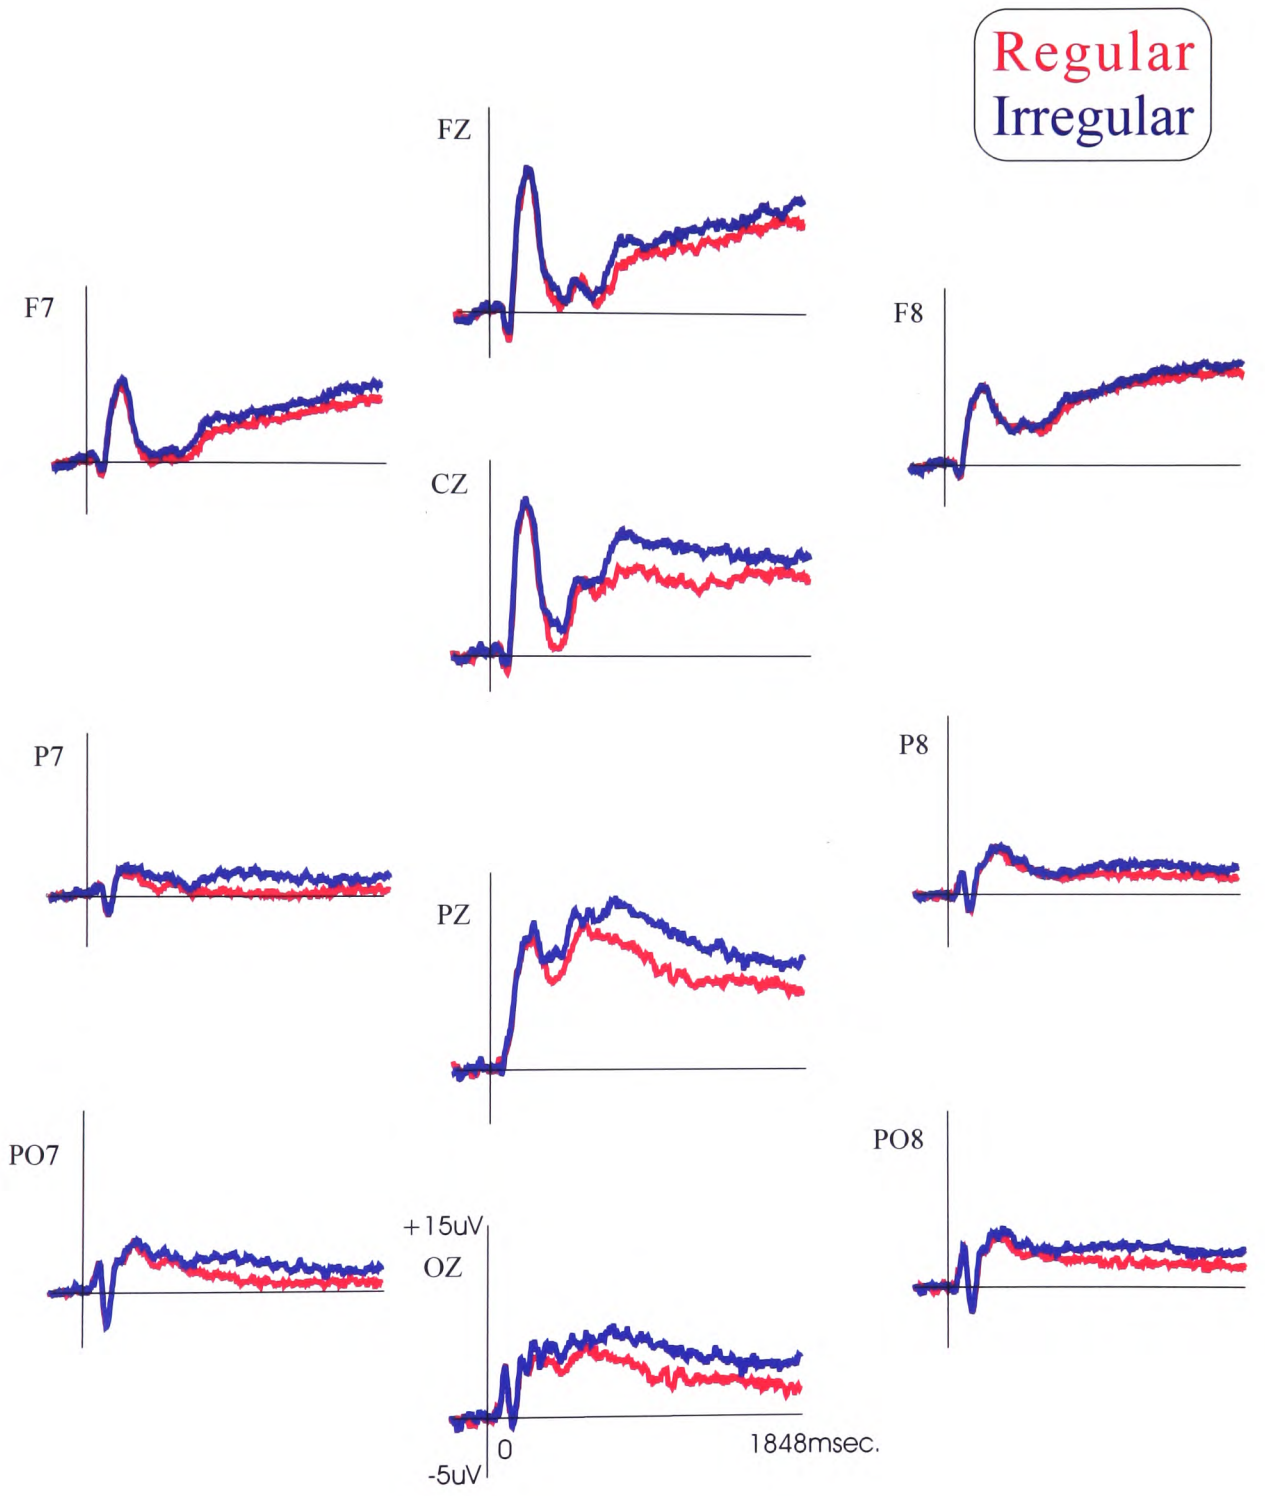


Plate 2: Grand-averaged waveforms for real verbs

topographies and waveforms give a representative impression of the brain activation for each condition, which is useful for understanding the specific effects detailed in this section.

The grand-averaged epoch will be divided into three phases for the purposes of explanation: 240–320msec.; 360–560msec.; and 600–1000msec. Note that no verb type-related differences were found before 240 msec. or between the three phases. Within each phase, significant results will be discussed by area in (as far as is practicable) the following order: midline; frontal; lateral; midlateral; posterior[13].

Phase 1: 240–320 msec.

In this phase, a clear event-related positive peak (hereafter P_2) was found in all but occipital areas, maximal around 280 msec. in frontal areas, diminishing in magnitude towards posterior sites.

A significant difference in amplitude was found between verb types at left hemisphere **lateral** sites at the zenith of the peak (timebin 240–280 msec.; $F(1, 19) = 4.821; p = 0.041$). This verb type × hemisphere interaction showed greater negativity in LH for regular verbs ($\bar{x} = 4.431\mu$V in LH) than for irregulars ($\bar{x} = 5.197\mu$V in LH), but no difference was found in RH. This interaction is depicted in Figure 6.2[14].

Between 320–360 msec., no significant differences were found between waveforms for regular and irregular stimuli.

[13] Area names will be emboldened for ease of reference.

[14] In this and all subsequent interaction plots in this chapter, microvoltage change is plotted on the y axis, and hemisphere on the x axis.

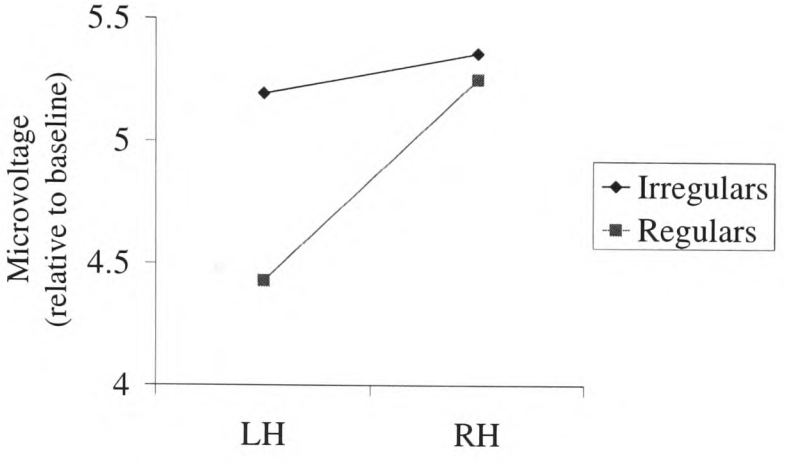


Figure 6.2: Lateral verb type × hemisphere interaction between 240–280msec.

Phase 2: 360–560 msec.

Waveforms diverged again at the bottom of the P_2 (c. 360–480 msec.), as can clearly be seen in Plate 2. Again, a relative positivity for irregular over regular verbs was found; this difference was maximal over the **midline** ($\bar{x}$(regular) = 3.603μV; $\bar{x}$(irregular) = 4.824μV), in particular at sites CZ and PZ (timebin 360–400 msec.; $F(4, 76) = 3.172; p = 0.018$ ($p = 0.047$ after ϵ adjustment)). This difference resulted from a prominent negative deflection for regulars around the end of the language-related peak. Verb type-related differences can be seen across the whole scalp during this phase (between c. 360–480 msec.): a verb type × hemisphere interaction at **frontal** sites (see Figure 6.3) showed relative irregular positivity in LH ($\bar{x} = 1.925\mu$V; cf. 1.575μV for regulars) and relative regular positivity in RH ($\bar{x} = 3.773\mu$V; cf.

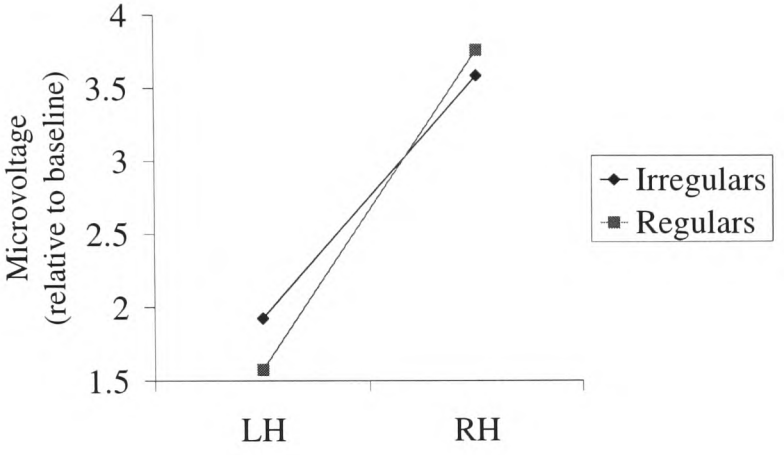


Figure 6.3: Frontal verb type × hemisphere interaction between 360–400msec.

3.598μV for irregulars)(timebin 360–400 msec.; $F(1,19) = 5.083; p = 0.036$). It should be noted that this was the only period during which verb type was found to have any effect on potentials in frontal areas of either hemisphere.

During the same period, significant irregular positivity was also evident in LH **lateral** sites (timebin 360–400 msec.; $F(1,19) = 6.452; p = 0.020$) and LH **midlateral** sites (timebin 360–400 msec.; $F(1,19) = 5.045; p = 0.037$). Finally an effect of verb type was found at **posterior** sites (without interactions), where again irregular stimuli evoked greater positivity ($\bar{x} = 4.298\mu$V; cf. 3.587μV for regulars)(timebin 360–400 msec.; $F(1,19) = 5.946; p = 0.025$). This effect was significant up to 480 msec.

After the second phase, there followed a period in which there was no effect of

verb type on the potentials at any site, between approximately 560–600 msec.

Phase 3: 600–1000 msec.

In the final phase (from 600 msec. onwards) verb type was found to have a strong effect on μV change at all nonfrontal sites. **Midline** sites showed a main effect of verb type from c. 600 msec. with irregular stimuli evoking greater potentials than regular stimuli, maximal at around 720–760 msec. ($\bar{x}$(regular) = 6.318μV; $\bar{x}$(irregular) = 8.693μV; $F(1, 19) = 9.267; p = 0.007$). A verb type × site interaction was found from c. 880 msec., increasing in significance towards the end of the analysed timescale, showing a greater influence of verb type on potentials at central and posterior midline sites (*i.e.* CZ, PZ, OZ) than at frontal midline sites (that is, sites FPZ and FZ).

Lateral sites yielded slightly more complex results in the final phase: a verb type × site interaction between 720–800 msec., maximal at site FC5 ($\bar{x}$(regular) = 2.442μV; $\bar{x}$(irregular) = 3.844μV between 760–800 msec.) developed into a verb type × hemisphere interaction from 800 msec., as shown in Figure 6.4 (maximal between 920–960 msec.; $F(1, 19) = 10.591; p = 0.004$). In this interaction, although irregular stimuli evoked greater potentials than regular stimuli in both hemispheres, and although RH was more active than LH, the difference between potentials in LH ($\bar{x}$(irregular) = 2.839μV; $\bar{x}$(regular) = 1.553μV) was greater than the difference between potentials in RH ($\bar{x}$(irregular) = 4.541μV; $\bar{x}$(regular) = 4.105μV)[15].

All **midlateral** sites showed a strong relative positivity for irregulars over the

[15]Means taken from the period of the greatest statistical significance of the interaction, *viz.* 920–960 msec.

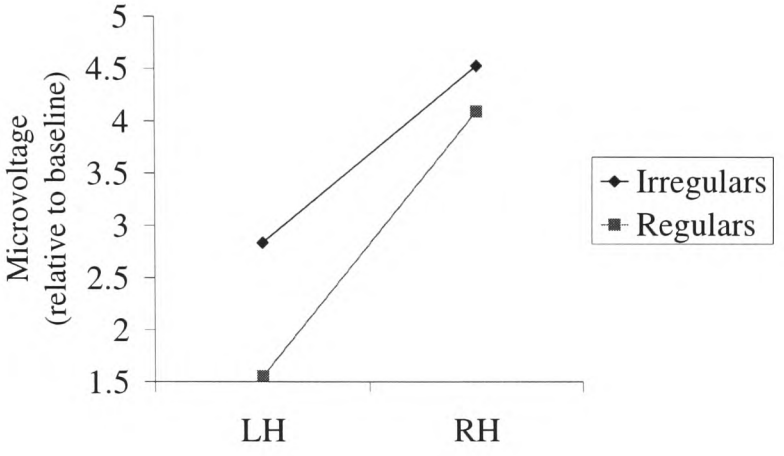


Figure 6.4: Lateral verb type × hemisphere interaction between 920–960msec.

whole of the final phase, and a verb type × hemisphere interaction was found from c. 720 msec. (see Figure 6.5) along similar lines to that found for lateral sites described in the preceding paragraph: in this case, means of 5.295μV (for irregulars) and 3.689μV (for regulars) were recorded in LH, and means of 6.057μV (for irregulars) and 5.148μV (for regulars) recorded in RH[16].

Finally, **posterior** sites also showed a strong effect of verb type on potentials from approximately 720 msec.; this effect increased in significance throughout the final phase ($\bar{x}$(irregular) = 3.803μV; $\bar{x}$(regular) = 2.173μV at 960–1000msec.). Verb type × hemisphere interactions were found between 800–960 msec. (maximal between 920–960 msec.; $F(1, 19) = 24.814; p < 0.001$) which were very similar to those found

[16]Means taken from the period of the greatest statistical significance of the interaction, *viz.* 920–960 msec.

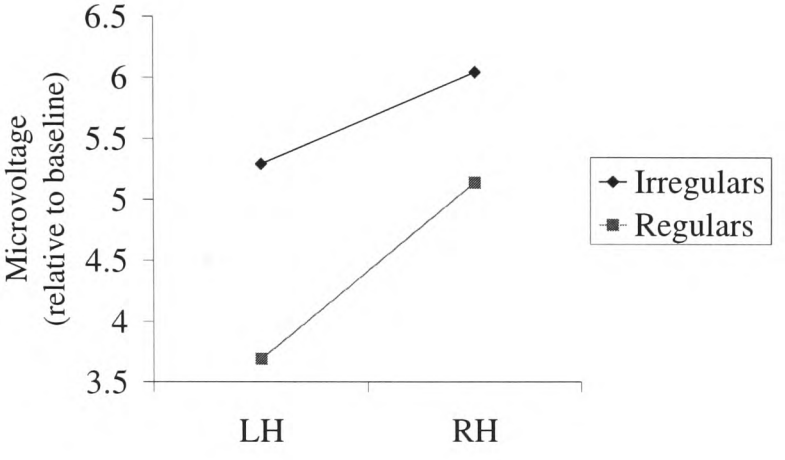


Figure 6.5: Midlateral verb type × hemisphere interaction between 920–960msec.

in lateral and midlateral areas during this phase (see Figure 6.6), with LH means of 3.776μV for irregulars and 1.931μV for regulars, and RH means of 4.045μV for irregulars and 3.151μV for regulars[17].

Normalised Grand-Averaged Results The grand-averaged data for all subjects were then normalised. That is to say, separate files were created for each timebin within each verb type; then, within each timebin, the mean microvoltage value was calculated for each site (across all subjects), the smallest and greatest mean ('*minval*' and '*maxval*') were calculated, and each data point revalued by application of the simple formula:

$$\frac{ACTUAL\mu V - minval}{maxval - minval}$$

[17]Means taken from the period of the greatest significance of the interaction, *viz.* 920–960 msec.

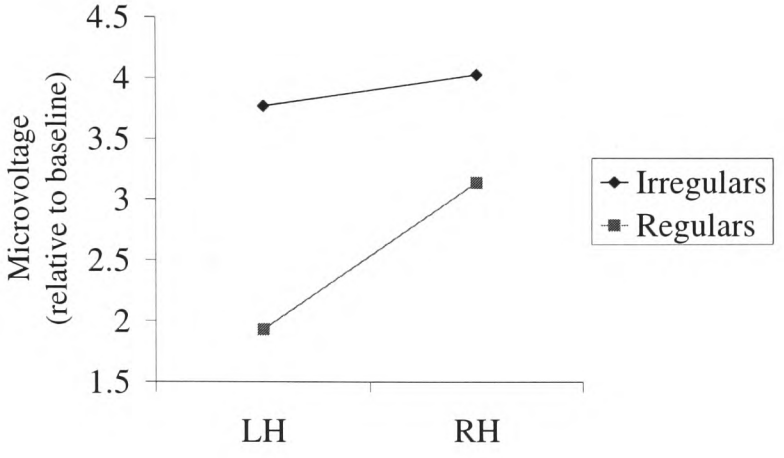


Figure 6.6: Posterior verb type × hemisphere interaction between 920–960msec.

Each individual data point was therefore scaled to a number generally falling between 0 and 1.

By this method of analysis, the amplitude of any particular data point therefore becomes less 'relevant' in comparison with the distribution of the waveforms.

The data from 240–1000msec. were then selected for further analysis (*i.e.* all the data containing verb type-related differences) and 80 further ANOVAs (identical in design to those conducted previously) were then carried out on these normalised data. The results were broadly similar, with one main difference: the period in which verb type seemed to have no significant effect on potential (between phases 2 and 3—described above as 560–600 msec.) was broader in the normalised dataset (approximately 560–680 msec.), suggesting that the phases from 360–560 msec. and

from 680–1000 msec. are indeed distinct from one another.

Subset Analyses

Since the stimuli had been carefully balanced in terms of neighbourhood density, it was possible to carry out statistical analyses on subsets of the data. The datasets for irregular and regular stimuli were separated, and four subset analyses conducted as follows:

1. Irregular stimuli: HighI vs. LowI comparison;
2. Irregular stimuli: HighR vs. LowR comparison;
3. Regular stimuli: HighI vs. LowI comparison;
4. Regular stimuli: HighR vs. LowR comparison.

In other words, these four comparisons asked: does the number of irregular neighbours of a irregular stimulus have an effect on the potentials evoked during inflection (item 1)? What about the number of regular neighbours belonging to irregular stimuli (item 2)? Similarly, are regular stimuli affected by their irregular neighbours (item 3)? Or their regular neighbours (item 4)? Put more formally, the four subset analyses are assessing the impact of Idensity (*i.e.* HighI or LowI) and Rdensity (*i.e.* HighR or LowR) on regular and irregular verb inflection, where Idensity is the irregular neighbourhood structure of a verb, and Rdensity its regular neighbourhood structure.

By performing these analyses, it was possible to determine in detail the impact of neighbourhood density on past tense inflection at a neurophysiological level. It

was not possible to subdivide the stimuli any further (to enable analyses comparing, for instance, HighILowR with HighIHighR stimuli) since there were not enough stimuli used in the experiment to produce statistically reliable analyses at that level of specificity.

One might expect that irregular stimuli would be affected by both their irregular and their regular neighbours, since dual-route and single-route theorists both hypothesise some sort of pattern-associative network for irregulars, in which all neighbours of a stimulus may have some impact on its inflection. More contentious is the prediction for regular stimuli (items 3 and 4 above): symbolic theories suggest that, since a regular verb is inflected by means of a default suffixation rule, its neighbourhood structure (whether in terms of regular or irregular neighbours) should have no effect on that process. Conversely, connectionists argue that, if regulars are inflected by a process similar to that for irregularisation, neighbourhood density *should* have an effect, and this effect may be reflected in language-evoked potentials.

Each of the four subset analyses consisted of 125 ANOVAs carried out (similarly to the grand-averaged analysis) on the data subset divided into 40 msec. timebins. These subset analyses will be treated in turn.

Irregular Stimuli—HighI vs. LowI Neighbourhood Density Analysis 18 of the 125 ANOVAs carried out in this analysis of irregular stimuli were found to show condition-related statistical differences in the waveforms, where the conditions were HighI and LowI neighbourhood density (hereafter Idensity). These differences were

found between 240–320 msec. in midline and midlateral sites, and from approximately 400–1000 msec. in frontal sites. Each of these effects will be described in turn.

Effect 1:

As noted above, the primary event-related positive peak was observed to be maximal around 280 msec. It was during this timeframe that differences were found between HighI irregulars and LowI irregulars in midline and midlateral sites. An Idensity $\times$ site interaction was found at the **midline** (timebin 280–320 msec.; $F(4,76) = 4.778; p = 0.002$ ($p = 0.015$ after ϵ adjustment)). In this interaction, HighI irregulars elicited a more positive peak than LowI irregulars at frontal/central midline sites FPZ, FZ and CZ, maximal at site FZ ($\bar{x}$(HighI) = 7.475μV; $\bar{x}$(LowI) = 5.368μV); this difference was not observed at sites PZ and OZ.

At around the same time, two interactions were found at **midlateral** sites: an Idensity $\times$ hemisphere interaction (timebin 240–280 msec.; $F(1,19) = 4.402; p = 0.050$), followed by an Idensity $\times$ site interaction (timebin 280–320 msec.; $F(4,76) = 4.774; p = 0.002$ ($p = 0.028$ after ϵ adjustment)). In the Idensity $\times$ hemisphere interaction (see Figure 6.7), LH midlateral sites showed greater positivity at the zenith of the peak for HighI stimuli than for LowI stimuli ($\bar{x}$(HighI) = 7.678μV; $\bar{x}$(LowI) = 7.134μV in LH); this difference was not observed to the same extent in RH ($\bar{x}$(HighI) = 7.501μV; $\bar{x}$(LowI) = 7.419μV in RH). In the Idensity $\times$ site interaction, the HighI-evoked potentials dropped slightly earlier from the peak of the P_2 than the LowI potentials at certain midlateral sites in both hemispheres (*viz.* FC1, FC2, FC3

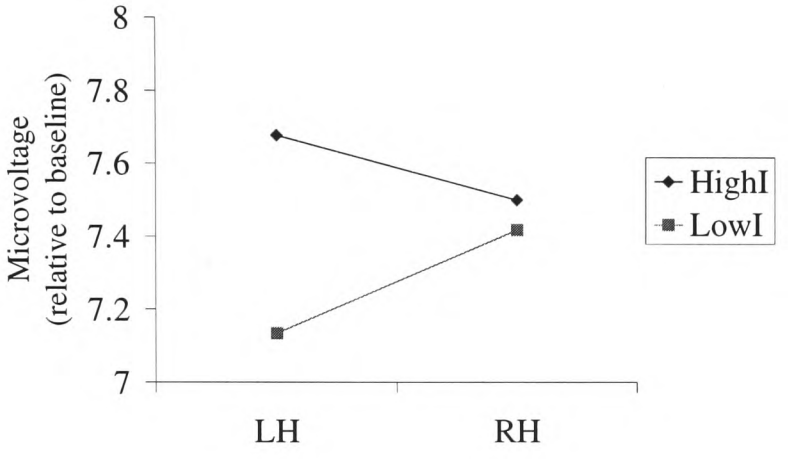


Figure 6.7: HighI vs LowI condition × hemisphere interaction for irregulars

and FC4). This effect was maximal at site FC1 ($\bar{x}$(HighI) = 3.826μV; $\bar{x}$(LowI) = 4.992μV between 280–320 msec.).

Effect 2:

The second effect was a HighI positivity at certain **frontal** sites, resulting in Idensity × site interactions from 400 msec. onwards, maximal at sites FP1 and FP2 (timebin 840–880 msec.; $F(4, 76) = 7.140; p < 0.001$ ($p = 0.001$ after ϵ adjustment); $\bar{x}$(HighI) across both sites = 8.285μV; $\bar{x}$(LowI) across both sites = 6.888μV).

In summary, during the process of irregularisation Idensity was found to have a noticeable effect, in which HighI stimuli evoked greater potentials at frontal/central midline and midlateral sites during the P_2, and at frontal sites later in the epoch.

Irregular Stimuli—HighR vs. LowR Neighbourhood Density Analysis 15 of the 125 ANOVAs conducted in this subset analysis showed significant differences in terms of Rdensity (that is, between HighR and LowR irregular stimuli). These generally fell into one of three groups: an effect in the timeframe 280–320 msec.; an Rdensity × hemisphere interaction at frontal sites between 400–520 msec.; or a post-peak effect between 520–640 msec., spreading from midline and lateral sites to midlateral and posterior areas. Again, these will be discussed in turn.

Effect 1:

At or around the zenith of the P_2 (approximately 280 msec.), effects of Rdensity were found at frontal, midlateral and posterior sites. First, a 3-way Rdensity × hemisphere × site interaction was noted in **frontal** areas (timebin 280–320 msec.; $F(4,76) = 3.128; p = 0.020$ ($p = 0.043$ after ϵ adjustment)). This seemed to result from a momentary relative positivity in HighR potentials at site F7 in LH, contrary to all other frontal sites in that timebin, where LowR potentials were more positive than those for HighR stimuli.

Secondly, **midlateral** sites showed a Rdensity × hemisphere interaction (timebin 280–320 msec.; $F(1,19) = 6.023; p = 0.024$), in which LowR stimuli evoked much greater positivity than HighR stimuli in RH ($\bar{x}$(HighR) = 5.202μV; $\bar{x}$(LowR) = 6.372μV in RH), and only slightly greater positivity in LH ($\bar{x}$(HighR) = 4.654μV; $\bar{x}$(LowR) = 5.288μV in LH). This interaction is depicted in Figure 6.8.

Finally, a pure effect of Rdensity was noted in **posterior** areas (timebin 280–320

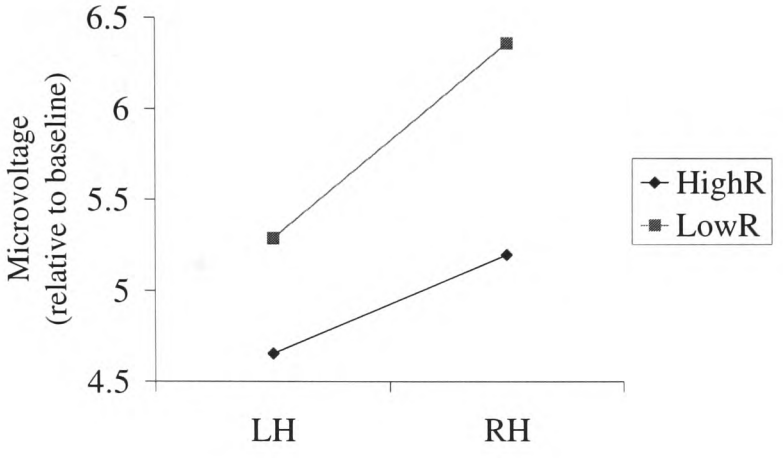


Figure 6.8: Rdensity × hemisphere interaction for irregulars

msec.; $F(1, 19) = 6.706; p = 0.018$), whereby LowR stimuli again elicited greater positivity than HighR stimuli, even though there was not a clearly-defined language-related peak at posterior sites.

Effect 2:

A rather confusing set of Rdensity × hemisphere interactions was found between 400–520 msec. in **frontal** areas (shown in Figures 6.9, 6.10 and 6.11). In all three 40 msec.-long timebins, RH activation was higher than LH activation for both HighR and LowR stimuli. However, between 400–440 msec. (Figure 6.9), both conditions yielded similar activation in LH ($\bar{x}$(HighR) = 1.618μV; $\bar{x}$(LowR) = 1.478μV in LH) but differing activation in RH ($\bar{x}$(HighR) = 2.883μV; $\bar{x}$(LowR) = 3.314μV in RH). Between 440–480 msec. (Figure 6.10) the reverse was true, where HighR positivity was

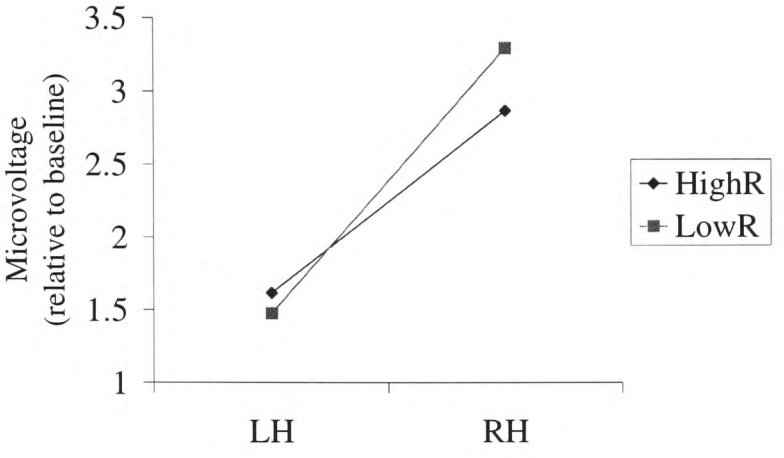


Figure 6.9: Rdensity × hemisphere interaction between 400–440msec.

found in LH only ($\bar{x}$(HighR) = 1.823μV; $\bar{x}$(LowR) = 1.478μV in LH, and $\bar{x}$(HighR) = 3.026μV; $\bar{x}$(LowR) = 3.055μV in RH). Finally, between 480–520 msec. (Figure 6.11), both conditions once again elicited similar activation in LH, and LowR positivity in RH ($\bar{x}$(HighR) = 2.052μV, $\bar{x}$(LowR) = 2.120μV in LH; and $\bar{x}$(HighR) = 3.436μV, $\bar{x}$(LowR) = 4.116μV in RH). The unstable nature of these interactions was a reflection of the fact that at the base of the P_2 (approximately 400–520 msec.) the waveforms of the two conditions were not particularly smooth, and seemed to interweave as they settled after the main burst of activity.

Effect 3:

Immediately after the main peak, as the activity levels out, LowR waveforms briefly settled around 1.5μV higher than HighR waveforms, in **midline** and **lateral**

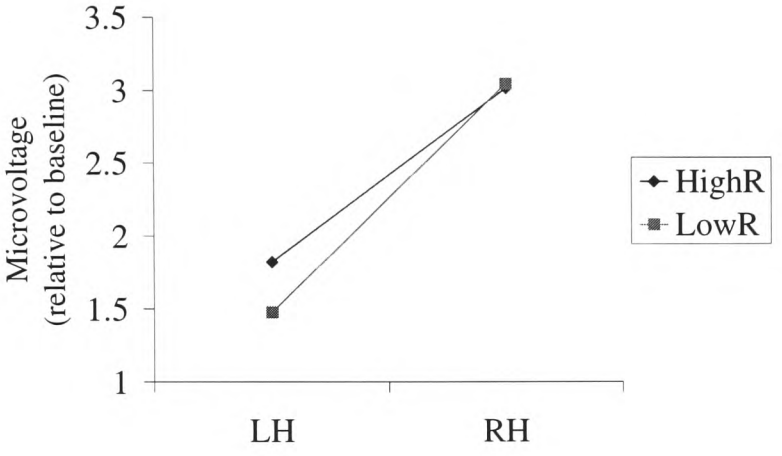


Figure 6.10: Rdensity × hemisphere interaction between 440–480msec.

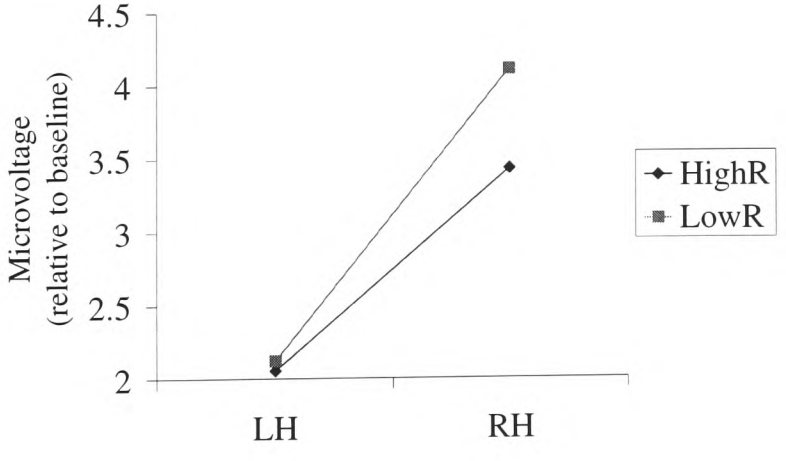


Figure 6.11: Rdensity × hemisphere interaction between 480–520msec.

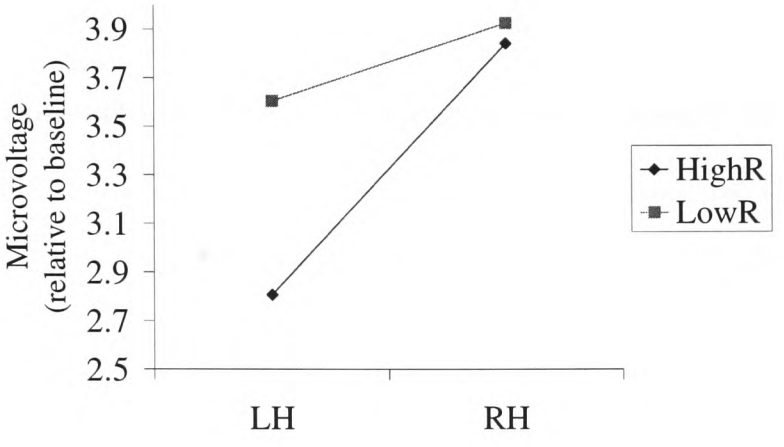


Figure 6.12: Posterior Rdensity × hemisphere interaction between 600–640msec.

areas (timebin 520–560 msec.; $F(1,19) = 4.759; p = 0.042$ on the midline; $F(1,19) = 4.830; p = 0.041$ at lateral sites), before blending together for the remainder of the epoch. This effect was also noticeable in the next timebin at **midlateral** sites C3, C4, CP1, CP2, CP3 and CP4, where an Rdensity × site interaction was observed (timebin 560–600 msec.; $F(4,76) = 3.982; p = 0.006$ ($p = 0.014$ after ϵ adjustment)). Finally, a subsequent Rdensity × hemisphere interaction in **posterior** areas indicated that this relative negativity of HighR waveforms was also significant at posterior LH sites (timebin 600–640 msec.; $F(1,19) = 5.524; p = 0.030$). This interaction is shown in Figure 6.12.

In summary, three effects were observed: at the zenith of the P_2; at the base of the P_2; and in the ensuing 120 msec. Although these effects took place at various

different sites, in all cases but one (see Figure 6.10 above for the exception) the results reflected a relative positivity for LowR over HighR stimuli.

Regular Stimuli—HighI vs. LowI Neighbourhood Density Analysis Although 16 of the 125 ANOVAs calculated in this subset analysis yielded significant results (similar to the proportion of significance in the previous 2 analyses), the effects were sporadic and generally momentary (that is, contained within one 40 msec. timebin), and are therefore difficult to interpret.

Apart from one 3-way Idensity × hemisphere × site interaction late in the epoch, where a relative positivity for LowI stimuli was found at site P8 (timebin 880-920 msec.; $F(4,76) = 3.086; p = 0.021$ ($p = 0.028$ after ϵ adjustment)), no effects of Idensity were found in frontal or posterior areas. The other significant results were found between 160–400 msec. in midline, lateral and midlateral areas, plus a late series of lateral interactions. These will be discussed in chronological order.

160–400msec.:

Two similar Idensity × hemisphere interactions were found between 160–200 msec., in **lateral** (timebin 160–200 msec.; $F(1,19) = 8.868; p = 0.008$) and **midlateral** areas (timebin 160–200 msec.; $F(1,19) = 8.839; p = 0.008$). In both cases, potentials for both conditions were similar in LH but showed a relative negativity for HighI stimuli in RH. This effect was maximal in midlateral areas ($\bar{x}$(HighI) = 3.293μV, $\bar{x}$(LowI) = 3.146μV in LH; and $\bar{x}$(HighI) = 2.724μV, $\bar{x}$(LowI) = 3.342μV in

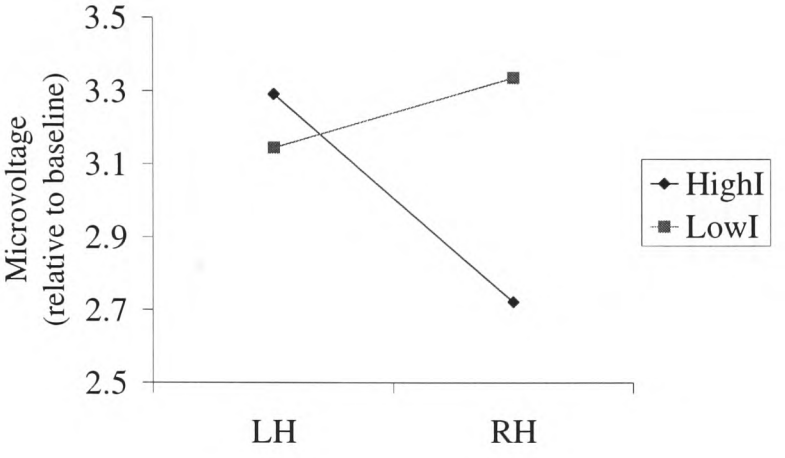


Figure 6.13: Idensity × hemisphere interaction for regulars

RH), as shown in Figure 6.13.

Idensity × site interactions were observed between 240–400 msec. on the **midline** (*i.e.* on the downward slope of the P_2), maximal between 360–400 msec. (timebin 360–400 msec.; $F(4, 76) = 5.154; p = 0.001$ ($p = 0.011$ after ϵ adjustment)). Interestingly, a relative positivity for LowI stimuli was found at frontal site FPZ, whereas at central/posterior sites CZ, PZ and OZ, HighI stimuli elicited relative positivity. No difference between conditions was found at FZ.

Finally, between 360–400 msec. (where the waveforms settle after the main peak), two interactions were found in **lateral** and **midlateral** areas. First, an Idensity × hemisphere interaction in lateral areas (timebin 360–400 msec.; $F(1, 19) = 5.153; p = 0.035$) showed relative positivity for HighI stimuli in LH and for LowI stimuli in RH.

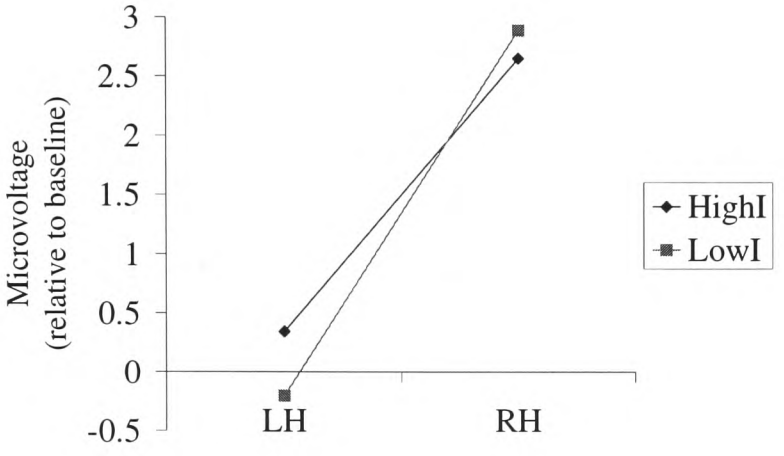


Figure 6.14: Lateral Idensity × hemisphere interaction for regulars

This interaction is depicted in Figure 6.14. Secondly, a 3-way Idensity × hemisphere × site interaction in midlateral areas showed a relative positivity for HighI stimuli at all sites in LH, but only at CP2 in RH (timebin 360–400 msec.; $F(4, 76) = 3.287; p = 0.015$ ($p = 0.023$ after ϵ adjustment)); the other 4 RH midlateral sites (*viz.* FC2, FC4, C4 and CP4) showed no effect of Idensity.

760–1000msec.:

The only *clear* effect was shown by a series of Idensity × site interactions occurring between 760–1000 msec. in **lateral** areas. This interaction was maximal between 880–920 msec. at sites CP5 and CP6 (timebin 880–920 msec.; $F(4, 76) = 5.850; p < 0.001$ ($p = 0.003$ after ϵ adjustment)). CP5 positivity was greater for LowI ($\bar{x}$(LowI) = 2.746μV) than for HighI stimuli ($\bar{x}$(HighI) = 2.033μV), and this relative positivity

was approximately twice as pronounced at CP6 in RH ($\bar{x}$(Lowl) = 4.563μV; $\bar{x}$(Highl) = 3.136μV).

In summary, apart from the late lateral ldensity × site interaction just mentioned above, the few effects found seemed to be small (albeit significant), brief and to follow no obvious pattern.

Regular Stimuli—HighR vs. LowR Neighbourhood Density Analysis A subset analysis of the HighR versus LowR regular stimuli yielded much more coherent results than the previous analysis. In total, 30 of the 125 ANOVAs carried out on these data produced significant findings, approximately twice as many as for any of the other three data subsets. The findings fall into three phases: a strong effect of Rdensity between 120–320 msec. (*i.e.* on the upward slope and at the zenith of the P_2) in all areas; a group of interactions between 520–640 msec. in all areas; and a cluster of interactions from 720–1000 msec. at midlateral and posterior sites. Between each of these three phases were periods in which no main effects or interactions were found. Each group of effects will be treated in turn.

Phase 1: 120–320 msec.

A main effect of Rdensity was observed between 120–320 msec., maximal between 240–280 msec. (*i.e.* at the zenith of the P_2) in **midline**, **frontal**, **lateral**, **midlateral** *and* **posterior** areas (just strongest in posterior areas: timebin 240–280 msec.; $F(1, 19) = 11.423; p = 0.003$)[18]. In all cases, HighR stimuli were found to elicit

[18]Although all 18 significant ANOVAs during this phase showed strongly significant F values.

potentials which were more positive than those for LowR stimuli.

Phase 2: 520–640 msec.

In the second phase, Rdensity × site interactions were found in **midline**, **lateral** and **midlateral** areas. Sites CZ, PZ and OZ (central and posterior midline), sites CP5, CP6, CP7 and CP8 (lateral) and sites CP1, CP2, CP3 and CP4 (midlateral) showed significant effects of Rdensity: the greatest single difference in microvoltage at any site occurred at PZ (timebin 560–600 msec.; $\bar{x}$(HighR) = 12.560μV; $\bar{x}$(LowR) = 9.799μV), and the most significantly affected zone was the midlateral area (timebin 560–600 msec.; $F(4,76) = 6.850; p < 0.001$ ($p = 0.006$ after ϵ adjustment)). Again, in every case, HighR stimuli elicited potentials which were more positive than those produced by LowR stimuli. Towards the end of this phase, one effect was found: an Rdensity × hemisphere interaction in **frontal** areas (timebin 600–640 msec.; $F(1,19) = 4.480; p = 0.048$), where an effect of Rdensity was observed in RH, again showing greater positivity for HighR potentials (see Figure 6.15).

Phase 3: 720–1000 msec.

In the third phase, 5 Rdensity × hemisphere interactions were found in **midlateral** and **posterior** areas; no other areas showed any effect of Rdensity after approximately 640 msec. In all cases (and contrary—it will be noted—to all effects earlier in the epoch) higher positivity for *Low*R stimuli was found in RH; in LH the differences were almost entirely insignificant except for the final timebin (960–1000 msec.),

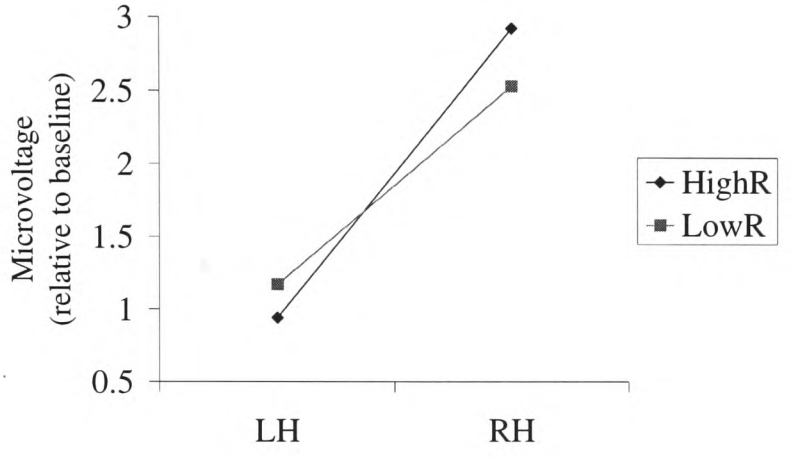


Figure 6.15: Frontal Rdensity × hemisphere interaction for regulars

wherein HighR waveforms were slightly more positive in both the midlateral and the posterior interactions. The strongest interaction was in midlateral areas (timebin 960–1000 msec.; $F(1, 19) = 19.286, p < 0.001$), and is shown in Figure 6.16. This interaction reflects the general trends of the others: LowR dominance in RH, HighR dominance in LH, with RH showing greater levels of microvoltage change compared to baseline values throughout the phase.

In summary, substantial and prolonged effects and interactions were found in this subset analysis, showing a strong relative positivity for HighR stimuli at and around the P_2, and more complex interactions (particularly affecting midlateral and posterior brain regions) later in the epoch.

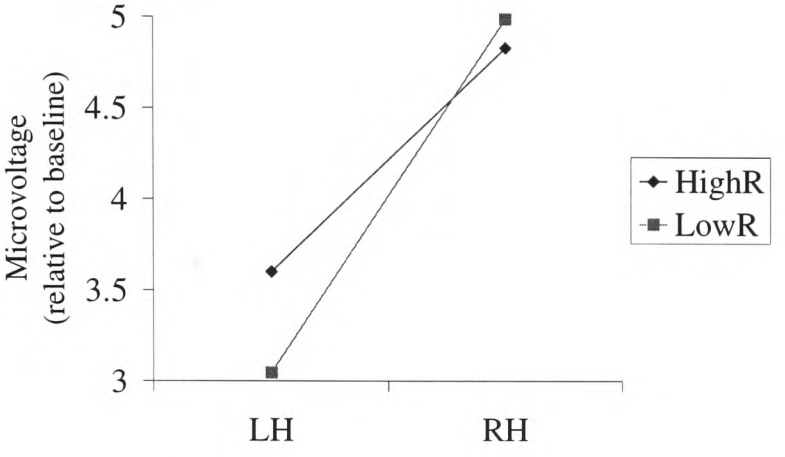


Figure 6.16: Late midlateral Rdensity × hemisphere interaction for regulars

6.1.3 Discussion

What do these results tell us? The grand-averaged data will be discussed first, followed by the four subset analyses.

Grand-Averaged Data

No evidence of an N400 was found in the data. Since this appears to be elicited in response to semantic violations or the online process of integrating sentential information (Kutas & Hillyard, 1983) this is not surprising. There was, however, evidence of syntax-dependent effects, *viz.* left anterior negativity (LAN) and, to a lesser extent, a syntactic positive shift (SPS or P600). LAN was found for irregulars and regulars between 400–600msec. after stimulus onset, and was more negative for regulars than for irregulars. A P600-like effect was also found, maximal over posterior midline sites

over a long period (c. 500–1000msec.). Positivity was much more pronounced for irregular stimuli than for regulars. It should be noted that these effects are usually found in response to violations, and may not even be syntax-specific. It is perhaps dangerous, therefore, to classify these results in terms of a P600 given the extent of its phase, and its virtual restriction to midline sites.

It is clear that regular and irregular verbs elicited substantial differences in terms of activation across much of cortex. These differences seemed to divide into three distinct phases: during the P_2 (c. 240–320 msec.); the approximately 200msec. period thereafter (360–560 msec.); and the late period of the epoch (from c. 600 msec. onwards). These phases are separated by periods in which no verb type-related differences were found.

During the **first** phase, a relative positivity for irregular verbs was noted in left hemisphere lateral areas.

During the **second** phase, regular verbs elicited a prominent negative deflection on the midline and in left anterior/lateral areas (replicating Breyer's findings); LH frontal, lateral and midlateral sites showed widespread positivity for irregular verbs.

In the **final** phase, relative negativity for regular verbs at LH lateral and midlateral sites was balanced by a strong positivity for irregulars, particularly at central/posterior midline sites.

With regard to the structures underlying these differences, it seems clear that Broca's area and/or nearby cortex is playing an active part in the inflection process, supporting the results of the Jaeger et al. PET study. This finding is also consistent

with studies into patients with Broca's aphasia, which have suggested that Broca's aphasics cannot understand the syntactic information contained in some verb inflections (Friederici, Wessels, Emmorey & Bellugi, 1992; Marslen-Wilson & Tyler, 1997; Ullman et al., 1997). However, this experiment (like that of Jaeger et al.) found activity in/around Broca's area for both verb types, which may conflict with Ullman et al.'s assertion that Broca's area 'subserve[s] the computation of regulars [as opposed to irregulars]' (see page 126 above). However, differences are so widespread that it is difficult to suggest with any accuracy what cytoarchitectonic structures may be involved.

Subset Analyses

The subset analyses yielded particularly interesting results. Generally, it was found that both Idensity and Rdensity had an effect on the neural activation levels produced by the inflection of both regular and irregular stimuli. In other words, the answer to all four questions asked on page 168 above is 'Yes', although the differences were often small and fleeting.

Looking first at irregular stimuli, it was found that Idensity affected potentials during the P_2 (at frontal/central midline and midlateral sites) and later in the epoch at frontal sites. In both cases, stimuli with an above-median number of irregular neighbours (HighI stimuli) evoked greater potentials than LowI stimuli. More widespread (but less long-lasting) effects were found for Rdensity during and for around 400 msec. after the P_2. In these cases (with one exception) a relative positivity for LowR stimuli

was found. So it appears that the number of irregulars and the number of regulars in the neighbourhood of an irregular stimulus does effect the potentials evoked during inflection, with HighI verbs eliciting more activity than LowI stimuli, and LowR verbs evoking higher levels of activity than HighR.

Turning now to the regular stimuli, Idensity seemed to have a small effect on potentials: a relative negativity for LowI regulars was found in RH lateral and midlateral areas and at central/posterior midline sites; furthermore, LowI *positivity* was found at lateral sites late in the epoch. However, Rdensity seemed to have a much more significant impact on neural activation: prominent differences between HighR and LowR stimuli were found in all cortical areas during the P_2, and (later in the epoch) in midline, lateral and midlateral areas. In every case, relative positivity for HighR stimuli was found. Interestingly, towards the end of the epoch, LowR-dominant potentials were observed in midlateral and posterior RH.

Dual-route theories do not predict the data found for regular neighbourhood density: according to the principles of explicit rule-routes, Rdensity should have *no* effect on the regularisation process, and yet many significant occurrences of relative positivity for HighR stimuli were found.

It is, perhaps, surprising that the number of regulars in the neighbourhood of a regular verb should have such an effect on the levels of neural activity during inflection, and particularly that the difference between HighR and LowR regulars was greater than any of the other factors studied in the subset analyses. Single-route theorists might suggest that the Rdensity of a regular verb would have an effect

on its inflection, but would not necessarily suppose that it would have *more* of an effect than Idensity, or that regular verbs would be more susceptible to Rdensity than irregular verbs.

It is possible that the Rdensity effect results from the fact that regular verbs with many regular neighbours may belong to a very large cluster of other verbs; this is not the case for HighI irregulars, let alone—of course—for LowR or LowI verbs. A 'HighR' stimulus may belong in a cluster of several hundred friends, whereas no HighI stimulus has more than about 14: this may be leading to the effect of Rdensity being greater than the effect of Idensity.

In summary, clear distinctions were observed between regular and irregular inflectional processes at an electrophysiological level; these differences support the widely-held view of the importance of linguistic areas such as Broca's area in grammatical processing, although by the very nature of the technique, it is not possible to say whether these differences reflect the involvement of different brain regions, or the involvement of the same brain regions to different extents. Subset analyses revealed that both Idensity and Rdensity affected microvoltage change during past tense processing: this finding contradicts the prediction of dual-route theories that Rdensity is not a factor directly available to the processing system.

6.2 Experiment Two — Novel Verbs

6.2.1 Method

This experiment followed a very similar methodology to that used in the real verb experiment in section 6.1 above, with the only major difference being that novel rather than real verbs were used as stimuli.

Subjects

20 subjects were used in this experiment, the same as those used in experiment 1 above. Each subject underwent both experiments on the same day in order to save on preparation time[19]. Half of the male and half of the female subjects participated in experiment one first and then—after a short break—experiment two; the other half did the experiments in the other order.

Stimuli

160 monosyllabic novel verb stem forms were created for use as stimuli in this experiment. As explained in previous chapters, novel verb stems (like real verbs) can be classified in terms of neighbourhood density; that is, by counting the number of real verbs sharing the rime of the novel verb, and assessing the paradigms by which those real verbs are inflected.

It should be reiterated (see chapter 4 above) that any novel verb may be classified as pseudoirregular (*i.e.* classified in terms of the median densities of all neighbourhoods containing at least one irregular verb) or pseudoregular (*i.e.* classified in terms

[19] Approximately 45 minutes per subject.

Density Type		Pseudoirregular	Pseudoregular
Friends	Enemies	Density	Density
High	Low	3.33-0.17/4.00-0.17-10.5	0.00-0.00-126.33
High	High	4.00-0.50/6.33-0.50-35.17	1.00-2.67-19.00
Low	Low	1.00-0.00/1.00-0.00-9.33	0.00-0.00-1.00
Low	High	1.00-0.00/1.83-0.00-72.17	1.50-3.17-2.83

Table 6.4: Means of stimulus types used in experiment 2

of the median densities of all neighbourhoods containing at least one regular verb). In this experiment, 80 stimuli were created for classification in each paradigm, using the same 2 × 2 factor matrix as before: HighIHighR; HighILowR; LowILowR; and LowIHighR. Once again, pseudoirregular items were classified with reference to the medians of the 87 neighbourhoods containing 1 or more irregular verbs[20], and pseudoregulars were classified with reference to the 388 neighbourhoods containing at least one regular verb[21]. The 160 selected stimuli are shown in Tables F.1–F.4 in appendix F; a summary table of means is provided here (Table 6.4).

It was originally intended that each cell of 20 stimuli would contain 20 novel verbs, each from a different neighbourhood. However, it transpired that there were fewer than 20 usable neighbourhoods in some cells (*e.g.* only 6 in the HighILowR pseudoregular cell—Table F.4, columns 4–6) and so, for balance, 6 neighbourhoods were used in each cell.

Neighbourhoods containing conflicting irregular paradigms were not used, as far

[20]Median density for pseudoirregulars = 1–1/3–1–17.
[21]Median density for pseudoregulars = 0–0–6.

as was possible; this was only unachievable in 3 cases: /ɛd/[22], /iːl/[23] and /ɪt/[24]. Note also that (in Tables F.1–F.4) clusters are shown together with the most likely irregularisation paradigm (except for those cases where no such paradigm exists within that neighbourhood, in which case just the rime neighbourhood is shown).

Finally, there was no need to match stimuli in terms of frequency, since all novel verbs *de facto* have a Kučera-Francis frequency of zero. Nvalues were calculated for each stimulus, however, to ensure that some cells did not contain significantly more stimuli which were more like real words (orthographically) than the others. A 2-way ANOVA was conducted, where paradigm (pseudoregular/pseudoirregular) and Nvalue were the between-items factors, and Nvalues were found not to be significantly different between paradigms ($F(1, 158) = 0.345; p = 0.558$). The full data are shown in appendix J.

Procedure

The procedure for this experiment was exactly as for the previous ERP experiment, except that subjects were asked to say aloud the 'most appropriate' inflected form of each stimulus. Again, subjects were seated in a soundproof booth, and stimuli were presented on screen. After 2000 msec. the stimulus disappeared, which was the subject's cue to say aloud the most appropriate inflected form of that stimulus. Responses were conveyed by loudspeaker to the experimenter. As there were no correct or incorrect inflections, subject responses were coded as regularised if the

[22]Irregulars map predominantly to /ɛd/, but also ↦ /ɒd/.
[23]Irregulars map predominantly to /ɛlt/, but also ↦ /əʊl/.
[24]Irregulars map predominantly to /ɪt/, but also ↦ /æt/.

appropriate suffixation was appended to the stem form of the verb, otherwise as irregularised.

ERP Recording

ERPs were recorded from the 20 subjects in exactly the same manner as for the previous experiment. The only difference in this experiment was that every individual epoch was coded on computer as belonging to a 'regularised' or 'irregularised' response, according to the same criteria as those used in experiments 1–4 in chapter 5 above.

6.2.2 Results

3200 trials were recorded, 160 for each of 20 subjects. 486 trials were rejected owing to the presence of artefact-related potentials (15.19% of total trials; $\bar{x}$ per subject: 26; range: 1–72). Of the 2714 accepted trials, 1855 (68.3%) were regularised ($\bar{x}$ per subject: 92.75; range: 49–126) and 859 (31.7%) were irregularised ($\bar{x}$ per subject: 42.95; range: 18–68).

Preparation of Data for Analysis

The data were re-referenced and normalised in an identical manner to the data from the previous experiment. Again, data damaged by early responses, non-responses and artefact-related potentials were eliminated, although not, of course, 'incorrect' responses, and averaged files were created in the same way as before.

A grand-averaged file comparing all regularised with all irregularised stimuli was

compiled, and waveform graphs of microvoltage change against time created as before. A summed-area analysis was carried out similarly to before, with one difference: the 'condition' factor consisted of the values 'regularised' or 'irregularised' rather than 'regular' or 'irregular'.

Again, 125 2- or 3-way repeated-measures ANOVAs were calculated with condition, site and (where relevant) hemisphere as the within-subjects factors.

Grand-Averaged Results

Visual ERPs Early visual-evoked potentials were similar across conditions, consisting of a positive peak (P_1) maximal at approximately 90 msec. across occipital sites OZ, O1 and O2, and a negative deflection (N_1) most pronounced at approximately 120 msec. at sites PO7 and PO8. No condition-related effects were found during the first 120 msec. after stimulus onset.

Language-Related ERPs Plates 3 and 4 overleaf show sample grand-averaged topographies and waveforms for regularised and irregularised novel stimuli.

The overall structure of the waveforms was similar to those found in the previous experiment: a clear positive deflection (P_2) maximal around 280 msec. was found in frontal midline sites, diminishing in magnitude towards posterior areas. However, very few significant differences were found between the waveforms for regularised and irregularised stimuli: only 17 of the 125 ANOVAs calculated produced a significant condition-related result. Broadly speaking, two effects were revealed by the analysis: a frontal effect between c. 120–320 msec., and a later lateral effect between c. 440–600

Irregularised

Regularised

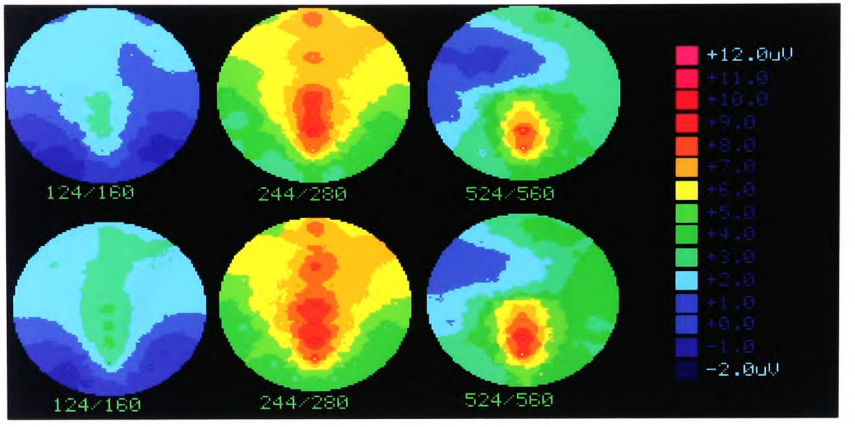


Plate 3: Grand-averaged ERP topographies for novel verbs

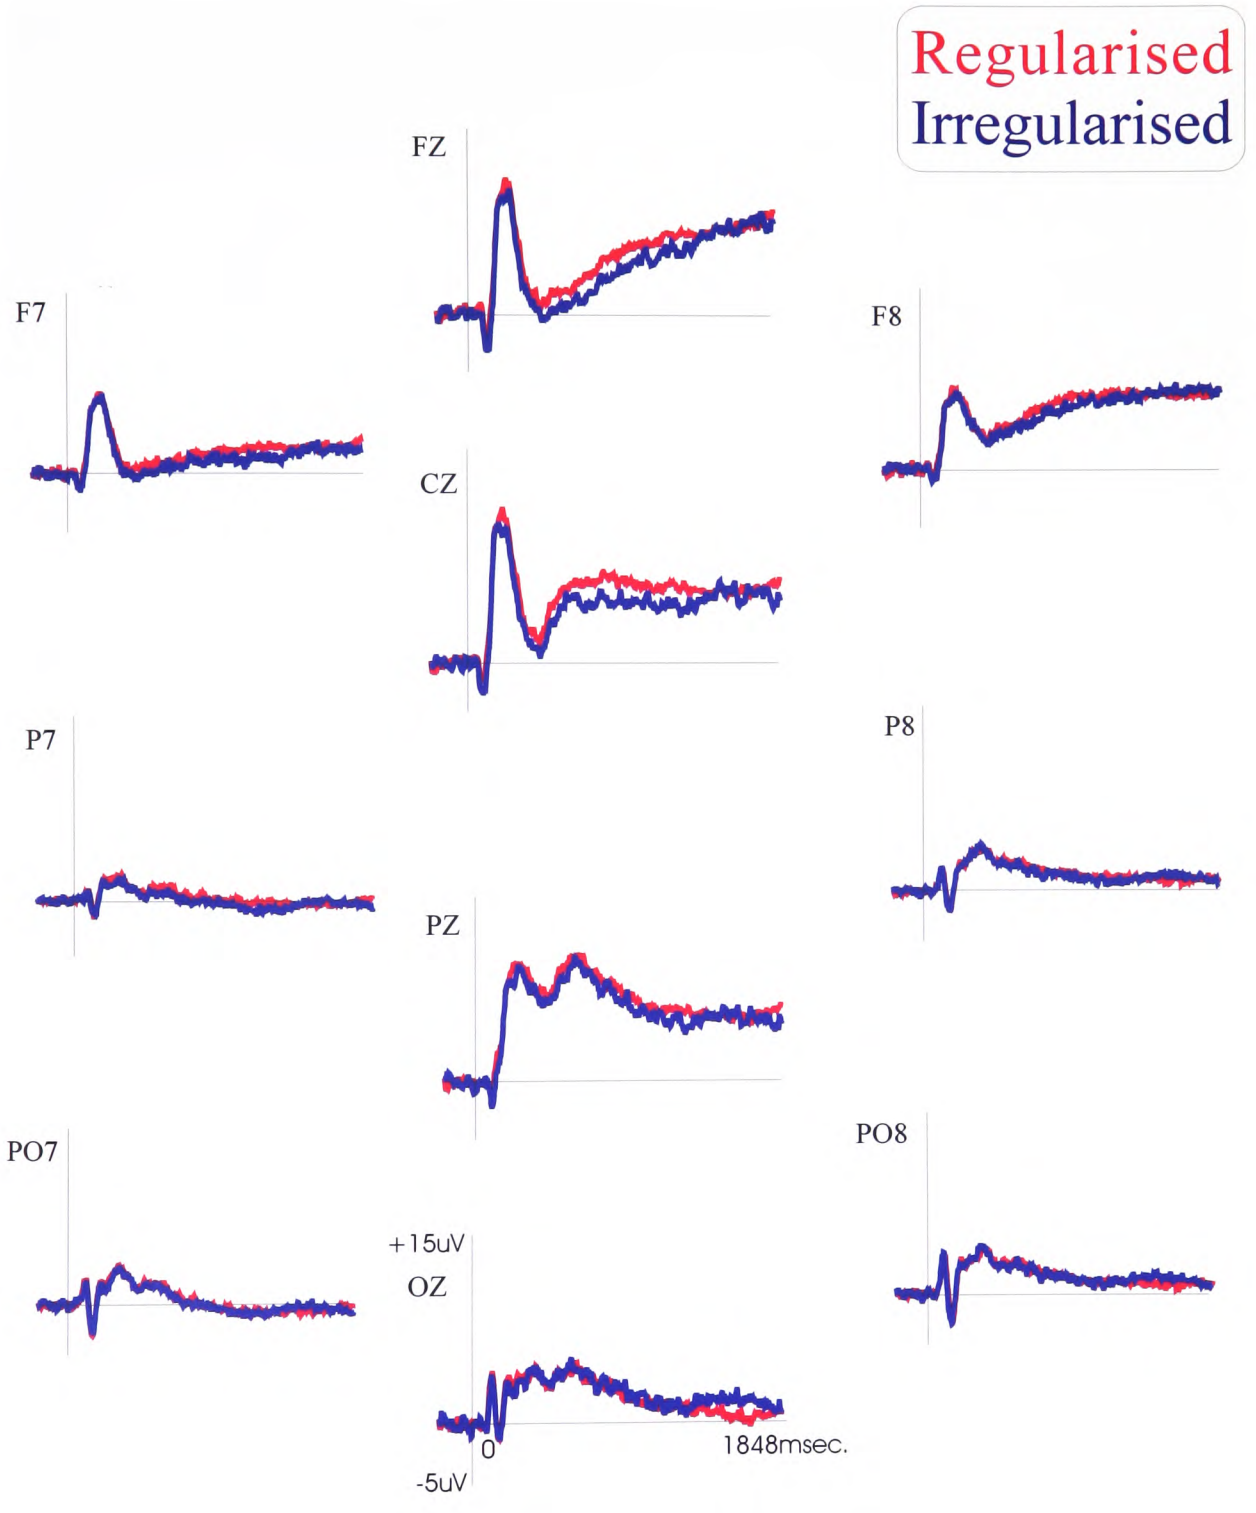


Plate 4: Grand-averaged waveforms for novel verbs

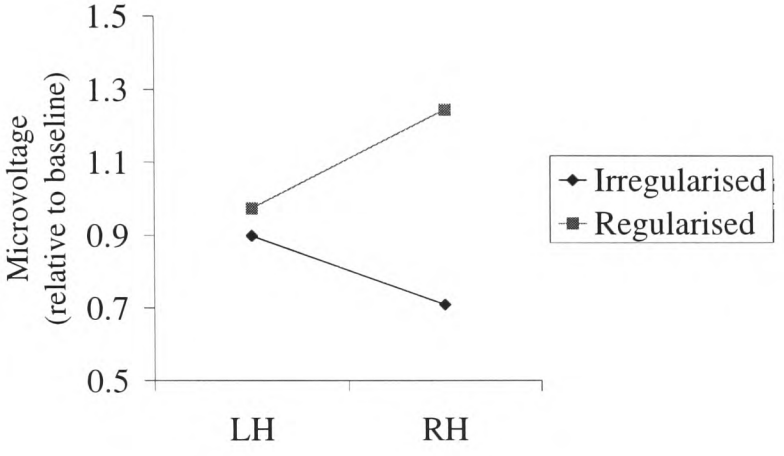


Figure 6.17: Frontal condition × hemisphere interaction (120–160 msec.)

msec. No other differences of note were found. These will be explained in turn.

Effect 1:

Two condition × hemisphere interactions were found in **frontal** areas: the first between 120–160 msec., the second between 160–200 msec. after stimulus onset (depicted in Figures 6.17 and 6.18). The first showed a relative disparity between the regularised and irregularised conditions in RH only, where regularised stimuli evoked more positive potentials than irregularised stimuli (timebin 120–160 msec.; $F(1, 19) = 5.393; p = 0.032$). The second interaction showed the reverse disparity in LH only, wherein *irregularised* stimuli evoked more positive potentials (timebin 160–200 msec.; $F(1, 19) = 4.559; p = 0.046$).

Condition × site interactions were also found throughout the phase 120–320 msec.

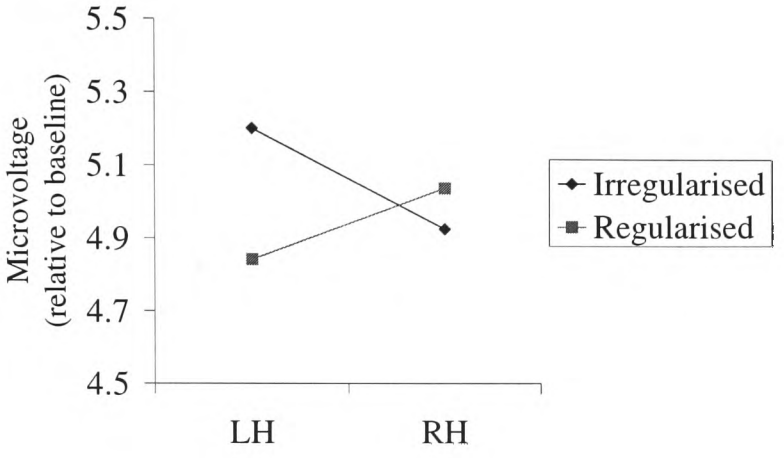


Figure 6.18: Frontal condition × hemisphere interaction (160–200 msec.)

in frontal areas, maximal between 240–280 msec. at sites FP1 and FP2 (timebin 240–280 msec.; $F(4, 76) = 4.265; p = 0.004$ ($p = 0.020$ after ϵ adjustment)). These interactions resulted from the fact that irregularised stimuli produced a higher peak at the P_2 at certain frontal sites (*viz.* FP1, FP2, AF7 and AF8).

Effect 2:

Turning to the later effect: a pure effect of condition was observed in **lateral** areas between approximately 440–600 msec. after stimulus onset, maximal between 480–520 msec. (timebin 480–520 msec.; $F(1, 19) = 5.648; p = 0.028$), where regularised stimuli evoked greater potentials than irregularised stimuli ($\bar{x}$(regularised) $= 2.273\mu$V; $\bar{x}$(irregularised) $= 1.820\mu$V). No interactions were found.

This effect seemed to be echoed briefly by midlateral areas, in which a main

effect of condition was found (timebin 520–560 msec.; $F(1, 19) = 4.788; p = 0.041$) where, again, regularised stimuli evoked greater potentials than irregularised stimuli ($\bar{x}$(regularised) = 3.366μV; $\bar{x}$(irregularised) = 2.550μV).

In sum, novel stimuli seemed to produce very similar (although not identical) waveforms regardless of which way they were inflected. The differences found suggest a general positivity for regularised stimuli, although irregularised stimuli did provoke a pattern of positive activity in LH frontal areas during the P_2.

Subset Analyses

Four subset analyses were then conducted, in a manner somewhat similar to those carried out in the previous experiment, in order to ascertain whether effects could be attributed to certain neighbourhood density types. The large number of regularised novel stimuli, both pseudoregular and pseudoirregular, meant that several subset analyses were possible, in order to see whether Idensity or Rdensity had an effect on the processing of novel stimuli:

1. HighI vs. LowI regularised pseudoirregulars;
2. HighR vs. LowR regularised pseudoirregulars;
3. HighI vs. LowI regularised pseudoregulars;
4. HighR vs. LowR regularised pseudoregulars;

Note that pseudoirregulars and pseudoregulars could not be compared directly, since the definitions of 'High' and 'Low' were different within each class. Note also that only

regularised novel stimuli were compared: the predominance of regularised responses was such that there were not enough irregularised responses per subject for subset analyses of irregularised stimuli to be statistically valid.

In each case, averaged files of each condition were created and divided into 40msec. timebins. 125 2- or 3-factor ANOVAs were then conducted on the different cortical areas, in a manner similar to those described on page 158 above, with condition, site and (where applicable) hemisphere as within-subjects variables.

The results for each of these subset analyses will be discussed in turn.

HighI vs. LowI Regularised Pseudoirregulars 394 HighI and 458 LowI regularised responses to pseudoirregular stimuli were compared in this analysis. Of the 125 ANOVAs calculated, 22 yielded significant condition-related findings. These seem to reflect three effects: a main effect of Idensity at posterior sites from about 240–880msec.; an effect of Idensity at RH frontal, lateral and midlateral sites at the bottom of the P_2; and an effect of Idensity at lateral, midlateral and midline sites around 640–680 msec. In all cases, HighI stimuli elicited greater positivity than LowI stimuli. These effects will be discussed in turn.

Effect 1:

A main effect of Idensity was found at **posterior** sites between approximately 240–880 msec. (maximal between 560–600 msec.; $F(1,19) = 8.073; p = 0.010$), in which HighI stimuli evoked more positive potentials than LowI stimuli in posterior

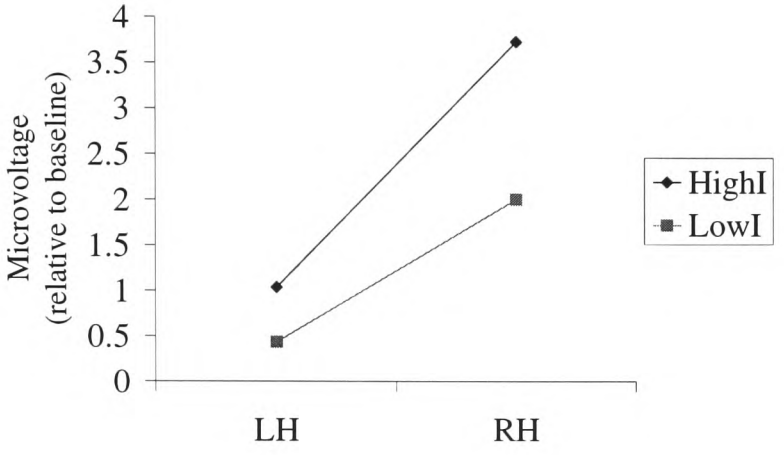


Figure 6.19: Frontal Idensity × hemisphere interaction

areas ($\bar{x}$(HighI) = 4.584μV; $\bar{x}$(LowI) = 2.761μV).

Effect 2:

Main effects of Idensity were also found in **lateral** and **midlateral** areas between 440–480 msec. (timebin 440–480 msec.; $F(1, 19) = 5.276; p = 0.033$ in lateral areas; $F(1, 19) = 5.293; p = 0.033$ in midlateral areas). Again, relative positivity for HighI over LowI stimuli was found ($\bar{x}$(HighI) = 3.537μV; $\bar{x}$(LowI) = 1.482μV at midlateral sites). Idensity × hemisphere interactions were found in **frontal** areas around this time, maximal between 400–440 msec. ($F(1, 19) = 10.399; p = 0.004$). In these interactions, a significant positivity for HighI stimuli was found in RH but not in LH; this is depicted in Figure 6.19.

Effect 3:

Finally, a main effect of Idensity was noted in **midline**, **lateral** and **midlateral** areas between 640–680 msec., maximal at midlateral sites (timebin 640–680 msec.; $F(1,19) = 7.576; p = 0.013$). Again, HighI stimuli yielded greater positivity than LowI stimuli.

In summary, several significant differences between HighI and LowI regularised pseudoirregulars were found, most noticeably and consistently in posterior areas. In all cases, potentials were found to be more positive for HighI than for LowI stimuli.

HighR vs. LowR Regularised Pseudoirregulars When HighR (N=450) and LowR (N=402) stimuli were compared in this analysis, only 8 of the 125 ANOVAs yielded significant results[25]. The only effect that seemed in any way robust was found in **lateral** and **midlateral** areas at the base of the P_2, where an Rdensity $\times$ hemisphere interaction showed that potentials for HighR stimuli were significantly more positive than those for LowR stimuli in LH (timebin 480–520 msec.; $F(1,19) = 5.550; p = 0.029$ in lateral areas). This interaction is depicted in Figure 6.20.

HighI vs. LowI Regularised Pseudoregulars 471 HighI and 568 LowI pseudoregular stimuli were compared in this subset analysis: 24 of the 125 ANOVAs were found to show significant condition-related differences, in two distinct phases: one at the peak of the P_2 in all nonposterior areas; and one later in the epoch, in all areas.

[25]Note that, at a significance level of 5%, one would expect around 6 of the 125 ANOVAs to show falsely significant, even if the conditions were virtually the same as each other.

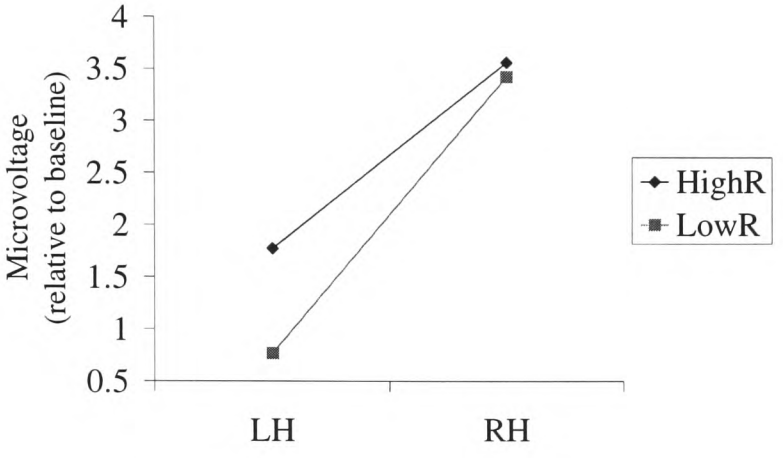


Figure 6.20: Lateral Rdensity × hemisphere interaction

These will be discussed in turn.

Phase 1: 240–320 msec.

A main effect of ldensity was found in all nonposterior areas between 240–320 msec., maximally significant in **frontal** areas (timebin 240–280 msec.; $F(1,19) = 8.147; p = 0.010$). In all cases, Lowl stimuli were found to peak more positively at the zenith of the P_2 ($\bar{x}$(Highl) = 6.165μV; $\bar{x}$(Lowl) = 7.839μV in frontal areas). ldensity × site interactions were also found during this period (between 240–280 msec.) in **midline**, **lateral** and **midlateral** areas, maximal at midlateral sites (timebin 240–280 msec.; $F(4,76) = 6.593; p < 0.001$ ($p = 0.010$ after ϵ adjustment)). On the midline, frontal/central sites FPZ, FZ and CZ showed relative positivity for Lowl stimuli: such positivity was also found at lateral sites FC5, FC6, FC7 and FT8, and

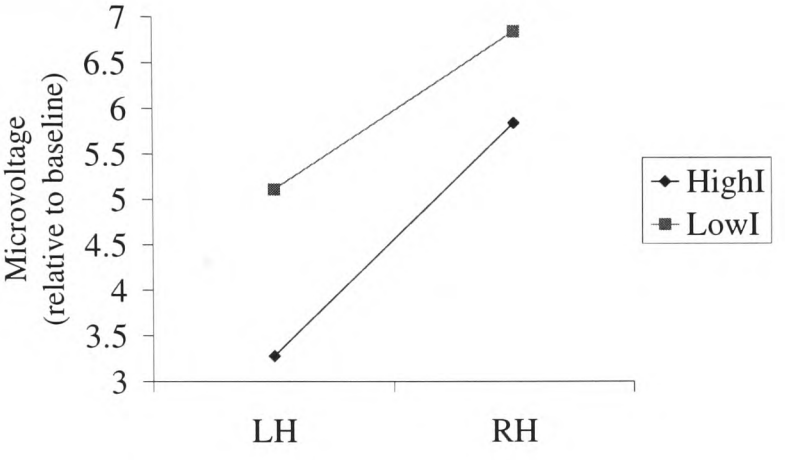


Figure 6.21: Frontal Idensity × hemisphere interaction

at midlateral sites FC1, FC2, FC3 and FC4.

Phase 2: 680–840 msec.

Between 680–800 msec., Idensity × hemisphere interactions were found in **frontal** areas (maximal between 760–800 msec.; $F(1,19) = 6.003; p = 0.024$). In these interactions, a relative negativity for HighI stimuli was found in LH, as depicted in Figure 6.21. This interaction was also found between 760–800 msec. in **lateral** areas (timebin 760–800 msec.; $F(1,19) = 5.234; p = 0.034$). In the subsequent 40 msec., a main effect of condition was found in all areas, maximal on the **midline** (timebin 800–840 msec.; $F(1,19) = 7.747; p = 0.012$). In all areas, relative positivity for LowI stimuli was found.

In summary, an effect of Idensity was found at the peak of the P_2 at certain

frontal, midline, lateral and midlateral sites. A similar widespread effect was also noted later in the epoch. In all cases of significance, LowI stimuli evoked greater positivity than HighI stimuli.

HighR vs. LowR Regularised Pseudoregulars When 496 HighR responses were compared to 543 LowR responses in this subset analysis, a mere 6 out of 125 ANOVAs showed any significant condition-related results. Each of these 6 results was a momentary interaction resulting from the activity of a single site. No two findings appeared in the same timebin, nor did any last longer than one timebin (40 msec.); these results therefore appear to be arbitrary (albeit statistically expected) and are not discussed further here.

6.2.3 Discussion

Grand-Averaged Data

Only two effects—both small—were found for the grand-averaged data in this novel verb ERP experiment. The first was a frontal effect around the P_2, during which a dominance for regularised stimuli was found in RH, and (in the subsequent timebin) a dominance for *irregularised* stimuli was found in LH. The second effect was a relative positivity for regularised stimuli in lateral areas around 500 msec. No other condition-related differences were found. The most notable finding was the fact that irregularised and regularised novel verbs elicited very similar potentials throughout the epoch, when one might have expected to find differences of neural activity similar in magnitude or distribution to those evoked by the inflection of real verbs. Although

the majority of stimuli were regularised, enough were irregularised (N=859) for the results to be robust.

Again, no evidence was found for an N400 and, again, strong evidence of an LAN was found—this was widespread over left anterior and temporal sites between c. 400–600msec. Furthermore (as for the real verb data) there was positivity around posterior midline sites CZ and PZ, beginning as the LAN subsided. This was, however, much less pronounced than for real verbs, and the results showed similar amplitude and distribution for both irregularised and regularised novel verb stimuli.

It would appear that a similar process is taking place in the production of both regular and irregular past tense inflections of novel verbs. This has certain implications for the single-route/dual-route debate. If inflection occurs via a single-route system, one would expect to find little or no epiphenomenal difference between the regularisation or irregularisation of a novel verb, since the process is the same regardless of the outcome. If, on the other hand, inflection were taking place via a dual-route system, the operational processes would be quite different, since a regularised verb is simply passed through the default suffixation 'explicit rule' route, while an irregularised verb is subjected to the irregular route involving lookup, search for similar real irregulars, and inflection by the appropriate irregular subparadigm. By this system, one might therefore expect regularised and irregularised novel stimuli to produce substantially different patterns of neural activation; this, however, does not appear to be the case.

The grand-averaged novel verb results therefore tend towards support for single-

route explanations of inflectional processing.

Subset Analyses

The four subset analyses on regularised novel stimuli yielded particularly interesting results. Similar results were found for pseudoregulars and pseudoirregulars, and produced two main findings of interest. First, Idensity (that is, the number of irregulars in the neighbourhood of a stimulus) was found to have a pronounced effect on the waveforms resulting from the (regular) processing of novel stimuli; such effects were particularly noticeable at and around the peak of the P_2. Secondly, and conversely, Rdensity (that is, the number of regulars in the neighbourhood of a stimulus) seemed to have virtually no effect on the evoked potentials of the regularised novel stimuli. In other words, the regularisation process of a novel verb was affected by the number of irregulars in that verb's neighbourhood, but *not* by the number of regulars in its neighbourhood—at least, as far as was discernible using ERP techniques.

Hybrid dual-route theories, with a rule component for regularisation and a pattern associator for known irregulars, would predict these results: Idensity should affect regularisation since partial matches with known irregulars may affect the process of regularisation; however, Rdensity should not have any such effect, as was found in these analyses. The situation for single-route theory is less clear: Idensity would be expected to have an effect on processing when regularising novel verbs since, again, partial matches with known irregulars may be reflected in processing. However, it is unclear what effect single-route theory would predict that Rdensity would have.

The data from this experiment therefore support dual-route theories of past tense processing, but do not necessarily argue for or against a single-route perspective.

6.3 Comparisons Between Real Verb and Novel Verb Findings

It has been seen (in section 6.2.2 above) that regularised and irregularised novel verb stimuli produced broadly similar waveforms—but how do novel verbs compare with *real* regularised or irregularised verbs? In order to address this question, two sets of comparative analyses were performed:

1. Regular real verbs vs. regularised novel verbs; and

2. Irregular real verbs vs. regularised novel verbs.

Unfortunately, *irregularised* novel verbs were not plentiful enough (per subject) to enable valid analyses to be performed comparing regular real verbs with irregularised novel verbs or, more pertinently, comparing irregular real verbs with irregularised novel verbs. However, the two analyses which *were* performed analysed the similarity between regularised novel verbs and regular or irregular real verbs: since regularised and irregularised novel verbs were found to be very similar in terms of scalp potentials, the results of these analyses reflected the trends of a comparison between novel verbs in general with regular or irregular real verbs in particular.

Once again, 125 2- or 3-way repeated-measures ANOVAs were conducted in each of the two analyses (1 ANOVA for each of 5 scalp areas per 40 msec. timebin).

In the analyses verbclass, site and (where applicable) hemisphere were the within-subjects factors: in the first analysis, the verbclass factor consisted of the conditions 'RealRegular' and 'NovelRegularised' and, in the second, of 'RealIrregular' and 'NovelRegularised'. Each of these analyses will be treated in turn.

6.3.1 Real Regular Verbs vs. Regularised Novel Verbs

25 of the 125 ANOVAs conducted in this comparison yielded significant results, separable into three effects: a verbclass × site interaction in frontal areas at the peak of the P_2; a lateral effect of verbclass between approximately 560–720 msec.; and a prolonged interaction between verbclass and site in posterior areas from approximately 400–1000 msec. These effects will be described in turn.

Effect 1:

A verbclass × site interaction was found in **frontal** areas between 200–280 msec. (maximal between 240–280 msec.; $F(4, 76) = 3.331; p = 0.014$ ($p = 0.042$ after ϵ adjustment)). This interaction showed a relative positivity for real regularised stimuli over regularised novel stimuli at the peak of the P_2 at sites FP1, FP2, AF3, AF4, AF7 and AF8.

Effect 2:

A main effect of verbclass was found in **lateral** areas between about 560–720 msec. (maximal between 600–640 msec.; $F(1, 19) = 10.284; p = 0.005$). Here, real regularised stimuli elicited a relative negativity at the base of the P_2 ($\bar{x}$(RealRegular)

$= 1.346\mu V$; $\bar{x}(\text{NovelRegularised}) = 2.450\mu V$).

Effect 3:

A very significant and long-lasting verbclass × site interaction was found in **posterior** areas from 400 msec. onwards (maximal around 680–720 msec.; $F(4, 76) = 11.918; p < 0.001$ ($p < 0.001$ after ϵ adjustment)). This showed a significant relative positivity for real regularised stimuli at posterior sites PO3, PO4, O1 and O2 from about 400 msec. until the end of the recorded period. This was by far the most pronounced difference found between the two stimulus types in this analysis, even though the actual microvoltage difference was—at most—approximately $1\mu V$.

In summary, regularised real and novel verbs elicited generally similar potentials in nonposterior areas. A slight frontal positivity was found for real regularised stimuli at the height of the P_2, and a relative negativity for the same stimuli was found in lateral areas after the P_2. However, the major difference between waveforms was a significant positivity for real over novel regularised stimuli at certain posterior sites for the entire post-P_2 phase of the epoch.

6.3.2 Real Irregular Verbs vs. Regularised Novel Verbs

Real irregular verbs and regularised novel verbs differed significantly in all areas for much of the post-P_2 phase of the epoch—in total, 48 out of 125 ANOVAs showed significant F-values, of which 44 occurred after 400 msec. Each area will be discussed in turn.

A main effect of verbclass was found on the **midline** from 680–1000 msec. (maximal between 760–800 msec.; $F(1,19) = 9.426; p = 0.006$), in which irregularised real verbs elicited much greater positivity than regularised novel stimuli ($\bar{x}$(RealIrregular) = 8.917μV; $\bar{x}$(NovelRegularised) = 5.682μV). This effect developed into a verbclass × site interaction from 880 msec., increasing in significance as the epoch developed (timebin 960–1000 msec; $F(1,19) = 5.177; p = 0.001$ ($p = 0.021$ after ϵ adjustment)). This interaction showed substantial positivity for real irregular verbs at central and posterior midline sites CZ, PZ and OZ, but not at frontal midline sites FPZ or FZ. This difference between waveforms was maximal at site PZ, with a microvoltage disparity of nearly 5μV ($\bar{x}$(RealIrregular) = 11.425μV; $\bar{x}$(NovelRegularised) = 6.453μV at site PZ).

Lateral areas showed an interaction between verbclass and hemisphere from 680–1000 msec. (maximal between 840–880 msec.; $F(1,19) = 24.932; p < 0.001$). These interactions showed a relative positivity for novel regularised stimuli in RH, and relative negativity for the same stimuli in LH, as depicted in Figure 6.22. This pattern of activity was also found in **frontal** areas between approximately 680–880 msec. (again maximal between 840–880 msec.; $F(1,19) = 6.727; p = 0.018$), as shown in Figure 6.23.

A main effect of verbclass was found between 720–1000 msec. in **midlateral** areas (maximal between 760–800 msec.; $F(1,19) = 6.473; p = 0.020$). Again, relative positivity for real irregularised verbs was found. From 880–1000 msec., verbclass × site interactions showed that this relative positivity was most pronounced at sites

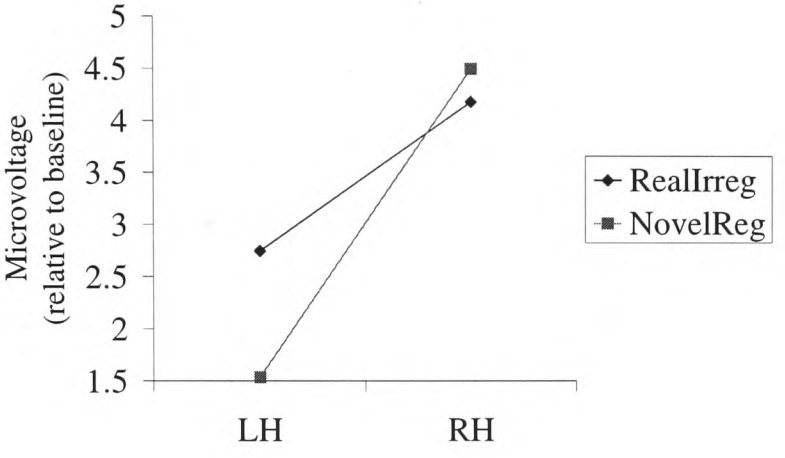


Figure 6.22: Lateral verbclass × hemisphere interaction

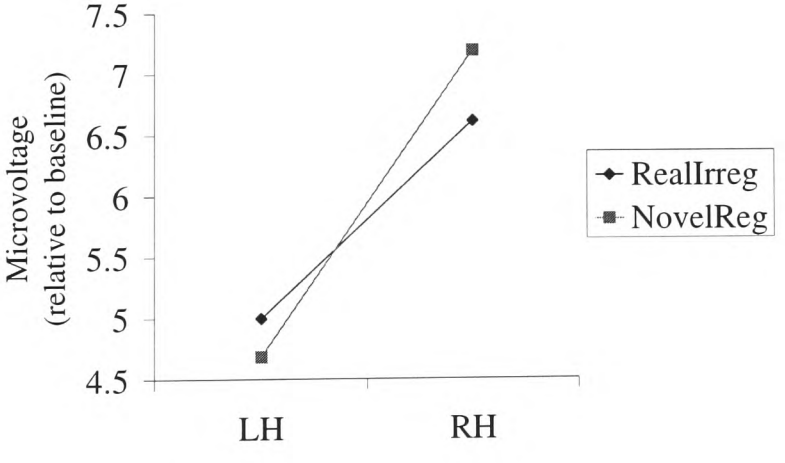


Figure 6.23: Frontal verbclass × hemisphere interaction

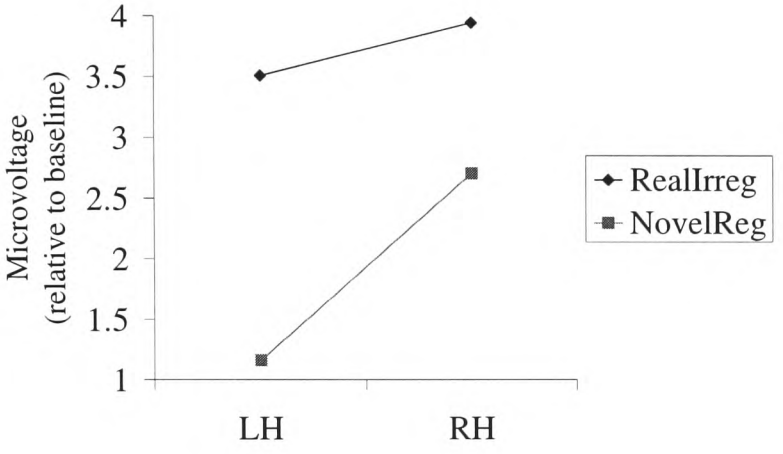


Figure 6.24: Posterior verbclass × hemisphere interaction

CP1, CP2, CP3 and CP4 (timebin 960–1000 msec.; $F(4,76) = 8.080; p < 0.001$ ($p = 0.001$ after ϵ adjustment)).

Finally, **posterior** sites showed substantial significant and prolonged differences between irregular real verbs and regularised novel verbs from approximately 400 msec. onwards. A main effect of verbclass from 440–520 msec. and from 720–1000 msec. showed a relative positivity for real irregulars over regularised novel verbs. This effect increased in significance as the epoch progressed (timebin 960–1000 msec.; $F(1,19) = 22.524; p < 0.001$). Verbclass × hemisphere interactions from 680–1000 msec. showed a relative negativity for novel stimuli which was particularly marked in LH, as depicted in Figure 6.24. This interaction was maximal between 840–880 msec. (timebin 840–880 msec.; $F(1,19) = 24.445; p < 0.001$).

In summary, real irregular verbs and regularised novel verbs elicited similar potentials between stimulus onset and approximately 400 msec., after which the waveforms differed significantly in all areas until the end of the epoch. In all cases of divergence, real irregular verbs evoked more positive microvoltage change than regularised novel stimuli.

6.3.3 Discussion

250 ANOVAs were calculated in order to compare regularised novel stimuli with regular and irregular real verbs. Regularised novel stimuli differed from *regular* verbs in three ways:

1. A superiority for real verbs at some frontal sites at the zenith of the P_2;

2. A lateral inferiority for real verbs from 560–720 msec.;

3. Posterior verbclass $\times$ site interactions from 400–1000 msec.

However, the general shapes of the waveforms for each class of verb were broadly similar in all nonposterior areas.

The same regularised novel stimuli differed from *irregular* verbs in posterior areas from around 400 msec. onwards, but also differed significantly in all other areas from approximately 700 msec. onwards.

Comparative analysis of these stimulus subsets did not yield significant findings as to the underlying causes of the general LAN and late posterior midline shift found across real and novel verb stimuli.

It will be recalled, from the real verb ERP experiment above, that irregular verbs tended to evoke greater positivity than regular verbs, in general terms. Where differences occurred in these two comparative analyses, potentials for irregular verbs were more positive than those for regular verbs, and these were both positive in comparison to regularised novel verbs. This resulted in the second comparative analysis (comparing real irregulars with regularised novel verbs) yielding more significant results than the first. Before and during the P_2, however, the only effect found in either comparative analysis was a brief frontal effect at certain sites at the peak of the P_2, showing a relative positivity for real regular verbs over regularised novel stimuli.

6.4 Conclusion

Two experiments have been described in this chapter. The first, using real verbs, found substantial differences in neural activity elicited by the inflection of regular and irregular verbs. No N400 evidence was found, although left anterior negativity and late posterior midline positivity were clearly visible from the data.

These findings broadly supported Breyer's (1997) results, although in this case more differences were found. This may be due to the use of a larger number of stimuli, with neighbourhood density carefully controlled and balanced, and an improved presentation procedure. In the subset analyses, both ldensity and Rdensity were found to affect waveforms for real verbs, contrary to the theoretical expectations of dual-route proponents.

The second experiment (using novel verbs as stimuli) found similar patterns of

left anterior negativity and late positive shift as for real verbs. It was also found that, whether an irregular or regular paradigm was used to inflect a novel stimulus, few differences could be discerned in terms of neural activity. This suggests that a similar process may be taking place regardless of the inflectional paradigm used, again supporting a general single-route theory of inflectional morphology. However, detailed subset analyses found that while irregular neighbourhood density had a marked impact on the potentials evoked in the regularisation process, regular neighbourhood density did not. Comparison should be drawn with the results of the connectionist modelling described in detail in the next chapter: when testing the artificial systems, significant differences were found for Rdensity and Idensity, at least in terms of their effect on the probability of regularising or irregularising stimuli. These ERP results therefore sit somewhat uncomfortably with a single-route approach, which may have difficulty accounting for the data.

Chapter 7

Modelling Inflectional Morphology Using Connectionist Networks

7.0 Introduction

Previous chapters have shown how—by using stimuli carefully balanced in terms of their neighbourhood structure—adults demonstrated a clear influence of regular (as well as irregular) neighbourhood density when inflecting novel verb stems to the past tense (chapter 5), suggesting that regularisation may take place by some sort of analogy system rather than the explicit, default, rule-based component posited by some theorists (*e.g.* Pinker, 1991; Marcus et al., 1992; Pinker, 1994). Similarly balanced stimuli have also been used to elicit distinct patterns of electrophysiological activation for regular and irregular verbs, novel verbs, and for subsets of regulars and irregulars varied in terms of neighbourhood structure (chapter 6). This chapter describes how five simulations were run in order to train a connectionist network to inflect English monosyllabic verb stems to their past tense forms; success on the task was tested, as was the model's ability to generalise inflectional patterns to previously unseen stimuli. Further investigations evaluated whether the network modelled the 'regularity by frequency' interaction (Daugherty & Seidenberg, 1992; Ellis & Schmidt, 1998), and a regression analysis was performed on the system's internal representations to find out whether internal function for regular and irregular stimuli could be dissociated. Finally, the network was tested on the 48 novel verb stimuli used in the behavioural experiments from chapter 5 above, and the influence of neighbourhood density on network performance was compared with the human data.

These simulations were an attempt to test the ability of a network, given the stem

form of a verb as input, to produce the appropriate past tense form. More specifically, the network was presented with a binary code of the phonological representation of the sound of the stem form of a verbal stimulus, from which it was required to produce as output the corresponding binary phonological representation of the sound of the past tense form of that stimulus.

Unlike many other connectionist models of inflectional morphology (*e.g.* Plunkett & Marchman, 1991, 1993; Plunkett & Juola, 1999), these simulations were not attempting to model the developmental profile of acquisition, but rather the end state of the system (that is to say, the 'adult' state). Therefore, the primary issues being investigated were not, for example: which types of verb are learnt most easily; what sort of errors are made during the training phase, and so on—but rather, how effective is the solution which the network finds? How do the results of testing the 'adult' network compare with the data on human adults described earlier in the thesis?

Certain assumptions were made in this model: in particular, it was assumed that the adult has some sort of mental lexicon, containing phonological representations of words; concepts such as the notion of 'verb' and 'stem' were also implicitly assumed. Problems such as that created by the inclusion in the training set of phonologically identical stems with different inflections (the homophony problem) were not considered here, although this would be solved easily enough with the addition of a simple semantic module to enable differentiation between the stimuli at the input level (MacWhinney & Leinbach, 1991).

7.1 Method

7.1.1 Network Configuration

The model was a standard 3-layer[1], feedforward artificial neural network with 128 input units, 200 hidden units and 160 output units. All units were fully connected: in other words, all units in the input layer were connected to all hidden units, and all hidden units were connected to all output units. There were no connections *within* each layer.

7.1.2 Training Corpus

As far as technically possible, the training set contained all monosyllabic English verbs (together with their corresponding past tense forms) taken from the CELEX database (Baayen et al., 1995), with the exception of 12 verbs retained for testing (see section 7.2.2 below). However, the fact that stimuli were represented phonologically immediately imposed certain constraints on the stimulus set. For instance, the network could not be expected to learn the verb pairs 'ring'↦'rang', 'ring'↦'ringed' and 'wring'↦'wrung', since all three would be represented to the network at input as the binary code for /rɪŋ/, and there would therefore be no way for the network (possessing no orthographic or semantic information) to ascertain which of the three equally correct outputs should be produced. In all such cases of conflict, the verb with the highest Kučera-Francis (hereafter 'KF') stem frequency was contained in the training set, and all other phonologically-identical stems withheld from the net-

[1] 3 layers of nodes, 2 layers of connections.

work. The training set thus only contained phonologically unique stimuli (at least, before frequency of each individual verb was taken into account—see page 218 below), and the homophony problem was thereby avoided. In cases where two or more phonologically-identical verb stems have the same KF frequency (*e.g.* 'ring' (↦'rang') and 'ring' (↦'ringed')), the stem-past pair with the highest product (*i.e.* inflected) frequency was selected for inclusion[2]. It was considered sensible always to use the higher frequency stimulus in cases of conflict since a human subject, if presented with the sound of a word (say, /rɪŋ/) and asked to inflect it to the past tense, might reasonably be expected to think of the verb with the higher frequency in the language when considering his or her response.

First, the orthographic and phonological (DiSC coded[3]) stems and past tense inflections of all English words classified as verbs were extracted from the CELEX corpus (N = 5866). Then all polysyllabic stems were removed, leaving 1970 monosyllables. Token frequencies of all stems were taken from Kučera & Francis (1967)[4] and in all cases where more than one stem had the same phonological coding, only the stem-past pair with the highest stem token frequency was included in the stimulus set. This left 1871 stimulus/product pairs, which were then checked by hand to eliminate errors[5], such as the verb 'wound' (↦'wounded') being incorrectly coded as */waʊnd/↦*/waʊndɪd/. Finally, possible linking 'r's[6] were also removed in order to

[2]In this case, 'ring'↦'rang'.

[3]'Distinct Single Character' coding, in which each phonetic segment (here meaning consonant, affricate, short vowel, long vowel or diphthong) is represented by a single ASCII character.

[4]Used as an approximate measure of the frequency of the stem in the language.

[5]Which are, unfortunately, rife in the CELEX database.

[6]Coded as /*/ in IPA and as /R/ in DiSC coding.

replicate Southern British English standard pronunciation.

Twelve verb stems were then removed from the corpus, to be used as 'real verb' test stimuli after training.

The resultant corpus consisted of 1859 phonologically unique stem-past pairs (N(reg) = 1710, N(irreg) = 149), with token stem frequencies ranging from 0 to 3941 (μ=6)[7].

In order to incorporate token frequency into the training set, each stimulus did not simply appear $freq$ times[8], or even $freq+1$ times (allowing stimuli with a token frequency of zero to be presented once per epoch to the network), since some verbs would still appear nearly 4000 times more frequently than most of the others, thereby in effect 'hijacking' the network. In order to flatten the top end of the frequency distribution, the function $l = log_e(freq^2 + 1)$ was therefore applied to the Kučera-Francis frequency of each stimulus in the corpus, yielding a log distribution of the set. Using $freq^2$ allows for a wider spread of log frequencies than if $freq$ had been used—frequency can therefore still have an impact on the training set, without being 'overloaded' by a few very high-frequency verbs. Note also that by using '+1' in the equation, even verbs with a frequency of zero are still presented to the network once per epoch.

Each stimulus was then repeated l times in the final training set (with l ranging from 1–17, $\mu = 4$). In other words, all stimuli appeared at least once in every epoch,

[7]Where μ signifies the median.

[8]Where $freq$ is its token frequency in the language.

and the stimulus with the highest token frequency was presented to the network 17 times per epoch. This process resulted in a full training corpus of 8840 present and past tense verb pairs. The full input corpus is listed (with log frequencies) in appendices K and L.

7.1.3 Representation

The training corpus then had to be translated into a numerical representation intelligible to the network. This was done in three stages. First, the DiSC-coded phonetic representations of the stimuli (stems and products) were converted to the Moby phonetic alphabet (Ward, 1997) using a specially written Perl script. The equivalent codes for International Phonetic, DiSC and Moby alphabets are shown in Table 7.1. It should be stated that slight alterations to standard Moby alphabetic representations were made in order to ensure that all long vowels and diphthongs were represented using two symbols, and all short vowels and consonants were represented using only one. Note also from the Table that affricates were treated as stop-plus-fricative sequences rather than as individual phonetic segments.

Secondly, the Moby-coded stimulus stems were then aligned into a right-justified CCCVVCCC input template (where 'C' represents a slot for a consonantal phoneme, and 'V' a slot for a vocalic phoneme). Thus, for example the verb 'bring', coded as /brIN/ in the Moby alphabet, becomes '#br#I##N' in the template, where '#' represents a blank segment. Note that phonemes are only right-justified *within* their respective consonantal or vocalic cluster. Similarly, the corresponding past tenses

Vowel Type	IPA	DiSC	Moby
Consonants	p	p	p
	b	b	b
	t	t	t
	d	d	d
	k	k	k
	g	g	g
	ŋ	N	N
	m	m	m
	n	n	n
	l	l	l
	r	r	r
	f	f	f
	v	v	v
	θ	T	T
	ð	D	D
	s	s	s
	z	z	z
	ʃ	S	S
	ʒ	Z	Z
	j	j	j
	x	x	x
	h	h	h
	w	w	w
Affricates	ʧ	J	tS
	ʤ	-	dZ
Linking 'r'	*	R	(removed)
Short vowels	ɪ	I	I
	ɛ	E	E
	æ	{	&
	ɒ	Q	O
	ʌ	V	@
	ʊ	U	U
Long vowels	iː	i	ii
	ɑː	#	AA
	ɔː	$	OO
	uː	u	uu
	ɜː	3	[@][@]
Diphthongs	eɪ	1	eI
	aɪ	2	aI
	ɔɪ	4	OI
	əʊ	5	oU
	aʊ	6	aU
	ɪə	7	I@
	ɛə	8	(@)(@)
	ʊə	9	U@

Table 7.1: IPA, DiSC and Moby alphabetic equivalents

were aligned into a CCCVVCCCVC template, the two extra segments at the end allowing for representation of any of the three types of regular suffix ('#d', '#t' and 'Id') within the output template. So, for instance, 'wanted' would be aligned into the template as '##w#O#ntId'.

Finally, the stem–past verb pairs could then be transformed into binary sequences interpretable as input and target output patterns. Each of the 8 segments of the template was converted using the PGPFone binary alphabet (Juola & Zimmerman, 1996) into a 16-bit sequence, composed as follows:

Bits 1–3: Height of articulation;

Bit 4: Tenseness;

Bit 5: Nasality;

Bit 6: Laterality;

Bit 7: Voicing;

Bits 8–10: Manner of articulation;

Bits 11–16: Place of articulation.

Single-bit features had binary values; multi-valued features were represented by right-justified thermometer coding (using strings of 1s) into 3- or 6-bit segments of the sequence. The complete 16-bit sequence is thus composed as in Figure 7.1.

The Moby-to-PGPFone translations for all phonemes used are shown in Table 7.2.

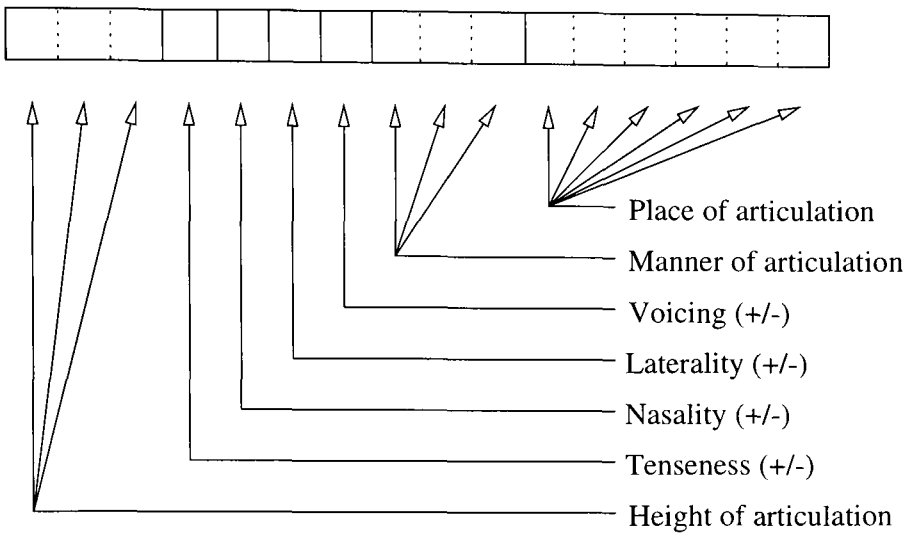


Figure 7.1: Composition of 16-bit binary sequence representation

After each segment of the template had been translated, these 16-bit sequences were then strung together to make one long sequence representing the whole template. Resultant sequences were thus 128 bits in length (16×8) for input patterns, and 160 bits long (16×10) for target output patterns.

7.1.4 Training

The system was trained using a mass (rather than incremental) training set. In other words, all stimuli were presented every epoch. The network used a supervised learning mechanism, learning taking place by backpropagation of error (Rumelhart et al., 1986). Input patterns were presented individually to the network, in random order (without replacement), and connections updated after presentation of each pattern. Learning rate (η) was set to 0.1 and momentum (α) to 0.

Every time the network was presented with an input pattern, it calculated and

Moby code	Approximate Pronunciation	Binary Sequence
&	"a" in "dab"	0001001111000111
(@)	"a" in "air"	0010001011000111
@	"a" in "ado"	0010001111000011
A	"a" in "far"	0000001111000001
a	first half of "i" in "ice"	0001001111000001
b	"b" in "nab"	1110001000111111
C	"ch" in "ouch"	1110000000000011
D	"th" in "the"	1110001001001111
d	"d" in "pod"	1110001000000111
E	"e" in "red"	0010001111000111
e	first half of "a" in "day"	0011001111000111
f	"f" in "elf"	1110000001011111
g	"g" in "fig"	1110001000000001
h	"h" in "had"	1110000011000001
I	"i" in "hid"	0110001111000111
i	"e" in "see"	0111001111000111
J	"g" in "vegetable"	1110001000000011
j	"y" in "you"	0110001011000000
k	"c" in "act"	1110000000000001
l	"l" in "ail"	1110011011000111
m	"m" in "aim"	1110101000111111
N	"ng" in "bang"	1110101000000001
n	"n" in "and"	1110101000000111
O	"o" in "dog"	0010001111000001
o	first half of "o" in "boat"	0011001111000001
p	"p" in "imp"	1110000000111111
[@]	"u" in "burn"	0010001011000011
r	"r" in "ire"	1110001011000111
S	"sh" in "she"	1110000001000011
s	"s" in "sip"	1110000001000111
T	"th" in "bath"	1110000001001111
t	"t" in "tap"	1110000000000111
U	"oo" in "book"	0110001111000001
u	"oo" in "too"	0111001111000001
v	"v" in "average"	1110001001011111
w	"w" in "win"	0110001011000001
Z	"s" in "vision"	1110001001000011
z	"z" in "zoo"	1110001001000111
#	silence/vacant	0000000000000000

Table 7.2: Binary encoding based on Moby Pronunciator

produced its 'best guess' of the appropriate output pattern; this was compared with the target (*i.e.* correct) output pattern, and error was calculated for the output units in proportion with the Root Mean Squared (RMS) error; appropriate alterations to the strength of connection weights were then backpropagated through the network to increase the probability of the network producing the correct output pattern on subsequent presentation of the same stimulus. After every epoch the weights matrix was saved in a computer file.

The model was built on a UNIX platform using version 5.7 of the PlaNet simulator (Miyata, 1991). Five simulations were run[9], all identical except for the strengths of the initial weight matrix, which were randomised in the range ± 0.5 by means of different starting seeds for each simulation. Each simulation was trained for 250 epochs. Unless specifically stated otherwise, all results refer to mean findings across all five simulations.

7.2 Results

7.2.1 Error Rates for Training Set

After every epoch of training, the entire training set was presented to the network (without error propagation) in order to test how much of the corpus had been learnt. Network responses were fed through a computer program and each 16-bit segment of the output pattern was compared with a list of all the 16-bit phoneme representations in the Moby pronunciator: the phoneme representation which most closely matched

[9]See appendix M for the training script.

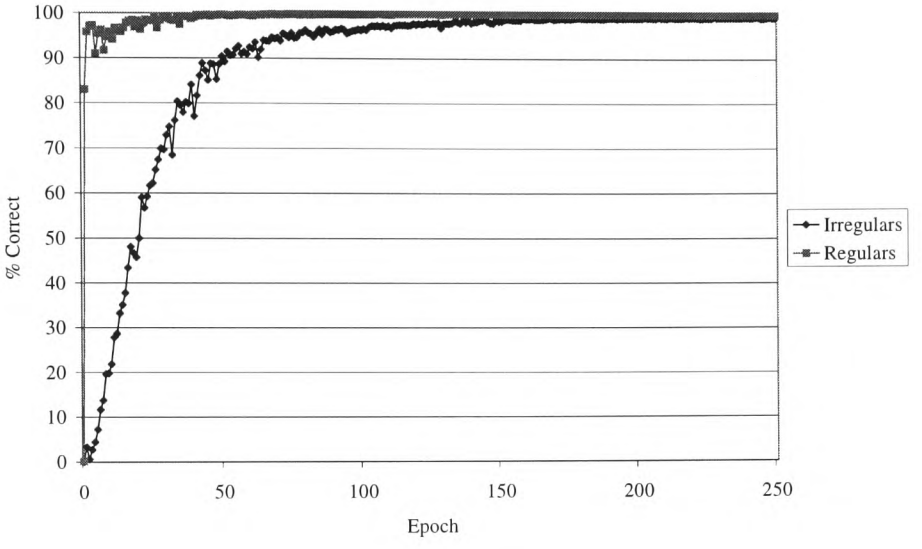


Figure 7.2: Percentage of tokens correct by verb type

the output segment was selected; once the whole output pattern had been parsed in this way, the selected phonemes were strung together in sequence and recorded in DiSC-coded format in a computer file. Responses were classified as correct if the DiSC-coded network output exactly matched the DiSC-coded target output from the training set.

Results for all simulations showed that the network was highly proficient at learning to inflect English verb stems to their past tense forms. Figure 7.2 shows the percentage of correct token inflections against epoch, averaged across the five simulations.

Regular Errors

As Figure 7.2 shows, after only one epoch, the network produced correct suffixation responses for 83% of the regular stimuli; this increased to 96% after the second epoch. By epoch 50, regularisation was consistently over 99% correct. After 250 epochs, correct regularisation was found to occur in 99.98% of cases. In fact, when all five simulations were tested at the end of training, only two regular verbs were found to have been incorrectly inflected. These were 'neigh' (incorrectly inflected by simulations 1, 2 and 4) and 'rive' (incorrectly inflected by simulations 1, 4 and 5). Both words have a Kučera-Francis frequency of 0, and were thus presented to the network just once per epoch.

Interestingly, in all 3 cases where 'neigh' was processed incorrectly it was inflected to */næiːd/ instead of /neɪd/, in spite of the fact that the vocalic cluster /æiː/ did not appear anywhere in the training set input or output patterns. Note also that, although (incorrectly) containing a vowel-change, the word was nonetheless consistently (correctly) suffixed.

'Rive' was also incorrectly inflected in three simulations, although this time in three different ways: to */rɔiːvd/, */rauːvd/ and */rəuːvd/. It is interesting that the network demonstrated a tendency to (over-)lengthen the second vowel-segment of diphthongs to produce the above outputs rather than */rɔɪvd/, */raʊvd/ or */rəʊvd/, which would at least incorporate vowel sounds on which the network had been trained. Again, despite the internal vowel changes, the stimulus was consistently suffixed.

A quick survey of 20 University students showed that, although most knew the meaning of the word 'rive', only 25% knew that it is regular in the past tense. This may, however, be a confound due to the fact that it is irregular in its most common form, the past participle 'riven'—although this was not information available to the network.

Irregular Errors

Irregular verbs elicited a different pattern, producing a rather gentler learning curve. Successful irregular inflection reached 50% after approximately 20 epochs, and 90% by epoch 50. 95% success was reached after about 70 epochs, and irregular inflection was consistently above 99% after 190 epochs. After 250 epochs, the success rate for irregular verbs was 99.32%, an average of 5.2 type errors per network.

All five simulations produced incorrect output on the irregular verbs 'gird', 'slay' and 'smite' (KF frequencies of 1, 0 and 0 respectively). All five simulations zero-marked 'gird' to */gɜːd/; 'slay' was inflected to */sleɪd/, */slɛ/, */sliː/ and */sleiː/ (twice); and 'smite' was inflected to */smauːt/ four times and */smaiːt/[10] once. It is worth noting that 'slay' and 'smite' are not only 0-frequency verbs, but also have neighbourhood competition from other irregulars as well as regulars.

Three simulations failed to inflect 'flee' correctly, two producing */flɪd/ and one the overregularisation */fliːd/. Two simulations zero-marked 'rend', two produced */krɛt/ instead of /krɛpt/, and two inflected 'cleave' to */klɒuːv/ rather than to

[10] Almost zero-marked as /smaɪt/.

/kləʊv/; finally, one simulation (simulation 5) produced */sɛt/ instead of /sæt/ and */swɛm/ instead of /swæm/.

It is interesting that there was such a considerable overlap of errors between simulations even though in training, 547 verbs from the corpus had a KF frequency of 0 or 1, and were thus only presented once per epoch. It would appear that a few irregular verbs were simply harder for the networks to learn than the rest.

These results show that the networks were extremely successful at learning to map English verb stems to their past tense forms and, by the end of training, produced results at least as good as would be expected for adult humans.

7.2.2 Generalisation to Unseen Stimuli

After being trained for 250 epochs, the generalisation ability of the network was tested, by presenting it with 12 real verb stems which had been removed from the corpus before training. 6 of these verbs had regular past tense forms and 6 were irregular. Of the regulars, 2 ended in each of the 3 regular suffixes, *viz.*: 'blab' and 'furl' (+/d/); 'clamp' and 'doff' (+/t/); and 'drat' and 'jilt' (+/ɪd/). All 6 had been selected from a pool of monosyllabic regulars with a KF frequency of zero. The 6 irregulars were all exemplars of different neighbourhood clusters: 'shrive' (↦/ʃrəʊv/, KF=0), 'bleed' (↦/blɛd/, KF=2), 'grind' (↦/graʊnd/, KF=2), 'kneel' (↦/nɛlt/, KF=5), 'swear' (↦/swɔː/, KF=10) and 'sweep' (↦/swɛpt/, KF=15). All had a KF frequency below the median for monosyllabic irregular verbs (where $\mu = 37$). Note also that 2 of the 6 irregulars and 2 of the 6 regulars ended in alveolar stops.

These 12 test stimuli were each presented once to each of the five simulations. When network output was decoded it was found that every simulation had correctly inflected all twelve stimuli.

7.2.3 Regularity by Frequency Interaction

Background

A key behavioural dissociation between regular and irregular verbs is found in the so-called 'regularity by frequency interaction'. Reaction-time experiments on inflectional tasks have shown that response latencies are influenced by the frequency of past tense forms when inflecting irregular stimuli to the past tense, but *not* when inflecting regulars (Prasada et al., 1990; Seidenberg & Bruck, 1990). Furthermore, this finding appears to be valid even when stem frequencies are equated. Regulars have an overall tendency to be inflected faster than irregulars; however, inflection of irregular verbs with high-frequency past tense forms is faster than for those with low-frequency past tense forms, whereas past tense frequency for regular verbs seems to have no effect on response speed. This regularity by frequency interaction has also been found in non-native speakers, and has been used as evidence for a symbolic linguistic module used in first and second language acquisition (Eubank & Gregg, 1995; Pinker, 1991).

Pinker (1991) argued that, if irregulars are stored in an associative network, one would expect low-frequency irregular forms to take longer to be accessed than high-frequency forms and, conversely, that frequency should have no effect on a regular verb if it is inflected by a symbolic suffixation process (as frequency effects are a general

property of associative memory systems, but not of symbolic rule-based systems). Since the regularity by frequency interaction seems to show precisely those effects, Pinker used this dissociation as evidence to support the hybrid dual-route model of inflectional morphology.

Ellis & Schmidt (1998) countered this argument, however, by suggesting that the regularity by frequency interaction is also generated by the 'Power Law of Practice', a basic learning principle, and could thus be replicated in a connectionist network. The power law states that performance improves with practice, and that the amount of improvement decreases as a function of increasing practice or frequency (Anderson, 1982). The law has been investigated in relation to many tasks, including (in this context) linguistic functions such as lexical recognition (Kirsner, 1994) and comprehension using second-language morphosyntax[11] (DeKeyser, 1997). Ellis & Schmidt claimed that the power law of learning is ubiquitous and quantitative, and that a connectionist network should be able to simulate its effects.

They conducted two experiments in support of this claim: a behavioural experiment on humans, and a connectionist simulation.

In the behavioural experiment, they created a miniature artificial language (or MAL) containing twenty 'nouns'. Subjects spent an hour a day learning these words by associating them with picture stimuli of standard items such as an umbrella, a car, and so on (taken from Snodgrass & Vanderwart, 1980). The words were designed to

[11] See Ellis (1996) for a discussion of the implications of the power law for second-language acquisition.

be memorable by suggesting English cognates (*e.g.* 'pisc' for fish, 'brol' for umbrella, 'charp' for chair).

10 of the nouns were regular, plural forms being created by concatenation with the affix 'bu-' (*e.g.* 'bubrol'); the other 10 were irregular, each being inflected by a different affix (*e.g.* 'ropisc', 'zecharp'). Half of the 10 regulars and half of the 10 irregulars were high-frequency, being presented exactly five times as often as the remainder. The stimuli thus formed four 'cells', which Ellis & Schmidt called 'HighFreqReg', 'LowFreqReg', 'HighFreqIrreg' and 'LowFreqIrreg'.

Subjects were first presented with pictures of the stimuli—one at a time—on a computer monitor and required to say aloud the 'name' of the item in the picture as quickly as possible. Learning was 'supervised' in that the computer would then speak the correct word. All items were presented an equal number of times, and in the singular form (*i.e.* frequency and regularity were not being manipulated at this stage). This training continued until the subjects could begin to name the items within 2 sec. of stimulus onset, with 100% correct performance.

The second stage was similar, except that subjects were presented with either one picture or two (the same picture repeated) on screen. If there were two pictures, subjects were to say aloud the plural form of the 'name' of the item as quickly as possible. Over time, their performance improved until it was 100% on all items. In this task, 'high-frequency' nouns were presented 5 times more frequently than 'low-frequency' nouns.

Like Prasada et al. (1990), Ellis & Schmidt found significant effects of frequency

and regularity, where high-frequency items were inflected with greater accuracy than low-frequency items, and where regular plurals were learnt better than irregular plurals. A regularity by frequency interaction showed that the frequency effect was larger for irregular items than for regulars. However, contrary to Prasada et al., they also found that in the early stages of acquisition there was also a frequency effect for *regular* items, and that this effect diminished with learning.

They then built a connectionist network and exposed it to the same stimuli, in the same order, as the human subjects, in order to see whether they could parallel the human results. This was not the first connectionist network to address the 'regularity by frequency interaction' issue: Daugherty & Seidenberg (1992) had previously attempted to model the interaction in two simulations. However, they found no interaction in the first model (using 309 regular verbs and 104 irregulars), and the effect was only successfully simulated when they skewed the input set by reducing the number of irregulars from 104 to 24.

Ellis & Schmidt designed a multi-layer, feedforward network which used high-level representations rather than phonological coding, having just one input node for each picture (what Paivio (1986) would call an 'imagen') and one output node for each morpheme (what Morton (1979) would call a 'logogen'). The network was trained for 500 epochs on singular forms, followed by 700 epochs on singulars and plurals. By plotting RMS error against epoch, they found that the network output paralleled the human data. In both the human and the network experiments, they found that high-frequency items were learnt faster than low-frequency ones; that regulars were learnt

faster than irregulars; and that both humans and the network showed the regularity by frequency interaction, whereby frequency had less of an effect on regulars than on irregulars. They concluded that the regularity by frequency interaction was easily explained by an artificial neural network and, far from Pinker's arguments, was not "any reason to reject connectionist accounts of morphosyntax in favour of a hybrid model" (page 333).

Ellis & Schmidt's findings were significant and interesting. However, it must be remembered that they used an artificial language which, despite its obvious advantages in terms of accurate monitoring of acquisition, was nonetheless "a travesty of natural language" (page 328), to use their own words, and contained only a very small number of words. Furthermore, by using an imagen↦logogen network structure, they avoided the potential over-reliance on phonological cues (MacWhinney, 1994), but there was therefore no opportunity to reflect the cryptoembodiment of 'rules' within the phonological structure of stimuli, which may affect frequency or regularity effects.

It was therefore decided to attempt to replicate Ellis & Schmidt's findings here, using a network which had been successfully trained on phonological representations of virtually all monosyllabic English verbs.

Data

The 1859 stimuli from the original training set were divided into separate files according to whether they were regular or irregular, and whether they were high-frequency,

Frequency	Class	N
High	Irregular	120
High	Regular	750
Low	Irregular	16
Low	Regular	811

Table 7.3: Number (N) of verbs per category in regularity × frequency analysis

medium-frequency or low-frequency. The dividing-point for frequency was decided by calculating the median of all log frequencies used in the network training. Given that all stimuli were presented for training between 1 and 17 times per epoch (the l value shown in the Tables in appendices K and L), and given that $\mu(l) = 4$, low-frequency stimuli were defined as those verbs presented fewer than 4 times per epoch, and high-frequency stimuli defined as those which were presented more than 4 times per epoch. Stimuli with an l-value of exactly 4 (*i.e.* medium-frequency stimuli) were excluded from the analysis. This yielded the category distribution shown in Table 7.3, where N is the number of stimuli within each category.

Note that the total number of stimuli used was 1697, since all 13 medium-frequency irregulars and all 149 medium-frequency regulars were left out of the analysis. Four separate pattern files were created on computer for use as input to the network, each containing a phonologically-represented list of all of the stimuli within each of the four classes shown in Table 7.3.

Using a simple computer script, the weights matrices which had been saved after each epoch were loaded, one at a time, into the network (thereby returning it to the state it had been in at that point in training); the pattern files were then presented

separately, and the root mean square (RMS) error from the network recorded; this was repeated for each of the four stimulus types, and then for all 250 epochs, and then repeated again for the other four simulations. Results averaged across all five simulations are shown in Figure 7.3.

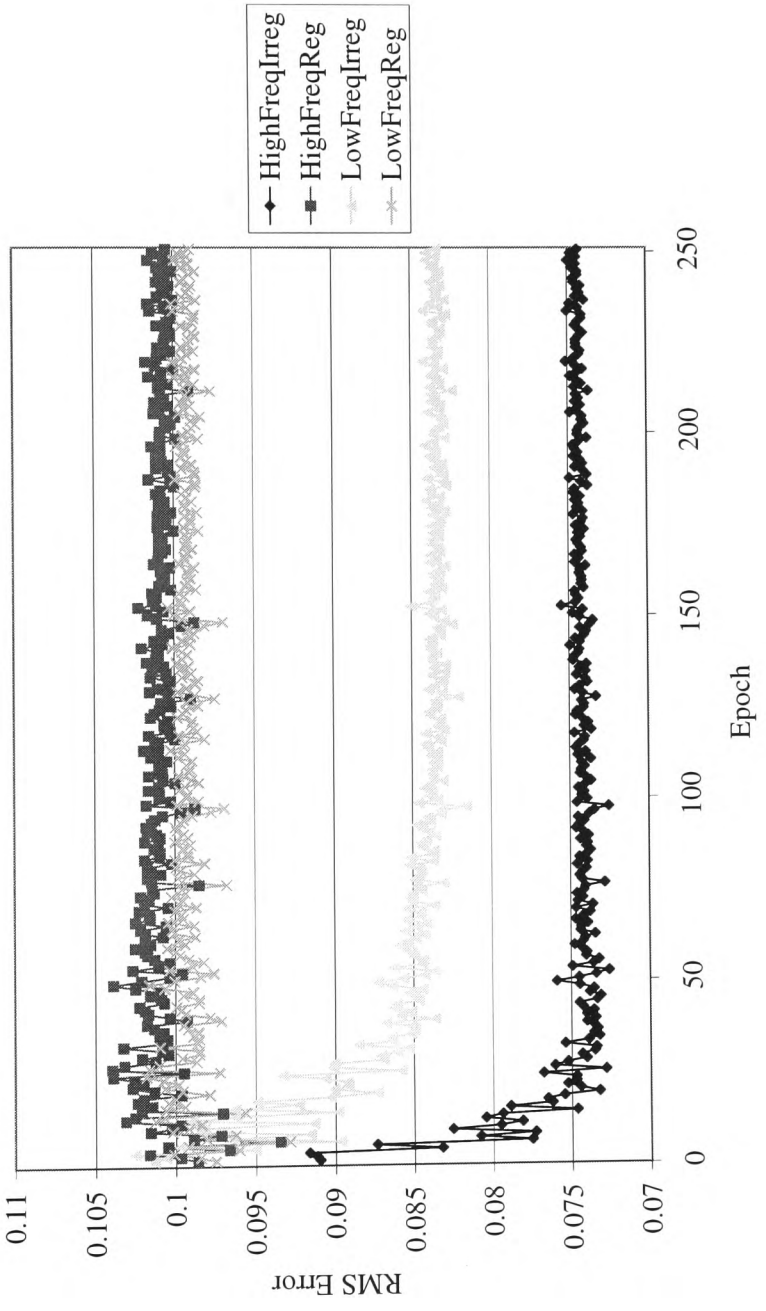


Figure 7.3: Regularity by frequency interaction across all five simulations

A main effect of frequency was found ($F(1,4) = 2766.361$; $p < 0.001$), whereby high-frequency verbs produced a lower error score than low-frequency stimuli; a main effect was also found for regularity ($F(1,4) = 17283.474$; $p < 0.001$) where, interestingly, irregular stimuli showed a lower error score than regulars—the opposite finding to the Ellis & Schmidt model. An interaction between the two was also found ($F(1,4) = 6688.446$; $p < 0.001$), whereby frequency had a significant effect on irregular verbs (Scheffé: $F = 11.33; p < 0.001$) but not on regulars ($F = 2.09; p > 0.05$). These results may be seen in Figure 7.3, which shows the averaged RMS error for high and low frequency regulars and irregulars plotted against training epoch. It is clear from the Figure that error for all four word types was stable after the first 50 or so epochs, and changed little for the remainder of the training process. Note the clear interaction between regularity and frequency: not only is the error for regulars and irregulars different, but high-frequency and low-frequency regulars are virtually indistinguishable on the graph, whilst a noticeable effect of frequency on irregulars is visible.

The results are fascinating in that a single-route connectionist system can accommodate the human data by simulating the regularity by frequency interaction found in human reaction time studies. But why does a neural network behave in this way? In essence, it seems that there is an initial effect of frequency and regularity; however, as the system learns and stabilises, type frequency for regular verbs is so high that token frequency does not appear to have any pronounced effect. Conversely, type frequency for irregulars remains relatively low and, in consequence, token frequency

continues to have a significant effect.

These findings therefore concur with the conclusion of Ellis & Schmidt, that the regularity by frequency interaction does not constitute evidence either *for* the 'rule-governed vs. rote-learnt' principle of dual-route models of inflectional morphology, or *against* single-system connectionist accounts. Importantly, and unlike Ellis & Schmidt's work, these results were generated:

- by a connectionist network trained on almost all monosyllabic real English verbs, rather than on 20 stimuli from an artificial language;
- using past tense morphology rather than pluralisation as the task being modelled; and
- with low-level phonological coding rather than arbitrary input-output representations.

7.2.4 Internal Representations

As has already been explained in detail in the previous chapter, Jaeger et al. (1996) claimed that the dual-route theory of inflectional morphology predicts that distinct brain areas would be involved in the processing of regulars and irregulars and, conversely, that a single-mechanism approach predicts identical patterns of activity in the processing of all verbs. Results of their PET study showed overlapping but distinct activation for regulars and irregulars (with more widespread activity elicited by irregular verbs), and they thus argued strongly against the connectionist approach,

and weakly in favour of dual-route theory.

Elman & Hare (1997) argued that Jaeger et al.'s findings "do not in fact distinguish between the hypotheses in the manner that is claimed". Elman & Hare's basis for this view was that a single-mechanism approach does *not* predict identical brain activity regardless of the type of verb being processed. To justify this assertion, they built a small (30 hidden unit) feedforward network, and trained it on 982 verbs (889 regulars, 93 irregulars) for 2000 epochs. They then tested the network on all non-homophonous forms and recorded the activations of the hidden units. By carrying out a regression analysis on these activations, they investigated whether the activity of the hidden units correlated with the type of verb being presented to the network. They found that the partial correlations, although small, were significant for 15 of the 30 hidden units. In other words, half of the hidden units seemed to be primarily active for either regular or irregular verbs, and half did not. They also found that more units seemed to be active while processing irregulars than during regular inflection, a result reflective of Jaeger et al.'s findings for humans.

From this, Elman & Hare argued that differential brain activity for different stimuli is expected and supported by dual-route *and* single-route theory, and that Jaeger et al.'s conclusions were thus unjustified. They concluded:

> Localization of data of the sort presented so far argue only that stimulus differences which produce differences in behavior must also produce differences in brain activity, and do not speak to (sic) the unitary or non-unitary nature of brain mechanisms.

In order to see whether their results could be found in a larger network, trained for

fewer epochs and on more verbs, a similar statistical analysis was conducted in this thesis, by studying the internal activations of one of the five network simulations (simulation 1 was chosen arbitrarily).

The weights matrix from simulation 1 (after 250 epochs of training) was loaded into the network, and all 1859 stem forms used in the training set were presented, once each, to the model. After each presentation of a stimulus, the activations for all 200 hidden units were recorded in a computer file, and the file was coded according to whether the stimulus had been regular or irregular. Once all 1859 verbs had been presented to the network, a spreadsheet was created consisting of 1859 rows (one per stimulus) and 201 columns (one per hidden node, plus one for the verb type code—regular or irregular). A regression analysis was then conducted on the data using SPSS version 8.0.

The software analysed to what extent each hidden node acted as a predictor of verb type. For example, if a particular hidden node were always highly active when a regular stimulus was presented at input, and always inactive when an irregular was presented (or vice versa), then that node could be said to be strongly correlated with verb type. If, on the other hand, a node were to show a similar level of activation regardless of whether an irregular or a regular was presented at input, then no significant correlation between that node and verb type would be found.

Every hidden node can therefore be assessed in terms of the correlation between its own activity, and the type of stimulus being presented. Regression analysis allows one—in essence—to divide the hidden nodes into three groups: those that tend to be

active only for regular stimuli; those that tend to be active only for irregular stimuli; and those that tend to be approximately equally active (or inactive) regardless of the type of verb being processed.

Significant partial correlations between hidden unit and verb type were found for 111 of the 200 hidden units (55.5%), a similar proportion to the correlations found in Elman & Hare's work. Of these 111 significant correlations, 56 were found to be correlated with irregular stimuli, and 55 with regulars. In other words, the proportion of units which were primarily active for regular verbs was approximately equal to the proportion of units which were primarily active for irregulars, the remainder of the hidden units being roughly equally active or inactive regardless of stimulus type. The full results of the analysis are shown in Tables N.1–N.4 in appendix N. In those Tables, for each node (shown in column one), unstandardised coefficients (B) are shown in column two, followed by their standard error (column 3), then standardised coefficients (β) in column 4, t-values and their significance in columns 5 and 6, and finally zero-order, partial and part correlations in columns 7, 8 and 9.

This analysis therefore replicates Elman & Hare's study in broad terms. The quantitative superiority of hidden units processing irregular verbs rather than regulars was not found; however, the results very much support the argument that a single-system mechanism does *not* predict identical patterns of activity during the processing of regular and irregular verbs.

7.2.5 Generalisation to Novel Forms

The networks were then tested on the 48 novel stimuli which had been used in the set of experiments with human subjects in chapter 5 above. Again, the stimuli were either HighI or LowI, and either HighR or LowR.

First, the stimuli were coded to a form which could be used as input to the network, using exactly the same method as for the training set (see section 7.1.3 above). The weights matrix for simulation 1 after 250 epochs was then loaded into the network, and the 48 stimuli were presented (in random order, without replacement) to the network. Outputs were decoded and then marked as regularised or irregularised according to the same criteria as were used for the human subjects in the earlier experiments. This process was repeated for the other 4 simulations.

Results for Pseudoirregular Stimuli

Results for pseudoirregular stimuli (for all networks) are shown in Figures 7.4 and 7.5. Figure 7.4 shows the combined total number of irregularisations made in each of the four verb type categories, for all five simulations (out of a potential maximum of 30 irregularisations in each category). Figure 7.5 shows the same data divided up for comparison between different network simulations—in this case, therefore, there was a possible maximum of 6 irregularisations per simulation in each category.

Looking at Figures 7.4 and 7.5, a trend of a similarity gradient is clear, from HighILowR stimuli to LowIHighR verbs. In order to test the significance of this trend, a two-factor repeated measures ANOVA was conducted on the results (F1), in

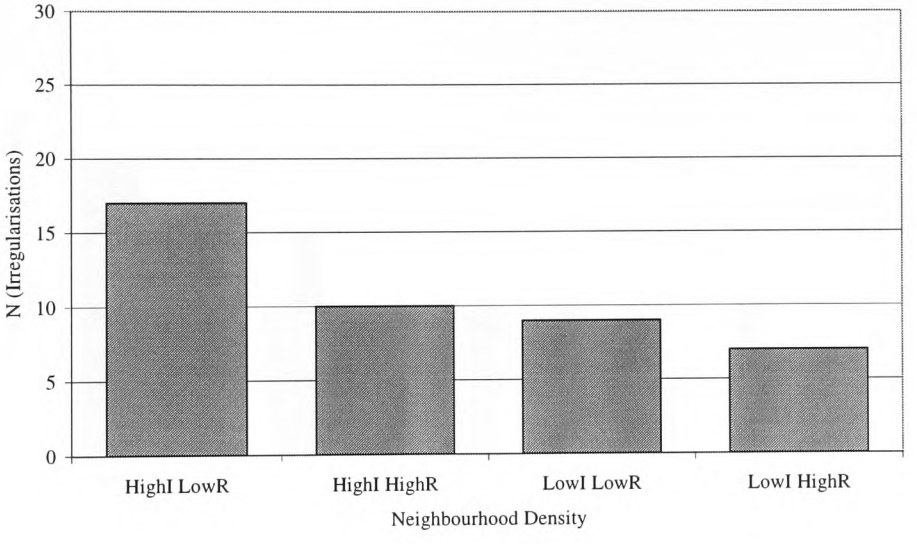


Figure 7.4: Total results for pseudoirregular verbs

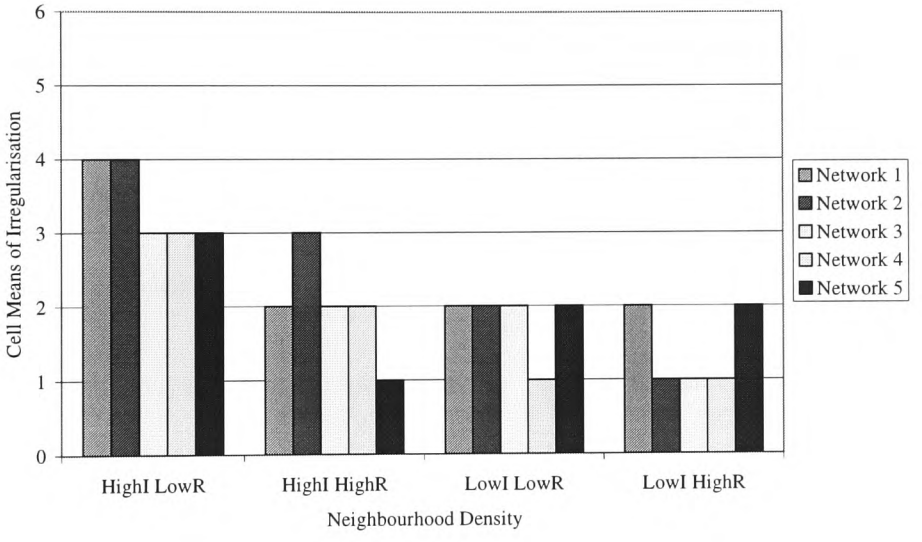


Figure 7.5: Results for pseudoirregular verbs by network

which Idensity (HighI or LowI) and Rdensity (HighR or LowR) were the two within-subjects factors. Each of the five simulations was treated as a separate subject for the purposes of the analysis. An items analysis was also conducted on the 24 pseudoirregular stimuli: a 2-factor repeated-measures ANOVA (F2), with Idensity and Rdensity as between-items factors.

A significant effect of Idensity was found, whereby HighI novel stimuli were more likely to be irregularised than were LowI stimuli ($F1(1,4) = 11.000$; $p = 0.0295$; $F2(1,20) = 6.135$; $p = 0.0223$). Rdensity effect was also significant, showing that LowR verbs were more likely to be irregularised than HighR verbs ($F1(1,4) = 81.000$; $p = 0.0008$; $F2(1,20) = 17.348$; $p = 0.0005$). No interaction was found.

Results for Pseudoregular Stimuli

Results for pseudoregular stimuli are shown in Figures 7.6 and 7.7. Figure 7.6 shows the total number of regularisations made by all five simulations in each of the four verb type categories (out of a possible 30 regularisations). The same results are shown in Figure 7.7, separated out by simulation. From these Figures it is immediately clear that a very high proportion of LowI stimuli were regularised (as shown by the two columns on the right in Figure 7.6); this is to be expected (since by definition LowI pseudoregulars contain no irregulars at all in their rhyming neighbourhood) and reflects the trends found in human data.

Again, a two-factor repeated-measures ANOVA was carried out on the results for the five simulations (F1), with Idensity and Rdensity as the 2 within-subjects factors.

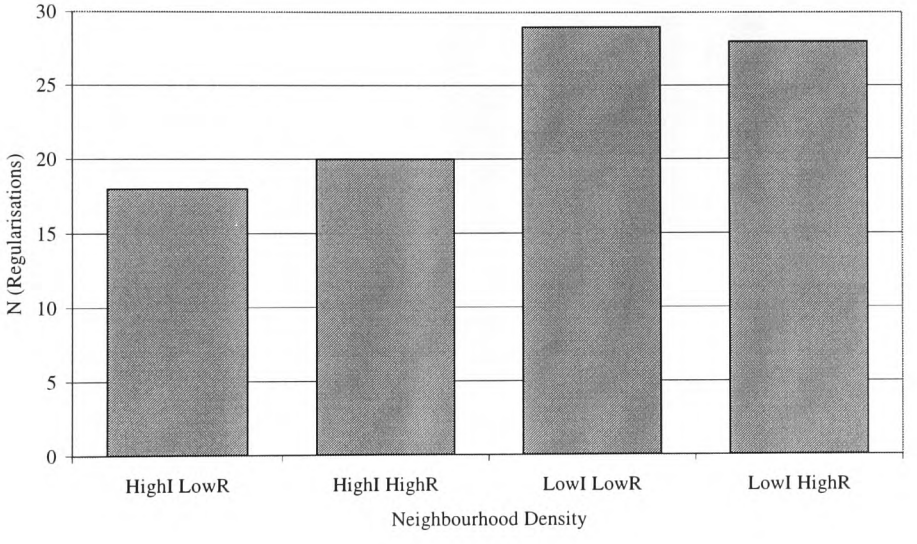


Figure 7.6: Total results for pseudoregular verbs

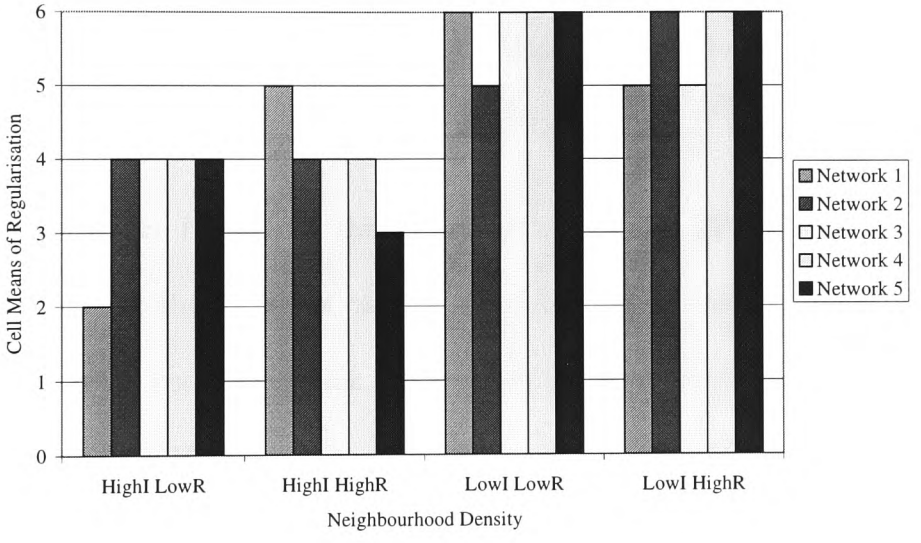


Figure 7.7: Results for pseudoregular verbs by network

A two-factor repeated-measures ANOVA was also conducted as an items analysis (F2), with Idensity and Rdensity as between-items factors.

A main effect of Idensity is clear ($F1(1,4) = 103.143$; $p = 0.0005$ in the subjects analysis), where LowI stimuli had a very strong tendency to be regularised ($\bar{x}$(regularised) $\approx$ 95%) compared to HighI stimuli ($\bar{x}$(irregularised) $\approx$ 63%). This was also found in the items analysis ($F2(1,20) = 10.939$; $p = 0.0035$). No other significant effects were found, either in the subjects analysis or in the items analysis. Of particular note is the fact that no effect of Rdensity was found.

Comparison with Human Data

Comparison with the Figures for the human data (on pages 117 *et seq.* above) shows that the results for the networks broadly mimicked the tendencies found in human subjects. That is, pseudoirregular verbs showed a gradual decrease in the rate of irregularisation from HighILowR to LowIHighR stimuli, with HighIHighR novel verbs being irregularised more than LowILowR stimuli. Furthermore, the results for pseudoregulars showed that HighILowR stimuli were the least likely to be regularised, then HighIHighR stimuli, and that both types of LowR novel verbs were highly likely to be regularised.

However, in order for more accurate comparisons to be made, separate ANOVAs were carried out on the pseudoirregular and pseudoregular stimuli, combining the data from the human and network experiments. For the analysis, the network data were compared with the data from the 'Auditory / No Context' experiment on humans,

to optimise the comparability between the network experiments and the human data in terms of methodology. Since the networks had been presented with verb stems out of any syntactic or semantic context, and since the stimuli had been represented phonologically, it made the most sense to compare the data with the 'Auditory / No Context' human experimental results, rather than with visually-presented stimuli, or stimuli presented in sentential context. Each network was treated as an individual subject in the analyses.

For both pseudoregular and pseudoirregular stimuli, a three-factor (2 within-subjects, 1 between-subjects) repeated-measures ANOVA was conducted on the human and network data together, where the 2 within-subjects factors were Idensity and Rdensity, and the between-subjects factor was subject class (coded as 1 for a human and 0 for a network).

In the **pseudoirregular** 'humans vs. networks' ANOVA, main effects of Idensity and Rdensity were found ($F(1,23) = 80.453$; $p < 0.0001$ and $F(1,23) = 19.374$; $p = 0.0002$ respectively), where HighI novel verbs were more likely to be irregularised than LowI stimuli, and LowR stimuli were more likely to be irregularised than HighR. More importantly, it was found that humans and networks behaved differently when presented with the same stimuli ($F(1,23) = 12.479$; $p = 0.0018$). The source of this difference was found in a 2-way Idensity $\times$ subject class interaction ($F(1,23) = 14.777$; $p = 0.0008$). Here it was found that although both humans and connectionist models tended to irregularise HighI stimuli more than LowI stimuli, humans did so more often than networks. This interaction is depicted in Figure 7.8. No other effects

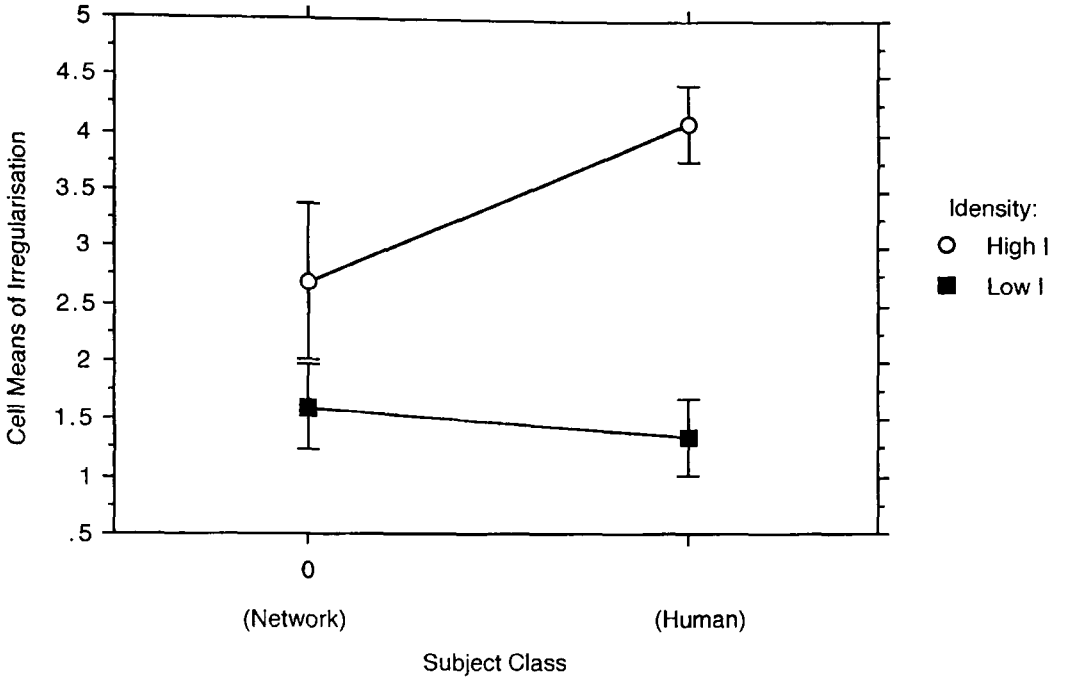


Figure 7.8: Idensity × subject class interaction for pseudoirregulars

or interactions were found.

In the **pseudoregular** 'humans vs. networks' ANOVA, a main effect of subject class was found, whereby networks were much more likely than human subjects to regularise pseudoregular stimuli, regardless of neighbourhood density ($F(1,23) = 45.067$; $p < 0.0001$). In fact, the mean percentage of regularisations by the networks was approximately 79%, compared to approximately 57% by human subjects. A main effect of Idensity was also found, where HighI novel verbs were less likely to be regularised than LowI stimuli ($F(1,23) = 140.375$; $p < 0.0001$); however, unlike the original human data, no main effect of Rdensity was found. Finally, there was an interaction between Idensity × Rdensity, showing that Rdensity has an impact on regularisation, but only for HighI (rather than LowI) stimuli ($F(1,23) = 7.926$;

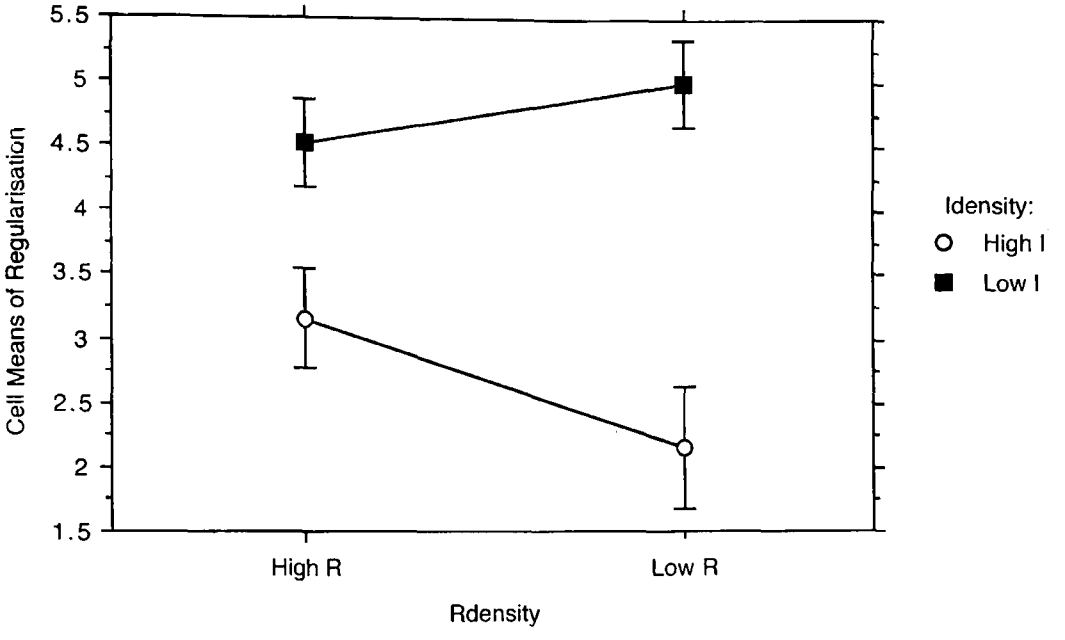


Figure 7.9: Idensity × Rdensity interaction for pseudoregulars

$p = 0.0098$). This interaction is depicted in Figure 7.9.

In summary, these analyses found that the networks' novel inflections showed effects of Idensity (for pseudoirregular and pseudoregular stimuli) and Rdensity (for pseudoirregular stimuli only). It is interesting to note that no effect of Rdensity on pseudoregular stimuli was found in the networks, given that:

1. such an effect might be expected in an analogy-based system; and

2. the (debatable) lack of regular neighbourhood effects in humans has been cited as evidence against connectionist networks, on the basis that connectionist models would not show such a lack (Prasada & Pinker, 1993).

When compared directly with the human data from chapter 5, it was found that humans and networks both tended to irregularise more HighI than LowI pseudoirreg-

ulars, humans did so more frequently. Humans were also more likely to regularise the pseudoregular stimuli, regardless of neighbourhood factors.

7.3 Discussion

All five network simulations described in this chapter were successful at modelling the past tense inflections of English monosyllabic verbs. Regular verbs tended to be learnt faster than irregulars but, within 200 epochs of training, each simulation produced correct responses to all but a handful of (very low-frequency) stimuli.

When tested on 12 real (regular and irregular) verb stems which had been omitted from the training corpus, all 5 simulations inflected the previously-unseen stimuli correctly.

When high- and low-frequency irregulars and regulars were compared with one another, a main effect of regularity was found, where irregular stimuli elicited a lower RMS error score than regulars; a main effect of frequency was also found, where high-frequency stimuli produced lower error scores than low-frequency stimuli. A two-way interaction between regularity and frequency showed that frequency had a significant effect on irregular verbs, but not on regulars. Although a lower error score was found for irregulars than for regulars, contrary to the regularity effect found by Ellis & Schmidt (1998), the findings nonetheless support their conclusions (and those of Daugherty & Seidenberg, 1992) that the regularity by frequency interaction, found in human data, can be modelled successfully by single-system connectionist networks, and should therefore not be taken as evidence in favour of the symbolic dual-route

theory of past tense morphology.

A regression analysis, conducted on the activations of the network's hidden units (in other words, on the 'internal representations' of the network), found that over half of the 200 hidden units were significantly partially correlated with the class (regular or irregular) of each stimulus. This finding supports the conclusion of Elman & Hare (1997) that connectionist networks do not predict identical patterns of activity regardless of the type of verb being inflected, and therefore that neurophysiological studies such as that of Jaeger et al. (1996) are wrong to distinguish between dual-and single-route theories simply on the basis of differential brain activity for regular and irregular verbs.

When tested on 48 novel verb stimuli, it was found that the network tended to irregularise HighI pseudoirregulars more than HighR pseudoirregulars. In other words, the number of regulars and the number of irregulars in the phonological neighbourhood of a pseudoirregular stimulus both had an effect on the probability of irregularisation for that stimulus. With pseudoregular stimuli, a clear effect of Idensity was found, where LowI stimuli were much more likely than HighI stimuli to be regularised, at a rate of approximately 95% (as opposed to about 63% for HighI stimuli). It will be recalled that, for pseudoregulars, 'LowI' means that the stimulus has *no* irregular neighbours in its rime neighbourhood, whereas HighI means that there are 1 or more rhyming irregulars in the neighbourhood of the novel stimulus. However, no effect of Rdensity was found, as might have been expected from human data (as described in chapter 5).

The networks' results for novel verbs were then compared statistically with the data for humans tested on the same stimuli. Although the broad trend of results for different verb types was found to be similar for networks and humans (compare Figure 5.1 on page 117 with Figure 7.4 on page 243, and Figure 5.3 on page 120 with Figure 7.6 on page 246), an interaction between Idensity and subject class showed that humans were much more likely than networks to irregularise Highl stimuli. Furthermore, when human and network data for pseudoregulars were compared, a main effect of subject class showed that networks were significantly more likely than humans to regularise pseudoregular stimuli.

7.4 Conclusion

In this chapter, it has been shown how a simple feedforward connectionist network was designed and constructed to map phonological representations of monosyllabic English verb stem forms onto phonological representations of their past tense inflected forms. It successfully modelled the 'regularity × frequency' interaction found in human behavioural data, and showed significant partial correlations between its internal state and the type of verb being processed. The networks' reactions to novel verb stimuli broadly paralleled earlier behavioural findings for humans, although significant differences were found. Despite these differences, the network accounted for much of the human data, and was particularly competent at modelling findings which have been used in the literature to argue *against* subsymbolic single-system theories of English past tense inflectional morphology.

Chapter 8

General Conclusion

8.1 Summary

This thesis has explored key issues in the debate between single- and dual-route theories of inflectional morphology, from the perspective of three different fields of cognitive science: experimental psychology, neurophysiology and artificial neural network modelling. Within these fields, factors of regular and irregular neighbourhood density have been controlled using novel and real verb stimuli, and found to have implications for the two dominant theories of past tense processing.

In chapter 1 some basic properties of verbs and early theories of verbal processing were described; chapter 2 explained the principles underlying parallel distributed processing; and chapter 3 covered the experiments and arguments in the inflectional morphology debate from the mid-1980s up to the mid-1990s.

One critical issue in the single-route/dual-route debate is that each theory makes different predictions about the impact that the number of regular verbs in the rhyming neighbourhood of a novel stimulus may have on its inflection. Steven Pinker's hybrid dual-route theory proposed two processing routes, one for regulars and one for irregulars. According to this theory, verbs are processed by an explicit, symbolic, default suffixation rule unless blocked by a match (or sometimes a partial match) with a known irregular in a separate associative mechanism. The model predicts, therefore, that irregular neighbourhood density will affect the inflection of a novel verb, but that the number of similar regular verbs will not have an effect. Conversely, connectionist single-route theories predict that, since verbs are processed by analogy to all

known items, both irregular *and* regular neighbourhood density will have an impact on novel verb inflection.

In particular, Prasada & Pinker (1993) and Lee (1995) attempted to test these conflicting predictions using novel verb elicitation tasks. Prasada & Pinker found no evidence for an effect of regular neighbourhood density, and concluded that their human data therefore supported dual-route theory; Lee, by contrast, found a clear effect of the number of regular verbs in the rime neighbourhood of novel stimuli, and concluded that his data supported single-route explanations of past tense inflection.

In chapter 4, the specific systems of neighbourhood classification used by Prasada & Pinker and Lee were discussed, analysed and found to be flawed. A new way of classifying neighbourhood density was proposed, and it was argued that this quantitative classification, in which a verb is assigned a measure of irregular and regular density in comparison with *all* verbs in the language, provides a more objective and methodologically sound basis for stimulus selection and control. This new classification formed the basis for the three major experimental components of the thesis.

In chapter 5, a set of four behavioural elicitation experiments was described, in which adult subjects were required to inflect stem forms of novel verbs to appropriate past tense forms. Stimuli were controlled for irregular and regular neighbourhood density, and it was found that both irregular and regular neighbourhood structure affected past tense processing: the probability of regularisation was greater for LowI than HighI verbs, and was also greater for HighR than for LowR stimuli. This finding was shown to conflict with Pinker's hybrid dual-route theory of inflectional morphol-

ogy, and to support single-route connectionist theory.

Chapter 6 described two ERP studies. The first used real verbs as stimuli to be inflected to the past tense. Statistical analysis found substantial differences in activation between responses to regular and irregular stem forms, thereby replicating Breyer (1997). Subset analyses found that differences in regular and irregular neighbourhood density each had noticeable effects on the inflection of both regular and irregular verbs.

The second ERP experiment recorded EEG from subjects while they were inflecting *novel* stimuli to appropriate past tense forms. Small frontal and lateral effects were found when the waveforms for regularised and irregularised stimuli were compared statistically, but otherwise the elicited waveforms were remarkably similar. However, when subset analyses were conducted on the regularised novel stimuli, it was found that the number of irregulars in the rime neighbourhood of a regularised stimulus had a marked effect on its waveforms but, conversely, that the number of *regulars* in its neighbourhood had almost no impact at all. These findings would seem to be more easily accounted for by dual-route than single-route theories.

Waveforms for regularised novel verbs were then compared with those for real regular and irregular verbs. Substantial differences were found after the main language-related potential peak, showing that novel stimuli tended to evoke less positive potentials than regular or irregular real verbs. Particularly significant effects were found in posterior areas.

In sum, the two ERP studies yielded a considerable amount of new data which

provide a detailed analysis of the effects of regular and irregular neighbourhood structure on the potentials elicited during real and novel past tense verb processing.

Finally, chapter 7 described a set of five artificial neural networks which were trained on the English past tense task. These networks successfully learnt to map verb stems onto their past tense forms. Behavioural effects which have been cited as evidence for a dual-route system of inflectional morphology (such as the 'regularity × frequency' interaction and internally-dissociable neural activation) were successfully modelled using a single-route methodology. Furthermore, the networks showed substantial similarities to the human data (from chapter 5) in their responses to novel verb stems.

Collectively, these three sets of studies—behavioural, electrophysiological and connectionist—have produced a considerable quantity of data, which provide detailed findings with regard to regular and irregular neighbourhood density. In particular, several new findings for novel verbs have been reported. Critically, some of these results provide evidence against Pinker's hybrid dual-route theory, which predicts that regular neighbourhood density has no effect on past tense inflection; on the contrary, these data, which showed that irregular *and* regular neighbourhood effects were found when inflecting verbs to the past tense, tend to support the single-route perspective of past tense morphology.

8.2 Future Work

The new classification of neighbourhood density (defined in chapter 4) has provided some significant results. However, there is certainly room to improve the structure further: in particular, future development could—perhaps, should—incorporate token frequency of a verb's neighbours into its quantification. In this way, for example, proximity to a high-frequency verb would lead to a greater density than proximity to a low-frequency verb. One difficulty in executing such a design would be the necessary subjectivity involved in the equations: how, for instance, should a verb with three low-frequency regular neighbours compare to a verb with one high-frequency regular neighbour? It is interesting to see that, even without taking frequency of neighbours into account, the classification produced neighbourhood effects of very large statistical significance[1]. Nonetheless, controlling for frequency would be a beneficial next step in developing a rigorous system of neighbourhood classification.

The ERP studies described in chapter 6 were the first of their kind to investigate neighbourhood effects on real and novel verb inflection at an electrophysiological level. As explained in chapter 6, ERP imaging lacks spatial precision but is temporally accurate, while the reverse is true of PET techniques. Future research in this area may produce more accurate results if new methodologies—such as magnetoencephalography—are employed, whereby the temporal precision of ERPs is matched by the spatial accuracy of PET. However, further ERP studies would also

[1] F-values of over 400, in several cases.

yield much useful new information. For instance, in the novel verb ERP study in this thesis, not enough novel stimuli were irregularised for subset analyses (of neighbourhood effects on irregularised novel verbs) to be statistically significant. Future experiments could rectify this by presenting many more novel stimuli to more subjects: results of such analyses would help to clarify the presence or absence of regular (and irregular) neighbourhood effects on irregularised novel stimuli.

It would also be useful to experiment with different designs and learning strategies of connectionist networks, to try to reduce the differences between the human and network data found in this thesis. Factors to manipulate might include the number of hidden nodes used in the networks, the log frequency function (see page 218 above) and the phonological encoding/decoding system used. On a more general level, as connectionist models develop in complexity, they may begin to counteract the oft-levelled criticism that they are overspecific, dealing with one small area of a complex field in order to simplify a problem to modellable proportions.

Since the experimental components of this thesis were conducted, Pinker's (1991) hybrid dual-route theory has been amended in response to criticisms that it fails to account for some behavioural findings (similar to those found in this thesis) that regular past tenses do not appear simply to be rule-generated. In his most recent book, 'Words and Rules', Pinker (1999) suggested that some regulars are stored in memory, presumably alongside irregulars in the 'pattern associator' route of the model. Pinker supposed that regulars which rhyme with irregulars are stored in memory, as are high-frequency regulars, and this argument has support from other theorists

(Marcus, *pers. comm.*). This change in approach has blurred the distinction between regular and irregular past tense processing, in contrast to the clearly delineated routes of earlier theories. Certain specifics of the new theory must be defined before testable hypotheses can be generated: for instance, how high-frequency must a regular verb be before inclusion in memory? Or how similar must it be to known irregulars? Once such theoretical questions have been answered, psychologists and linguists will embark on the next phase of what Pinker (1999) called "one of the most vigorous controversies in the modern study of the mind".

Appendix A

Irregular Verbs — List α

abide–abode	dwell–dwelt	mean–meant	spill–spilt
be–was	eat–ate	meet–met	spin–spun
bear–bore	fall–fell	put–put	spit–spat
beat–beat	feed–fed	read–read	split–split
begin–began	feel–felt	rend–rent	spoil–spoilt
bend–bent	fight–fought	rid–rid	spread–spread
bereave–bereft	find–found	ride–rode	spring–sprang
beseech–besought	flee–fled	ring–rang	stand–stood
bet–bet	fling–flung	rise–rose	steal–stole
bid–bade	fly–flew	run–ran	stick–stuck
bid–bid	forsake–forsook	say–said	sting–stung
bind–bound	freeze–froze	see–saw	stink–stank
bite–bit	get–got	seek–sought	stride–strode
bleed–bled	give–gave	sell–sold	strike–struck
blow–blew	go–went	send–sent	string–strung
break–broke	grind–ground	set–set	strive–strove
breed–bred	grow–grew	shake–shook	swear–swore
bring–brought	hang–hung	shed–shed	sweep–swept
build–built	have–had	shine–shone	swim–swam
burn–burnt	hear–heard	shit–shat	swing–swung
burst–burst	heave–hove	shoot–shot	take–took
buy–bought	hide–hid	shrink–shrank	teach–taught
cast–cast	hit–hit	shrive–shrove	tear–tore
catch–caught	hold–held	shut–shut	tell–told
chide–chid	hurt–hurt	sing–sang	think–thought
choose–chose	keep–kept	sink–sank	thrive–throve
cleave–clove	kneel–knelt	sit–sat	throw–threw
cling–clung	knit–knit	slay–slew	thrust–thrust
come–came	know–knew	sleep–slept	tread–trod
cost–cost	lead–led	slide–slid	wake–woke
creep–crept	lean–leant	sling–slung	wear–wore
cut–cut	leap–leapt	slink–slunk	weave–wove
deal–dealt	learn–learnt	slit–slit	wed–wed
dig–dug	leave–left	smell–smelt	weep–wept
dive–dove	lend–lent	smite–smote	wet–wet
do–did	let–let	speak–spoke	win–won
draw–drew	lie–lay	speed–sped	wind–wound
dream–dreamt	light–lit	spell–spelt	wring–wrung
drink–drank	lose–lost	spend–spent	write–wrote
drive–drove	make–made		

Table A.1: British English irregular verbs—list α

Appendix B

Irregular Verbs — List β

abide–abode	countersink–countersank	heave–hove
arise–arose	creep–crept	hide–hid
awake–awoke	crossbreed–crossbred	hit–hit
backbite–backbit	cut–cut	hold–held
backslide–backslid	deal–dealt	hurt–hurt
be–was	dig–dug	inbreed–inbred
bear–bore	dive–dove	indwell–indwelt
beat–beat	do–did	inset–inset
become–became	draw–drew	interbreed–interbred
befall–befell	dream–dreamt	interweave–interwove
beget–begot	drink–drank	keep–kept
begin–began	drive–drove	kneel–knelt
behold–beheld	dwell–dwelt	knit–knit
bend–bent	eat–ate	know–knew
bereave–bereft	fall–fell	lead–led
beseech–besought	feed–fed	lean–leant
beset–beset	feel–felt	leap–leapt
bespeak–bespoke	fight–fought	learn–learnt
bestride–bestrode	find–found	leave–left
bet–bet	flee–fled	lend–lent
betake–betook	fling–flung	let–let
bethink–bethought	floodlight–floodlit	lie–lay
bid–bade	fly–flew	light–lit
bid–bid	flyblow–flyblew	lose–lost
bind–bound	forbear–forbore	make–made
bite–bit	forbid–forbade	mean–meant
bleed–bled	forecast–forecast	meet–met
blow–blew	foreknow–foreknew	miscast–miscast
break–broke	foresee–foresaw	misdeal–misdealt
breastfeed–breastfed	foretell–foretold	misgive–misgave
breed–bred	forget–forgot	mishear–misheard
bring–brought	forgive–forgave	mishit–mishit
broadcast–broadcast	forgo–forwent	mislead–misled
browbeat–browbeat	forsake–forsook	misread–misread
build–built	forswear–forswore	misspell–misspelt
burn–burnt	freeze–froze	misspend–misspent
burst–burst	gainsay–gainsaid	mistake–mistook
buy–bought	get–got	misunderstand–misunderstood
cast–cast	give–gave	offset–offset
catch–caught	go–went	outbid–outbid
chide–chid	grind–ground	outdo–outdid
choose–chose	grow–grew	outfight–outfought
cleave–clove	hamstring–hamstrung	outgo–outwent
cling–clung	hang–hung	outgrow–outgrew
come–came	have–had	outride–outrode
cost–cost	hear–heard	outrun–outran

Table B.1: British English irregular verbs—list β (Part 1)

outsell–outsold	send–sent	sweep–swept
outshine–outshone	set–set	swim–swam
overbid–overbid	shake–shook	swing–swung
overcome–overcame	shed–shed	take–took
overdo–overdid	shine–shone	teach–taught
overdraw–overdrew	shit–shat	tear–tore
overfly–overflew	shoot–shot	telecast–telecast
overhang–overhung	shrink–shrank	tell–told
overhear–overheard	shrive–shrove	think–thought
overleap–overleapt	shut–shut	thrive–throve
overlie–overlay	sightsee–sightsaw	throw–threw
override–overrode	sing–sang	thrust–thrust
overrun–overran	sink–sank	tread–trod
oversee–oversaw	sit–sat	typecast–typecast
oversell–oversold	slay–slew	typeset–typeset
overshoot–overshot	sleep–slept	typewrite–typewrote
oversleep–overslept	slide–slid	unbend–unbent
overtake–overtook	sling–slung	unbind–unbound
overthrow–overthrew	slink–slunk	underbid–underbid
partake–partook	slit–slit	undercut–undercut
preset–preset	smell–smelt	undergo–underwent
proofread–proofread	smite–smote	underlie–underlay
put–put	speak–spoke	undersell–undersold
read–read	speed–sped	undershoot–undershot
rebind–rebound	spell–spelt	understand–understood
rebuild–rebuilt	spellbind–spellbound	underwrite–underwrote
recast–recast	spend–spent	undo–undid
redo–redid	spill–spilt	unlearn–unlearnt
rehear–reheard	spin–spun	unsay–unsaid
remake–remade	spit–spat	unstring–unstrung
rend–rent	split–split	unwind–unwound
rerun–reran	spoil–spoilt	uphold–upheld
reset–reset	spoonfeed–spoonfed	upset–upset
retake–retook	spread–spread	wake–woke
retell–retold	spring–sprang	wear–wore
rethink–rethought	stand–stood	weave–wove
rewrite–rewrote	steal–stole	wed–wed
rid–rid	stick–stuck	weep–wept
ride–rode	sting–stung	wet–wet
ring–rang	stink–stank	win–won
rise–rose	stride–strode	wind–wound
roughcast–roughcast	strike–struck	withdraw–withdrew
run–ran	string–strung	withhold–withheld
say–said	strive–strove	withstand–withstood
see–saw	sublet–sublet	wring–wrung
seek–sought	swear–swore	write–wrote
sell–sold		

Table B.2: British English irregular verbs—list β (Part 2)

Appendix C

Test Passage Used in Behavioural Studies

When the bell rang, Richie _____ his bag quickly and rushed out of the classroom. In his haste, he _____ a cable and fell flat on his face. Fortunately for him, there was a mat right in front of him and it cushioned his fall. He picked himself up and _____ to the field to play. He was particularly fond of hockey. He took out his stick and _____ the ball all over the field.

Appendix D

Cloze Passage Used in Behavioural Studies

Tom, the cat, knew what to do. He _____ along the garden path till he found what he was looking for. Apparently what he was looking for was a delicious mouse under the garden table, with her tail _____ in the mousetrap. His mistress, a certain Mrs. Beatrice Potter, had left the trap under the table the previous night, and Tom knew where he could find a mouse without trying too hard. The mouse _____ fearfully at the advancing cat, knowing the end of her life was near. Glancing around, the mouse quickly _____ a nearby piece of cheese and offered it to Tom, hoping to entice Tom with the cheese. As Tom advanced closer to the mouse, it became increasingly clear to him that the mouse was trying to distract him with the piece of cheese. "Should I take the cheese offering or gobble up the mouse?" Tom _____ to himself.

Should Tom take the cheese offering or gobble up the mouse? the suspense was killing. Will Macy, the mouse, be able to get out of this predicament? Tom _____ at Macy and her cheese offering one last time and made up his mind. Suddenly, Tom _____ towards Macy. Macy braced herself for the worst. If the cat chose to take the cheese-offering, she would be safe for at least the short while that the cat _____ on the cheese and she would have time to plan her next move. If the cat _____ to go for her, this would be her last breath. All these thoughts _____ through her mind as she looked at Tom's shadow looming over her. And then suddenly .... THUMPP!!!

Tom missed. He actually missed the mouse ... and the cheese as well! Macy _____ her lucky stars and said a quick prayer of thanks to

God as well. "Heh, heh, scared you, didn't I?" Tom _____ , "You didn't think for a moment I would miss, did you?" "Oh no," Macy _____ to herself, "not a sadistic cat, please. I'd rather die in one mouthful than be stripped of my fur and flesh in stages." Suddenly, Tom _____ his right paw at Macy and left four scratches on her cheek. Blood _____ out slowly down her delicate left cheek. "God, have mercy on me," Macy cried out in pain. "Good heavens, a talking mouse!" Tom exclaimed in surprise. Tom didn't think that a lower life form could talk like a cat can. He knew that his mistress could babble in some unknown language, but certainly not a mouse.

What was Tom supposed to do with a talking mouse? He can't eat a life-form with a language as advanced as his. Sure, humans can talk; but their language doesn't have the economy of sounds and words a cat's language has. Certainly, human speech was nowhere near as concise as a cat's miaow or a mouse's squeak. Macy's thoughts _____ furiously through her mind. What was she to do next? The cat was obviously impressed by her speech, but would this save her? "Wait," Macy tried her luck, "please don't eat me; I can tell you where the other mice are." Macy wasn't sure if this would save her life but that was all she could think of in that short span of time.

Tom _____ around to face the mouse, his head resting on his front paws, his hind propped up, ready to dive onto the mouse at any moment. He _____ his left paw forward, and then his right. By this time, he was no more than two feet from Macy. Macy _____ uncontrollably. What next? What next? She _____ Tom opening his mouth, as though stretching it as wide as it would take to accommodate her in one gulp. It was as though Tom did this slowly so that Macy could count the number of sharp canines in his mouth. She felt Tom's breath engulfing her and was about to faint from fear when suddenly Tom let out a powerful yelp and threw Macy and the mousetrap onto the pavement.

"How stupid of you," Bob _____ Tom. Apparently, Bob had hit Tom so hard with a broom that he let out a loud yelp. Bob was Tom's cousin. "Let the mouse lead us to the rest of them; that way we can have all the mice we want for the rest of the year." Sensing her chance of wriggling out of this predicament, Macy _____ her head vigorously at Bob's _____ suggestion. "I will show you where the rest of the mice are if you let me go," Macy _____ to her tail which was still trapped. "Yes, but if you lie to us, I promise you a

death so slow and painful you'll wish you hadn't been born. Macy couldn't even begin to imagine the kind of pain she would have to endure, but she knew that she needed more time to plan her escape. Even if she had led them to the other mice, she _____ they would release her anyway. She shook her head frantically as though this would help her to think better. She reached into the depths of her mind and suddenly _____ up this fabulous plan.

She told the cats that in order for her to lead them to the mice's hide-out, they first had to release her so that she could lead the way. She _____ the cats that even if she tried to run and escape, they would still be able to out-run her easily since she was already injured. She added that she was the only mouse who was able to use the cats' language, and would therefore be a valuable asset to them. The cats _____ Macy's proposal and decided that she was incapable of out-running both of them. So they set Macy free and allowed her to lead the way, but _____ that if she tried to trick them, she and her family would die a tragic death. Macy wasn't sure what she was supposed to do next, but at least she was set free from the trap. There was no way she could out-run the cats, but she _____ that the only way out of this predicament was to out-think them; and this was not too hard to do, considering the cats' evaluation of their own intelligence. She _____ around the garden, thinking of a possible trick.

Over on the other side of the fence, some grasshoppers were jumping about in the field, basking in the first rays of the Spring sun, oblivious to her predicament. Macy _____ over her shoulder and found the cats walking very closely behind her. How was she to escape? As they approached the shed at the southern end of the garden, she realised that she was running out of time and space. Suddenly, an idea _____ to her mind. In her excitement, she _____ to her left a little more abruptly than she ought to. This immediately _____ the cats as she could see the cats pulling back their upper lips, exposing their sharp fangs. She turned to face the eastern end of the garden and _____ slowly while she worked out the details of her plan to deceive the cats.

As they approached the eastern end of the garden, they saw the outline of what looked like a well. Very few animals dared to venture near the old well as legend had it that it was in fact the mouth of an evil beast. Objects like tree branches and leaves had been known to fall into it, never to be seen again. Macy had never approached the well

before out of fear, but this was a desperate situation calling for drastic measures. As they approached the mouth of the dreaded beast, they could hear a low and strange sound coming from it, as though it approved their approach. Perhaps it was the rumbling of its stomach as its digestive juices _____ in anticipation. (They didn't realise that it was merely the howling of the wind in the well.) It became increasingly clear that the cats were just as frightened as Macy. "Why are you leading us to the dreaded creature on the eastern edge of the garden?" the cats _____ loudly, trying to hide their fear. "That's where the mice are; they hide in the mouth of the beast." "Oh, why doesn't the beast eat them up then?" "It doesn't fancy mice and cats," was Macy's feeble attempt at a logical reason. "Hmm," said the cats, as though understanding. Macy _____ the bucket hanging over the well and said, "Look, its tongue is not moving, it must be asleep. We can sneak in quietly and find the rest of the mice." The two cats glanced at each other, _____ the proposal for a moment and then said, "If this is some trick of yours, we will rip your stomach apart and tear your guts out." Macy was beginning to wonder which was worse, being swallowed by the beast or ripped apart by the cats. She _____ at the thought of either consequence.

When they were next to the sleeping beast, Macy _____ this trick to get the stupid cats to jump into the mouth. "I will jump in first, find the mice and then lead you to them," Macy suggested as she peered down the beast's throat. "Do you think we are stupid or what? We're not going to let you jump in first and warn the rest of the mice. Tom will jump down the throat first and then you will follow. I will jump in after you," Bob _____ confidently. Before Tom could protest, Bob _____ Tom over the ledge and into the well. Both Bob and Macy _____ into the well and waited for Tom to confirm the mice's hideout. They couldn't see what was happening below because it was a very deep well. After a few seconds, they heard what sounded like a splash.

Sensing the cat's concern, Macy said "Maybe Tom has found the mice and decided to keep them for himself." This remark made Bob so angry that he _____ into the well immediately, in search of the mice. But Macy knew better: they would never survive the beast. Macy couldn't believe her luck. One mousetrap, two cats and yet she was still alive. Maybe if she started jumping about to wake the beast up, it would swallow the two cats up for good, she thought to herself. So, Macy _____ to jump about on the ledge of the well. She made such

a commotion that she could feel the beast rumbling from its sleep. She was about to jump onto the grass when she _____ a strong gust of air sucking her into the well. The draught was so strong that she lost her balance and _____ into the well. Two stupid cats, one arrogant mouse and one mousetrap, the beast said to itself. Yum!

Appendix E

Real Verb ERP Stimuli

Irregular	Density	Frequency	Regular	Density	Frequency
BIND	4-0/8-0-4	00004	FLIP	0-0-32	00004
CLING	7-4/9-4-4	00006	SHAVE	0-0-15	00006
CUT	2-0/3-0-10	00192	BOOK	0-0-10	00193
DRINK	4-2/5-4-11	00082	TRAIN	0-0-52	00082
DRIVE	5-0/5-0-10	00105	SERVE	0-0-12	00107
FIND	4-0/8-0-4	00399	LOOK	0-0-10	00399
FLING	7-4/9-4-4	00002	MOURN	0-0-15	00002
GRIND	4-0/8-0-4	00002	PLUCK	0-0-11	00002
KEEP	6-0/8-0-7	00264	MISS	0-0-23	00258
RID	2-1/5-2-2	00019	ROB	0-0-10	00019
SHAKE	3-3/8-5-13	00017	QUOTE	0-0-21	00017
SHED	3-1/3-1-12	00011	WARN	0-0-15	00011
SHRINK	4-2/5-4-11	00005	CLASH	0-0-20	00005
SHUT	2-0/3-0-10	00046	WAVE	0-0-15	00046
SLEEP	6-0/8-0-7	00065	JOIN	0-0-10	00065
SPREAD	3-1/3-1-12	00083	MARK	0-0-16	00083
STRIVE	5-0/5-0-10	00007	BLAZE	0-0-18	00007
SWING	7-4/9-4-4	00024	JUMP	0-0-14	00024
TAKE	3-3/8-5-13	00611	USE	0-0-19	00589
WEEP	6-0/8-0-7	00014	BARK	0-0-16	00014

Table E.1: High friend / low enemy real stimuli

Irregular	Density	Frequency	Regular	Density	Frequency
BITE	2-3/4-7-30	00010	RUST	1-1-11	00010
BLEED	6-0/14-0-21	00002	DRAWL	1-2-29	00002
FEED	6-0/14-0-21	00123	EYE	3-6-122	00122
GROW	4-1/8-3-56	00063	KILL	1-1-18	00063
HIT	4-2/5-2-89	00115	STAY	2-4-67	00113
KNOW	4-1/8-3-56	00683	LAST	1-8-8	00676
LEAVE	2-3/2-4-20	00205	PLAY	1-1-18	00200
LEND	5-0/7-0-36	00014	PINE	1-2-42	00014
LIGHT	2-3/4-7-30	00333	NEED	6-14-21	00360
SEND	5-0/7-0-36	00074	DIE	3-6-122	00073
SIT	2-4/2-5-89	00067	FLOW	5-11-56	00067
SPIN	2-1/2-1-26	00005	HAUL	1-2-29	00005
SPIT	2-4/2-5-89	00011	SIGH	3-6-122	00011
SPLIT	4-2/5-2-89	00030	ROOT	1-3-41	00030
SWEAR	4-0/6-0-27	00010	THUMB	1-3-13	00010
TEAR	4-0/6-0-27	00011	RAKE	6-13-13	00011
THROW	4-1/8-3-56	00042	EASE	1-1-10	00042
WEAVE	3-2/4-2-20	00004	BLINK	6-9-11	00004
WIN	2-1/2-1-26	00055	PICK	1-1-23	00055
WRITE	2-3/5-6-30	00106	REACH	2-2-11	00106

Table E.2: High friend / high enemy real stimuli

Irregular	Density	Frequency	Regular	Density	Frequency
COME	1-0/3-0-13	00630	WANT	0-0-1	00329
BUILD	1-0/2-0-1	00086	TOUCH	0-0-2	00087
BURST	1-0/1-0-2	00033	REV	0-0-1	00033
CATCH	1-0/1-0-12	00043	SMOOTH	0-0-2	00042
COST	1-0/1-0-5	00229	LOVE	0-0-2	00232
DIG	1-0/1-0-6	00010	LAUNCH	0-0-2	00010
FREEZE	1-0/1-0-10	00006	WOLF	0-0-1	00006
GIVE	1-0/3-0-5	00391	FIELD	0-0-4	00274
HANG	1-0/2-0-9	00026	WOUND	0-0-1	00028
HOLD	1-0/4-0-9	00169	EARTH	0-0-3	00150
RUN	1-0/4-0-5	00212	VOICE	0-0-3	00226
SEEK	1-1/1-2-15	00069	KNIFE	0-0-1	00076
PUT	1-0/1-0-4	00437	HELP	0-0-3	00311
SPEAK	1-1/2-1-15	00110	ASK	0-0-5	00128
SPOIL	1-0/1-0-14	00003	CARVE	0-0-3	00003
STAND	1-0/4-0-10	00148	CAUSE	0-0-2	00130
STRIKE	1-0/1-0-8	00050	PULL	0-0-3	00051
SWIM	1-0/1-0-9	00015	BOOST	0-0-2	00015
THRUST	1-0/1-0-11	00022	PAUSE	0-0-2	00021
TREAD	1-3/1-3-12	00005	PLUNGE	0-0-4	00005

Table E.3: Low friend / low enemy real stimuli

Irregular	Density	Frequency	Regular	Density	Frequency
BEAT	1-2/2-2-27	00068	FROST	1-1-5	00006
BUY	1-2/1-5-122	00070	WING	11-13-4	00018
CHOOSE	1-1/1-1-27	00050	SOOT	1-1-4	00001
DO	1-0/5-0-60	01363	MIND	4-8-4	00325
EAT	1-2/1-3-27	00061	THIRST	1-1-2	00004
FALL	1-0/2-0-29	00147	BLIND	4-8-4	00047
FIGHT	1-4/2-9-30	00098	WORST	1-1-2	00034
FLEE	1-2/1-5-19	00001	PING	11-13-4	00000
FLY	1-2/2-4-122	00033	GILD	1-2-1	00001
GO	1-4/3-8-56	00626	LIVE	1-3-5	00177
HEAR	1-0/4-0-40	00153	FOOT	1-1-4	00070
HURT	1-0/1-0-21	00037	PUN	1-4-5	00001
LOSE	1-1/1-1-27	00058	SKID	3-7-2	00002
MEET	1-2/1-3-27	00148	KID	3-7-2	00061
RISE	1-0/2-0-277	00049	SIEVE	1-3-5	00001
SHINE	1-0/2-0-42	00005	STUN	1-4-5	00000
SAY	1-1/3-1-67	00504	SUN	1-4-5	00112
SHOOT	1-0/3-0-41	00027	DUN	1-4-5	00001
SLAY	1-1/1-3-67	00000	TING	11-13-4	00000
STICK	1-0/1-0-23	00039	SHUN	1-4-5	00001

Table E.4: Low friend / high enemy real stimuli

Appendix F

Novel Verb ERP Stimuli

Cluster	Density	Pseudo - Irregular	Neighbourhood	Density	Pseudo - Regular
/aɪv/ ↦ /əʊv/	5-0/5-0-10	SCRIVE	/eɪt/	0-0-570	JATE
		CRIVE			NATE
		PLIVE			YATE
/ɜːn/ ↦ /ɜːnt/	2-0/3-0-13	BLURN	/ɛkt/	0-0-45	TECT
		SCURN			PLECT
		PURN			STRECT
/ɛd/ ↦ /ɛd/	3-1/3-1-12	PRED	/ɪʃ/	0-0-49	LISH
		SNED			SISH
		CLED			DRISH
/ʌt/ ↦ /ʌt/	2-0/3-0-10	SNUT	/ɪʤ/	0-0-29	FIDGE
		THRUT			CLIDGE
		PLUT			SLIDGE
		STUT			SPLIDGE
/iːʧ/ ↦ /ɔːt/	2-0/2-0-11	SLEECH	/ɪp/	0-0-32	MIP
		CLEECH			BRIP
		STEECH			GLIP
/iːp/ ↦ /ɛpt/	6-0/8-0-7	CLEEP	/æk/	0-0-33	BRACK
		SKEEP			DRACK
		MEEP			THACK
		FEEP			SPLACK

Table F.1: High friend / low enemy novel stimuli

Cluster	Density	Pseudo - Irregular	Cluster	Density	Pseudo - Regular
/ɛə/ ↦ /ɔː/	4-0/6-0-27	PLAIR	/ɔː/ ↦ /uː/	1-3-17	GLAW
		SLAIR			PLAW
		THRAIR			PRAW
		YAIR			TRAW
/ɛnd/ ↦ /ɛnt/	5-0/7-0-36	STREND	/ɪə/ ↦ /ɜːd/	1-4-40	KEER
		SHEND			PLEER
		JEND			ZEER
		PREND			BLEER
/iːd/ ↦ /ɛd/	6-0/14-0-21	VEED	/ɪl/ ↦ /ɪlt/	1-1-18	PRIL
		CLEED			SHIL
		PREED			SCRILL
/iːn/ ↦ /ɛnt/	2-0/2-0-20	DREEN	/ʌm/ ↦ /eɪm/	1-3-13	STUM
		FLEEN			SHRUM
		SHREEN			VUM
/iːl/ ↦ /ɛlt/	3-1/4-1-18	PREEL	/iːm/ ↦ /ɛmt/	1-1-16	BLEEM
		SCREEL			SHEEM
		SHEEL			CLEEM
/ɪt/ ↦ /ɪt/	4-2/5-2-89	THRIT	/ænd/ ↦ /ʊd/	1-4-10	SHAND
		BLIT			CLAND
		RIT			TRAND

Table F.2: High friend / high enemy novel stimuli

Cluster	Density	Pseudo - Irregular	Neighbourhood	Density	Pseudo - Regular
/aɪk/ ↦ /ʌk/	1-0/1-0-8	DRIKE	/ɔɪst/	0-0-1	LOIST
		SHRIKE			FROIST
		VIKE			THOIST
		SHIKE			TROIST
/ɪg/ ↦ /ʌg/	1-0/1-0-6	SHIG	/ɪlʧ/	0-0-1	HILCH
		SPLIG			PILCH
		SIG			RILCH
		TRIG			SILCH
/ɪm/ ↦ /æm/	1-0/1-0-9	THRIM	/ɪmps/	0-0-1	FIMPSE
		DRIM			MIMPSE
		PIM			DIMPSE
/ʌst/ ↦ /ʌst/	1-0/1-0-11	WUST	/ɒm/	0-0-1	BLOM
		PLUST			GROM
		FUST			STROM
/iːz/ ↦ /əʊz/	1-0/1-0-10	CREEZE	/ʌlʧ/	0-0-1	PULCH
		DREEZE			KULCH
		VEEZE			SULCH
/æʧ/ ↦ /ɔːt/	1-0/1-0-12	DATCH	/ædʒ/	0-0-1	RADGE
		STATCH			LADGE
		GATCH			PADGE

Table F.3: Low friend / low enemy novel stimuli

Cluster	Density	Pseudo - Irregular	Cluster	Density	Pseudo - Regular
/ɔːl/ ↦ /ɛl/	1-0/2-0-29	GRAWL	/aɪnd/ ↦ /aʊnd/	4-8-4	PLYND
		FRAWL			CHYND
		JAWL			FLYND
/aɪn/ ↦ /ɒn/	1-0/2-0-42	STINE			GLYND
		CLINE	/ɜːst/ ↦ /ɜːst/	1-1-2	HURST
		PRINE			BLURST
		TRINE			GURST
/aɪz/ ↦ /əʊz/	1-0/2-0-277	PLYSE	/ɪld/ ↦ /ɪlt/	1-2-1	SPRILD
		CRYSE			RILD
		TYSE			PRILD
		BLYSE	/ɪv/ ↦ /eɪv/	1-3-5	MIV
/ɜːt/ ↦ /ɜːt/	1-0/1-0-21	LURT			SNIV
		SURT			SPRIV
		CLURT			FLIV
/ɪk/ ↦ /ʌk/	1-0/1-0-23	RICK	/ʌn/ ↦ /æn/	1-4-5	CHUN
		THRICK			SCRUN
		DRICK			THRUN
/uːt/ ↦ /ɒt/	1-0/3-0-41	GLOOT	/æv/ ↦ /æd/	1-1-0	RAV
		PROOT			TAV
		STROOT			SPLAV

Table F.4: Low friend / high enemy novel stimuli

Appendix G

Familiarity Data for ERP Stimuli

Stimulus	Fam	Stimulus	Fam	Stimulus	Fam	Stimulus	Fam
ASK	000	FLING	000	NEED	589	SOOT	390
BARK	518	FLIP	000	PAUSE	517	SPEAK	600
BEAT	536	FLOW	518	PICK	524	SPIN	000
BIND	000	FLY	537	PINE	557	SPIT	000
BITE	593	FOOT	583	PING	000	SPLIT	514
BLAZE	000	FREEZE	563	PLAY	586	SPOIL	000
BLEED	000	FROST	546	PLUCK	000	SPREAD	000
BLIND	531	GILD	000	PLUNGE	441	STAND	564
BLINK	450	GIVE	595	PULL	565	STAY	573
BOOK	643	GO	618	PUN	000	STICK	528
BOOST	000	GRIND	486	PUT	539	STRIKE	000
BUILD	554	GROW	562	QUOTE	000	STRIVE	000
BURST	000	HANG	514	RAKE	476	STUN	000
BUY	575	HAUL	454	REACH	577	SUN	635
CARVE	000	HEAR	602	REV	000	SWEAR	000
CATCH	000	HELP	594	RID	000	SWIM	550
CAUSE	557	HIT	000	RISE	546	SWING	000
CHOOSE	534	HOLD	596	ROB	508	TAKE	555
CLASH	488	HURT	579	ROOT	554	TEAR	555
CLING	000	JOIN	544	RUN	000	THIRST	000
COME	608	JUMP	551	RUST	484	THROW	548
COST	589	KEEP	584	SAY	000	THRUST	000
CUT	581	KID	559	SEEK	506	THUMB	601
DIE	569	KILL	549	SEND	529	TING	190
DIG	000	KNIFE	573	SERVE	000	TOUCH	590
DO	609	KNOW	605	SHAKE	000	TRAIN	548
DRAWL	000	LAST	587	SHAVE	000	TREAD	000
DRINK	628	LAUNCH	000	SHED	524	USE	000
DRIVE	000	LEAVE	000	SHINE	000	VOICE	596
DUN	000	LEND	566	SHOOT	536	WANT	606
EARTH	580	LIGHT	575	SHRINK	000	WARN	491
EASE	519	LIVE	608	SHUN	000	WAVE	518
EAT	529	LOOK	607	SHUT	000	WEAVE	000
EYE	611	LOSE	534	SIEVE	000	WEEP	455
FALL	572	LOVE	619	SIGH	518	WIN	582
FEED	000	MARK	533	SIT	581	WING	510
FIELD	000	MEET	575	SKID	000	WOLF	537
FIGHT	561	MIND	591	SLAY	481	WORST	000
FIND	580	MISS	586	SLEEP	610	WOUND	474
FLEE	486	MOURN	000	SMOOTH	570	WRITE	560

Table G.1: Familiarity values for ERP real verb stimuli

Appendix H

Nvalue Data for ERP Stimuli

Stimulus	NVal	Stimulus	NVal	Stimulus	NVal	Stimulus	NVal
ASK	08	FLING	06	NEED	10	SOOT	19
BARK	23	FLIP	06	PAUSE	05	SPEAK	06
BEAT	23	FLOW	12	PICK	17	SPIN	07
BIND	17	FLY	10	PINE	25	SPIT	10
BITE	19	FOOT	14	PING	18	SPLIT	01
BLAZE	08	FREEZE	02	PLAY	10	SPOIL	01
BLEED	03	FROST	03	PLUCK	03	SPREAD	00
BLIND	04	GILD	12	PLUNGE	01	STAND	03
BLINK	07	GIVE	15	PULL	17	STAY	09
BOOK	17	GO	14	PUN	21	STICK	09
BOOST	04	GRIND	03	PUT	21	STRIKE	08
BUILD	03	GROW	11	QUOTE	03	STRIVE	08
BURST	06	HANG	17	RAKE	20	STUN	06
BUY	14	HAUL	10	REACH	09	SUN	19
CARVE	07	HEAR	22	REV	09	SWEAR	05
CATCH	08	HELP	11	RID	18	SWIM	06
CAUSE	07	HIT	24	RISE	18	SWING	08
CHOOSE	03	HOLD	21	ROB	21	TAKE	16
CLASH	07	HURT	09	ROOT	17	TEAR	23
CLING	08	JOIN	07	RUN	19	THIRST	00
COME	22	JUMP	12	RUST	17	THROW	04
COST	19	KEEP	12	SAY	31	THRUST	01
CUT	22	KID	16	SEEK	12	THUMB	02
DIE	19	KILL	21	SEND	15	TING	17
DIG	22	KNIFE	01	SERVE	06	TOUCH	06
DO	17	KNOW	08	SHAKE	12	TRAIN	07
DRAWL	06	LAST	21	SHAVE	13	TREAD	07
DRINK	04	LAUNCH	03	SHED	08	USE	03
DRIVE	02	LEAVE	07	SHINE	13	VOICE	01
DUN	22	LEND	22	SHOOT	07	WANT	17
EARTH	04	LIGHT	13	SHRINK	04	WARN	15
EASE	11	LIVE	19	SHUN	04	WAVE	19
EAT	21	LOOK	15	SHUT	09	WEAVE	05
EYE	17	LOSE	16	SIEVE	03	WEEP	13
FALL	15	LOVE	20	SIGH	08	WIN	17
FEED	14	MARK	23	SIT	23	WING	14
FIELD	06	MEET	13	SKID	06	WOLF	04
FIGHT	12	MIND	18	SLAY	17	WORST	02
FIND	14	MISS	16	SLEEP	05	WOUND	08
FLEE	11	MOURN	03	SMOOTH	01	WRITE	04

Table H.1: Nvalues for ERP real verb stimuli

Appendix I

ERP Presentation Program

```
program pdpcode;

uses Crt, vga, LPT, timeunit;

var
        exp_file, res_file : string;
        expfile, resfile : text;
        lptport, keyval, code, xcoord : integer;
        rnum, stimdur : integer;
        key : char;
        n_trial1, n_trial2, n_trial3 : integer;
        r_port1, r_port2 : integer;
        expitem : string;
        stime, rtime, rt : integer;


procedure information;
begin
clrscr;

                write( 'Test File ?: ');
                readln( exp_file);
                assign( expfile, exp_file);
                reset ( expfile);

                lptport := 1;
                code := LPTbase(lptport);

end;
```

```
procedure close_files;
begin
close(expfile);
        end;

procedure display;
begin
clrscr;
      r_port1:=0;

      repeat r_port1:=port[$300] until r_port1>0;

      for n_trial1 := 1 to 3 do

          begin;

          r_port1:=0;
          Randomize;
          rnum := (Random(5));
          stimdur := 1200 + (rnum*100);
          readln( expfile, expitem);
  xcoord:=LENGTH(expitem);
  gotoXY(38,12);
  writeln('+');
  gotoXY(80,25);
          vgaunblank;
          delay(stimdur);
vgablank;
  clrscr;
  gotoXY(38 - (xcoord DIV 2),12);
  writeln(expitem);
          gotoXY(80,25);
          vgaunblank;
          port[code]:= 1;
          delay(2000);
          vgablank;
          clrscr;
  repeat r_port1:=port[$300] until r_port1>0;
          port[code]:= r_port1;
          delay(1000);
          vgablank;
          clrscr;
         end;
```

```
end;

begin
        load_clock;
        Highvideo;
information;
        display;
        close_files;
        vgaunblank;
end.
```

Appendix J

Nvalue Data for Novel ERP Stimuli

Stimulus	NVal	Stimulus	NVal	Stimulus	NVal	Stimulus	NVal
BLEEM	02	FUST	13	PREEL	02	SPLAV	02
BLEER	04	GATCH	08	PREND	01	SPLIDGE	01
BLIT	09	GLAW	06	PRIL	07	SPLIG	02
BLOM	08	GLIP	08	PRILD	01	SPRILD	00
BLURN	03	GLOOT	03	PRINE	13	SPRIV	02
BLURST	00	GLYND	01	PROOT	02	STATCH	06
BLYSE	00	GRAWL	06	PULCH	05	STEECH	02
BRACK	09	GROM	09	PURN	07	STINE	11
BRIP	08	GURST	06	RADGE	05	STRECT	02
CHUN	08	HILCH	05	RAV	25	STREND	01
CHYND	00	HURST	06	RICK	20	STROM	04
CLAND	06	JATE	17	RILCH	04	STROOT	00
CLED	09	JAWL	09	RILD	07	STUM	07
CLEECH	01	JEND	12	RIT	33	STUT	11
CLEED	03	KEER	16	SCREEL	03	SULCH	04
CLEEM	01	KULCH	04	SCRILL	02	SURT	13
CLEEP	05	LADGE	07	SCRIVE	03	TAV	26
CLIDGE	00	LISH	12	SCRUN	03	TECT	10
CLINE	10	LOIST	04	SCURN	03	THACK	05
CLURT	02	LURT	06	SHAND	07	THOIST	00
CREEZE	03	MEEP	11	SHEEL	09	THRAIR	00
CRIVE	05	MIMPSE	00	SHEEM	04	THRICK	01
CRYSE	02	MIP	23	SHEND	04	THRIM	02
DATCH	10	MIV	16	SHIG	10	THRIT	00
DIMPSE	01	NATE	19	SHIKE	05	THRUN	01
DRACK	05	PADGE	04	SHIL	10	THRUT	01
DREEN	03	PILCH	06	SHREEN	01	TRAND	04
DREEZE	02	PIM	19	SHRIKE	04	TRAW	09
DRICK	05	PLAIR	05	SHRUM	07	TRIG	12
DRIKE	04	PLAW	08	SIG	22	TRINE	11
DRIM	08	PLECT	02	SILCH	04	TROIST	00
DRISH	01	PLEER	02	SISH	12	TYSE	05
FEEP	11	PLIVE	03	SKEEP	06	VEED	11
FIDGE	05	PLUST	02	SLAIR	04	VEEZE	01
FIMPSE	00	PLUT	11	SLEECH	01	VIKE	16
FLEEN	02	PLYND	00	SLIDGE	02	VUM	15
FLIV	02	PLYSE	00	SNED	04	WUST	12
FLYND	00	PRAW	09	SNIV	05	YAIR	06
FRAWL	05	PRED	12	SNUT	10	YATE	13
FROIST	01	PREED	04	SPLACK	00	ZEER	08

Table J.1: Nvalues for ERP novel verb stimuli

Appendix K

Irregular Verbs in Network Training Set

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
BE	00228	11	GROW	00063	9	SIT	00067	9
BEAR	00057	9	HANG	00026	7	SLAY	00000	1
BEAT	00068	9	HAVE	03941	17	SLEEP	00065	9
BEND	00024	7	HEAR	00153	11	SLIDE	00020	6
BET	00020	6	HEAVE	00002	2	SLING	00001	1
BID	00022	7	HIDE	00022	7	SLINK	00000	1
BIND	00004	3	HIT	00115	10	SLIT	00006	4
BITE	00010	5	HOLD	00169	11	SMELL	00034	8
(BLEED)	00002	2	HURT	00037	8	SMITE	00000	1
BLOW	00033	7	KEEP	00264	12	SPEAK	00110	10
BREAK	00088	9	(KNEEL)	00005	4	SPELL	00019	6
BREED	00017	6	KNIT	00010	5	SPEND	00053	8
BRING	00158	11	KNOW	00683	14	SPILL	00001	1
BUILD	00086	9	LEAD	00129	10	SPIN	00005	4
BURN	00015	6	LEAN	00020	6	SPIT	00011	5
BURST	00033	7	LEAP	00014	6	SPLIT	00030	7
BUST	00007	4	LEARN	00084	9	SPOIL	00003	3
BUY	00070	9	LEAVE	00205	11	SPREAD	00083	9
CAST	00045	8	LEND	00014	6	SPRING	00127	10
CATCH	00043	8	LET	00384	12	STAND	00148	10
CHIDE	00002	2	LIGHT	00333	12	STEAL	00005	4
CHOOSE	00050	8	LOSE	00058	9	STICK	00039	8
CLEAVE	00000	1	MAKE	00794	14	STING	00005	4
CLING	00006	4	MAY	01400	15	STINK	00003	3
COME	00630	13	MEAN	00199	11	STRIDE	00016	6
COST	00229	11	MEET	00148	10	STRIKE	00006	4
CREEP	00010	5	OUGHT	00068	9	STRING	00019	6
CUT	00192	11	PUT	00437	13	STRIVE	00007	4
DEAL	00142	10	QUIT	00015	6	(SWEAR)	00010	5
DIG	00010	5	READ	00173	11	(SWEEP)	00015	6
DIVE	00023	7	REND	00001	1	SWIM	00015	6
DO	01363	15	RIDE	00049	8	SWING	00024	7
DRAW	00056	9	RING	00047	8	TAKE	00611	13
DREAM	00064	9	RISE	00102	10	TEACH	00041	8
DRINK	00082	9	RUN	00212	11	TEAR	00011	5
DRIVE	00105	10	SAY	00504	13	TELL	00268	12
DWELL	00008	5	SEE	00772	14	THINK	00433	13
EAT	00061	9	SEEK	00069	9	THRIVE	00001	1
FALL	00147	10	SELL	00041	8	THROW	00042	8
FEED	00123	10	SEND	00074	9	THRUST	00022	7
FEEL	00216	11	SET	00414	13	TREAD	00005	4
FIGHT	00098	10	SHAKE	00017	6	WAKE	00023	7
FIND	00399	12	SHALL	00267	12	WEAR	00036	8
FLEE	00001	1	SHED	00011	5	WEAVE	00004	3
FLING	00002	2	SHINE	00005	4	WED	00002	2
FLY	00033	7	SHIT	00002	2	WEEP	00014	6
FREEZE	00006	4	SHOOT	00027	7	WET	00053	8
GET	00750	14	SHRINK	00005	4	WIN	00055	9
GIRD	00001	1	(SHRIVE)	00000	1	WIND	00063	9
GIVE	00391	12	SHUT	00046	8	WORK	00760	14
GO	00626	13	SING	00034	8	WRITE	00106	10
(GRIND)	00002	2	SINK	00023	7			

Table K.1: Irregular verbs used in network training

Appendix L

Regular Verbs in Network Training Set

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
ACHE	00004	3	BIKE	00000	1	BOOZE	00004	3
ACT	00283	12	BILK	00000	1	BOP	00000	1
ADD	00088	9	BILL	00143	10	BORE	00024	7
AGE	00227	11	BIRCH	00002	2	BOSS	00020	6
AID	00130	10	BITCH	00006	4	BOTCH	00000	1
AIL	00000	1	(BLAB)	00000	1	BOUNCE	00008	5
AIM	00037	8	BLACK	00203	11	BOUND	00042	8
AIR	00257	12	BLAME	00034	8	BOW	00015	6
APE	00003	3	BLANCH	00000	1	BOW	00015	6
ARC	00000	1	BLARE	00000	1	BOWL	00023	7
ARCH	00013	6	BLAST	00015	6	BOX	00070	9
ARM	00094	10	BLAZE	00007	4	BRACE	00014	6
ASK	00128	10	BLEACH	00000	1	BRAG	00002	2
AWE	00005	4	BLEAT	00001	1	BRAID	00000	1
AXE	00006	4	BLEEP	00000	1	BRAIN	00045	8
BACK	00967	14	BLENCH	00000	1	BRAISE	00000	1
BAG	00042	8	BLEND	00009	5	BRANCH	00033	7
BAIL	00007	4	BLESS	00009	5	BRAND	00017	6
BAIT	00002	2	BLIGHT	00002	2	BRAVE	00024	7
BAKE	00012	5	BLIND	00047	8	BRAWL	00001	1
BALK	00000	1	BLINK	00004	3	BRAY	00000	1
BAN	00007	4	BLITZ	00003	3	BREACH	00006	4
BAND	00053	8	BLOAT	00008	5	BREAST	00011	5
BANG	00007	4	BLOCK	00066	9	BREATHE	00007	4
BANK	00083	9	BLOOD	00121	10	BREEZE	00014	6
BAR	00082	9	BLOOM	00012	5	BREW	00004	3
BARGE	00007	4	BLOT	00006	4	BRIBE	00001	1
BARK	00014	6	BLUE	00143	10	BRIDGE	00098	10
BASH	00000	1	BLUFF	00008	5	BRIEF	00073	9
BASK	00000	1	BLUNT	00009	5	BRIM	00004	3
BASTE	00000	1	BLUR	00003	3	BROACH	00001	1
BAT	00018	6	BLUSH	00002	2	BROIL	00002	2
BATH	00026	7	BOARD	00239	11	BRONZE	00011	5
BATHE	00004	3	BOAST	00008	5	BROOD	00009	5
BAWL	00000	1	BOAT	00072	9	BROOK	00003	3
BAY	00057	9	BOB	00040	8	BROWN	00176	11
BEACH	00061	9	BODE	00000	1	BROWSE	00000	1
BEAD	00001	1	BOIL	00012	5	BRUISE	00003	3
BEAM	00021	7	BOLT	00010	5	BRUSH	00044	8
BEARD	00026	7	BOMB	00036	8	BUCK	00020	6
BED	00127	10	BOND	00046	8	BUD	00009	5
BEEF	00032	7	BONE	00033	7	BUDGE	00003	3
BEG	00011	5	BOO	00001	1	BUFF	00005	4
BELCH	00002	2	BOOB	00000	1	BUG	00004	3
BELT	00029	7	BOOK	00193	11	BULGE	00005	4
BERTH	00004	3	BOOM	00008	5	BULL	00014	6
BEST	00351	12	BOOST	00015	6	BUM	00007	4
BIB	00002	2	BOOT	00013	6	BUMP	00005	4

Table L.1: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
BUNCH	00017	6	CHECK	00088	9	CLOT	00003	3
BUNG	00000	1	CHEEK	00020	6	CLOTHE	00001	1
BUNK	00018	6	CHEEP	00000	1	CLOUD	00028	7
BUOY	00000	1	CHEER	00008	5	CLOUT	00001	1
BURKE	00007	4	CHEW	00002	2	CLOWN	00003	3
BURP	00000	1	CHILL	00014	6	CLOY	00000	1
BURR	00007	4	CHIME	00000	1	CLUB	00145	10
BUS	00034	8	CHINK	00000	1	CLUCK	00003	3
BUSK	00000	1	CHIP	00017	6	CLUMP	00004	3
BUTT	00012	5	CHIRP	00000	1	CLUNK	00000	1
BUZZ	00013	6	CHOCK	00000	1	CLUTCH	00005	4
CADGE	00000	1	CHOKE	00009	5	COACH	00024	7
CAGE	00009	5	CHOMP	00001	1	COAL	00032	7
CAKE	00013	6	CHOP	00003	3	COAST	00061	9
CALL	00188	11	CHUCK	00014	6	COAT	00043	8
CALM	00035	8	CHUG	00000	1	COAX	00001	1
CAMP	00075	9	CHUM	00001	1	COCK	00005	4
CAN	01772	15	CHURCH	00348	12	COD	00006	4
CANE	00012	5	CHURN	00000	1	CODE	00040	8
CANT	00001	1	CINCH	00003	3	COIL	00006	4
CAP	00027	7	CITE	00007	4	COIN	00010	5
CARD	00026	7	CLACK	00000	1	COKE	00004	3
CARE	00162	11	CLAIM	00098	10	COMB	00006	4
CARP	00000	1	CLAM	00003	3	CON	00007	4
CART	00005	4	(CLAMP)	00000	1	CONK	00000	1
CARVE	00003	3	CLANG	00001	1	COO	00000	1
CASE	00362	12	CLANK	00000	1	COOK	00047	8
CASH	00036	8	CLAP	00001	1	COOL	00062	9
CAUSE	00130	10	CLASH	00005	4	COP	00015	6
CEASE	00015	6	CLASP	00000	1	COPE	00021	7
CHAFE	00001	1	CLASS	00207	11	CORD	00006	4
CHAFF	00000	1	CLAW	00001	1	CORE	00037	8
CHAIN	00050	8	CLEAN	00070	9	CORK	00009	5
CHAIR	00066	9	CLEANSE	00000	1	CORN	00034	8
CHALK	00003	3	CLEAR	00219	11	COSH	00000	1
CHAMP	00001	1	CLENCH	00001	1	COUCH	00012	5
CHANCE	00131	10	CLERK	00034	8	COUGH	00007	4
CHANGE	00240	11	CLEW	00000	1	COUNT	00049	8
CHANT	00002	2	CLICK	00002	2	COURSE	00465	13
CHAP	00005	4	CLIMB	00012	5	COURT	00230	11
CHAR	00001	1	CLINCH	00002	2	COW	00029	7
CHARGE	00122	10	CLINK	00000	1	COWL	00000	1
CHARM	00026	7	CLIP	00006	4	COX	00005	4
CHART	00022	7	CLOAK	00003	3	CRAB	00000	1
CHASE	00018	6	CLOCK	00020	6	CRACK	00021	7
CHAT	00005	4	CLOG	00002	2	CRAFT	00023	7
CHAW	00001	1	CLOP	00000	1	CRAM	00000	1
CHEAT	00003	3	CLOSE	00234	11	CRAMP	00002	2

Table L.2: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
CRANE	00005	4	DARN	00003	3	DRIP	00001	1
CRANK	00001	1	DART	00000	1	DRONE	00003	3
CRAP	00003	3	DASH	00011	5	DROOL	00000	1
CRASH	00020	6	DATE	00103	10	DROOP	00001	1
CRATE	00002	2	DAUB	00000	1	DROP	00059	9
CRAVE	00002	2	DAUNT	00001	1	DROWN	00003	3
CRAWL	00011	5	DAWN	00028	7	DROWSE	00000	1
CRAZE	00002	2	DAZE	00000	1	DRUB	00000	1
CREAK	00001	1	DECK	00023	7	DRUDGE	00000	1
CREAM	00020	6	DEEM	00001	1	DRUG	00024	7
CREASE	00001	1	DEIGN	00000	1	DRUM	00011	5
CREST	00012	5	DELVE	00000	1	DRY	00068	9
CREW	00036	8	DENT	00002	2	DUB	00000	1
CRIB	00005	4	DIB	00000	1	DUCK	00009	5
CRICK	00000	1	DICE	00014	6	DULL	00027	7
CRIME	00034	8	DIE	00073	9	DUMP	00004	3
CRIMP	00000	1	DIKE	00000	1	DUN	00001	1
CRINGE	00000	1	DIM	00019	6	DUNK	00000	1
CRISP	00008	5	DIN	00001	1	DUPE	00000	1
CROAK	00001	1	DINE	00002	2	DUST	00070	9
CROCK	00000	1	DIP	00006	4	DWARF	00003	3
CROOK	00003	3	DITCH	00010	5	EARN	00016	6
CROON	00000	1	DOCK	00008	5	EARTH	00150	11
CROP	00020	6	DODGE	00011	5	EASE	00042	8
CROSS	00055	9	(DOFF)	00000	1	EBB	00001	1
CROUCH	00007	4	DOG	00075	9	EDGE	00078	9
CROW	00002	2	DON	00023	7	EFF	00000	1
CROWD	00053	8	DOOM	00003	3	END	00410	13
CROWN	00019	6	DOPE	00002	2	ERR	00001	1
CRUISE	00002	2	DOSE	00011	5	ETCH	00000	1
CRUNCH	00002	2	DOT	00013	6	EYE	00122	10
CRUSH	00004	3	DOUBT	00114	10	FACE	00371	12
CRY	00048	8	DOUSE	00000	1	FADE	00002	2
CUBE	00001	1	DOWN	00895	14	FAG	00000	1
CUFF	00001	1	DOWSE	00000	1	FAIL	00037	8
CULL	00000	1	DOZE	00000	1	FAINT	00025	7
CUP	00045	8	DRAFT	00024	7	FAKE	00010	5
CURB	00013	6	DRAG	00015	6	FAN	00018	6
CURE	00028	7	DRAIN	00018	6	FARE	00007	4
CURL	00002	2	DRAPE	00000	1	FARM	00125	10
CURSE	00011	5	(DRAT)	00000	1	FART	00000	1
CURVE	00045	8	DRAWL	00002	2	FAST	00078	9
CUSS	00000	1	DREAD	00009	5	FAT	00060	9
DAB	00000	1	DREDGE	00000	1	FAULT	00022	7
DAMN	00034	8	DRENCH	00000	1	FAWN	00001	1
DAMP	00016	6	DRESS	00067	9	FAZE	00001	1
DANCE	00090	9	DRIFT	00018	6	FEAR	00127	10
DARE	00021	7	DRILL	00033	7	FEAST	00003	3

Table L.3: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
FEE	00016	6	FLOW	00067	9	GAG	00004	3
FEIGN	00000	1	FLUFF	00001	1	GAIN	00074	9
FELL	00092	10	FLUNK	00000	1	GALL	00007	4
FENCE	00030	7	FLUSH	00011	5	GAME	00123	10
FETCH	00006	4	FLUTE	00001	1	GAPE	00000	1
FETE	00003	3	FOAL	00002	2	GARB	00003	3
FEUD	00000	1	FOAM	00037	8	GAS	00098	10
FIB	00000	1	FOG	00025	7	GASH	00001	1
FIELD	00274	12	FOIL	00020	6	GASP	00003	3
FILCH	00000	1	FOLD	00007	4	GATE	00037	8
FILE	00081	9	FOOL	00037	8	GAUGE	00012	5
FILL	00050	8	FOOT	00070	9	GAWK	00000	1
FILM	00096	10	FORCE	00230	11	GAWP	00000	1
FINE	00161	11	FORD	00024	7	GAZE	00012	5
FIRM	00109	10	FORK	00014	6	GEAR	00026	7
FISH	00035	8	FORM	00370	12	GEL	00002	2
FIT	00075	9	FOUL	00004	3	GELD	00000	1
FIX	00014	6	FOUND	00536	13	GHOST	00011	5
FIZZ	00000	1	FOX	00013	6	GIBE	00001	1
FLAG	00016	6	FRAME	00074	9	GIFT	00033	7
FLAIL	00001	1	FRANK	00068	9	GILD	00001	1
FLAKE	00001	1	FRAY	00001	1	GIN	00023	7
FLAME	00017	6	FREAK	00004	3	GLANCE	00040	8
FLANK	00002	2	FREE	00260	12	GLARE	00007	4
FLAP	00000	1	FREIGHT	00028	7	GLAZE	00011	5
FLARE	00003	3	FRET	00001	1	GLEAM	00004	3
FLASH	00021	7	FRIGHT	00002	2	GLEAN	00001	1
FLAUNT	00000	1	FRINGE	00016	6	GLIDE	00002	2
FLAY	00000	1	FRISK	00000	1	GLINT	00002	2
FLECK	00001	1	FRIZZ	00000	1	GLOAT	00000	1
FLEDGE	00000	1	FRONT	00221	11	GLOSS	00001	1
FLEECE	00000	1	FROST	00006	4	GLOW	00016	6
FLEET	00017	6	FROTH	00001	1	GLUE	00008	5
FLEX	00002	2	FROWN	00001	1	GLUT	00000	1
FLICK	00002	2	FROWST	00000	1	GNASH	00000	1
FLIGHT	00046	8	FRUIT	00035	8	GNAW	00001	1
FLINCH	00000	1	FRY	00002	2	GOAD	00001	1
FLIP	00004	3	FUCK	00000	1	GOLF	00034	8
FLIRT	00001	1	FUDGE	00000	1	GONG	00000	1
FLIT	00000	1	FULL	00230	11	GOOF	00000	1
FLOAT	00003	3	FUME	00000	1	GOOSE	00004	3
FLOCK	00010	5	FUND	00062	9	GORE	00007	4
FLOG	00001	1	FUNK	00004	3	GORGE	00001	1
FLOOD	00019	6	FUR	00013	6	GOUGE	00001	1
FLOOR	00158	11	(FURL)	00000	1	GOWN	00016	6
FLOP	00001	1	FUSE	00005	4	GRAB	00016	6
FLOUNCE	00000	1	FUSS	00004	3	GRACE	00040	8
FLOUT	00000	1	GAFF	00000	1	GRADE	00035	8

Table L.4: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
GRAFT	00001	1	HARK	00003	3	HUFF	00010	5
GRANT	00047	8	HARM	00025	7	HUG	00003	3
GRAPH	00017	6	HASH	00001	1	HULL	00013	6
GRASP	00017	6	HATCH	00005	4	HUM	00005	4
GRASS	00053	8	HATE	00042	8	HUMP	00002	2
GRATE	00003	3	HAUL	00005	4	HUNCH	00007	4
GRAZE	00001	1	HAUNT	00004	3	HUNT	00010	5
GREASE	00009	5	HAW	00001	1	HURL	00003	3
GREET	00007	4	HAWK	00014	6	HUSH	00004	3
GREY	00012	5	HAZE	00007	4	HUSK	00000	1
GRIEVE	00000	1	HEAD	00424	13	HUT	00013	6
GRILL	00012	5	HEAL	00002	2	HYMN	00009	5
GRIME	00000	1	HEAP	00014	6	ICE	00045	8
GRIN	00013	6	HEAT	00097	10	INCH	00040	8
GRIP	00020	6	HEDGE	00002	2	INK	00007	4
GRIPE	00000	1	HEED	00008	5	IRK	00000	1
GROAN	00001	1	HELP	00311	12	ITCH	00005	4
GROIN	00004	3	HEM	00004	3	JAB	00001	1
GROOM	00005	4	HERD	00022	7	JAG	00001	1
GROOVE	00002	2	HEW	00000	1	JAIL	00021	7
GROPE	00001	1	HEX	00000	1	JAM	00006	4
GROSS	00066	9	HIE	00000	1	JAR	00016	6
GROUCH	00000	1	HIKE	00004	3	JAUNT	00000	1
GROUND	00186	11	HINT	00009	5	JAW	00016	6
GROUP	00390	12	HISS	00002	2	JAZZ	00099	10
GROUSE	00000	1	HITCH	00005	4	JEER	00000	1
GROWL	00004	3	HIVE	00002	2	JERK	00002	2
GRUB	00002	2	HOARD	00000	1	JEST	00001	1
GRUDGE	00007	4	HOAX	00000	1	JET	00029	7
GRUNT	00002	2	HOCK	00000	1	JEWEL	00001	1
GUARD	00048	8	HOE	00000	1	JIB	00000	1
GUESS	00056	9	HOG	00003	3	JIG	00008	5
GUEST	00039	8	HOIST	00001	1	(JILT)	00000	1
GUIDE	00036	8	HOLE	00058	9	JINX	00000	1
GULL	00001	1	HOME	00547	13	JIVE	00000	1
GULP	00002	2	HONE	00002	2	JOB	00238	11
GUM	00014	6	HONK	00000	1	JOG	00000	1
GUSH	00001	1	HOOD	00007	4	JOIN	00065	9
GUT	00001	1	HOOK	00005	4	JOINT	00039	8
GUY	00051	8	HOOP	00003	3	JOKE	00022	7
GYP	00006	4	HOOT	00009	5	JOLT	00004	3
HACK	00003	3	HOP	00002	2	JOSH	00000	1
HAIL	00010	5	HOPE	00178	11	JOT	00001	1
HALT	00010	5	HOSE	00009	5	JOUST	00001	1
HALVE	00000	1	HOST	00036	8	JUDGE	00077	9
HAM	00019	6	HOUND	00007	4	JUG	00006	4
HAND	00431	13	HOUSE	00591	13	JUICE	00011	5
HARE	00001	1	HOWL	00004	3	JUMP	00024	7

Table L.5: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
JUNK	00008	5	LIME	00013	6	MASK	00009	5
KEEN	00011	5	LIMN	00000	1	MASS	00110	10
KEN	00012	5	LIMP	00012	5	MAT	00005	4
KEY	00088	9	LINE	00298	12	MATCH	00041	8
KICK	00016	6	LINK	00016	6	MATE	00021	7
KID	00061	9	LISP	00000	1	MAUL	00000	1
KILL	00063	9	LIST	00133	10	MELD	00001	1
KINK	00000	1	LIVE	00177	11	MELT	00004	3
KIP	00000	1	LOAD	00045	8	MEND	00002	2
KISS	00017	6	LOAF	00004	3	MERGE	00010	5
KNEE	00035	8	LOAN	00046	8	MESH	00004	3
KNIFE	00076	9	LOATHE	00000	1	MESS	00022	7
KNIGHT	00018	6	LOB	00001	1	MEW	00001	1
KNOCK	00015	6	LOCK	00023	7	MIFF	00000	1
KNOT	00008	5	LODGE	00019	6	MILK	00049	8
LACE	00007	4	LOFT	00002	2	MILL	00011	5
LACK	00110	10	LOG	00011	5	MIME	00000	1
LADE	00000	1	LOLL	00000	1	MINCE	00001	1
LAG	00003	3	LONG	00755	14	MIND	00325	12
LAMB	00007	4	LOOK	00399	12	MINE	00059	9
LAME	00002	2	LOOM	00006	4	MINT	00007	4
LANCE	00003	3	LOOP	00021	7	MISS	00258	12
LAND	00217	11	LOOSE	00053	8	MIST	00014	6
LAP	00019	6	LOOT	00003	3	MIX	00013	6
LAPSE	00006	4	LOP	00001	1	MOAN	00001	1
LARD	00004	3	LOPE	00002	2	MOB	00010	5
LASH	00006	4	LOUNGE	00009	5	MOCK	00008	5
LAST	00676	14	LOVE	00232	11	MOIL	00000	1
LATCH	00005	4	LOW	00174	11	MOO	00000	1
LAUD	00000	1	LUFF	00000	1	MOOCH	00000	1
LAUGH	00028	7	LUG	00002	2	MOOR	00000	1
LAUNCH	00010	5	LULL	00002	2	MOOT	00001	1
LAVE	00000	1	LUMP	00007	4	MOP	00003	3
LAY	00139	10	LUNCH	00033	7	MOPE	00000	1
LAZE	00001	1	LUNGE	00004	3	MOULD	00001	1
LEACH	00000	1	LURCH	00003	3	MOULT	00000	1
LEAF	00012	5	LURE	00007	4	MOUNT	00026	7
LEAGUE	00069	9	LURK	00001	1	MOURN	00002	2
LEAK	00002	2	LYNCH	00000	1	MOUSE	00010	5
LEASE	00010	5	MAIL	00047	8	MOUTH	00103	10
LECH	00000	1	MAIM	00000	1	MOVE	00171	11
LEER	00000	1	MALT	00001	1	MOW	00000	1
LEG	00058	9	MAN	01207	15	MUCK	00001	1
LICK	00003	3	MAP	00013	6	MUD	00032	7
LIE	00059	9	MAR	00007	4	MUFF	00001	1
LIFT	00023	7	MARCH	00120	10	MUG	00001	1
LIKE	01290	15	MARK	00083	9	MULCH	00006	4
LILT	00001	1	MASH	00001	1	MULCT	00000	1

Table L.6: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
MULL	00000	1	PANT	00000	1	PLACE	00571	13
MUM	00001	1	PAR	00000	1	PLAGUE	00006	4
MUNCH	00001	1	PARCH	00000	1	PLAIT	00000	1
MUSE	00004	3	PARK	00094	10	PLAN	00205	11
MUSS	00000	1	PARSE	00000	1	PLANE	00114	10
MUTE	00003	3	PART	00500	13	PLANK	00007	4
NAB	00001	1	PASS	00089	9	PLANT	00125	10
NAG	00000	1	PASTE	00010	5	PLASH	00000	1
NAIL	00006	4	PAT	00035	8	PLATE	00022	7
NAME	00294	12	PATCH	00013	6	PLAY	00200	11
NAP	00004	3	PAUSE	00021	7	PLEACH	00000	1
NARK	00000	1	PAVE	00002	2	PLEAD	00005	4
NEAR	00198	11	PAWN	00002	2	PLEASE	00062	9
NECK	00081	9	PAY	00172	11	PLEAT	00000	1
NEED	00360	12	PEACH	00003	3	PLEDGE	00003	3
NEIGH	00000	1	PEARL	00009	5	PLIGHT	00007	4
NERVE	00012	5	PECK	00005	4	PLOD	00001	1
NEST	00020	6	PEE	00003	3	PLONK	00000	1
NET	00034	8	PEEK	00000	1	PLOP	00000	1
NICK	00025	7	PEEL	00003	3	PLOT	00037	8
NIP	00003	3	PEEP	00002	2	PLOUGH	00000	1
NIX	00000	1	PEER	00008	5	PLUCK	00002	2
NOD	00012	5	PEEVE	00000	1	PLUG	00023	7
NOOSE	00003	3	PEG	00004	3	PLUMB	00005	4
NOSE	00060	9	PELT	00000	1	PLUME	00002	2
NOSH	00000	1	PEN	00018	6	PLUMP	00004	3
NOTCH	00006	4	PERCH	00001	1	PLUNGE	00005	4
NOTE	00127	10	PERK	00001	1	PLUNK	00000	1
NUDGE	00002	2	PERM	00000	1	PLY	00000	1
NUMB	00004	3	PET	00008	5	POACH	00001	1
NURSE	00017	6	PHONE	00054	8	POD	00003	3
OIL	00093	10	PHRASE	00034	8	POINT	00395	12
OINK	00000	1	PICK	00055	9	POISE	00006	4
OUST	00003	3	PIECE	00129	10	POKE	00001	1
OUT	02096	16	PIERCE	00006	4	POLL	00009	5
OWE	00010	5	PIKE	00041	8	PONG	00000	1
OWN	00772	14	PILE	00025	7	POOL	00111	10
PACE	00043	8	PIMP	00003	3	POP	00008	5
PACK	00025	7	PIN	00016	6	POSE	00011	5
PAD	00008	5	PINCH	00006	4	POST	00084	9
PAGE	00066	9	PINE	00014	6	POT	00028	7
PAIN	00088	9	PING	00000	1	POUCH	00002	2
PAINT	00037	8	PINK	00048	8	POUNCE	00000	1
PAIR	00050	8	PIP	00023	7	POUND	00028	7
PALE	00058	9	PIPE	00020	6	POUR	00009	5
PALL	00004	3	PISS	00001	1	POUT	00001	1
PALM	00022	7	PIT	00014	6	PRAISE	00017	6
PAN	00016	6	PITCH	00022	7	PRANCE	00000	1

Table L.7: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
PRAWN	00000	1	QUEST	00016	6	ROAM	00006	4
PRAY	00012	5	QUEUE	00000	1	ROAR	00013	6
PREACH	00008	5	QUILT	00000	1	ROAST	00010	5
PREEN	00000	1	QUIP	00000	1	ROB	00019	6
PRESS	00127	10	QUIZ	00002	2	ROBE	00006	4
PRICE	00108	10	QUOTE	00017	6	ROCK	00075	9
PRICK	00002	2	RACE	00103	10	ROLL	00035	8
PRIM	00001	1	RACK	00009	5	ROMP	00001	1
PRIME	00045	8	RAFT	00004	3	ROOF	00059	9
PRIMP	00000	1	RAG	00010	5	ROOK	00000	1
PRINK	00000	1	RAGE	00016	6	ROOM	00383	12
PRINT	00018	6	RAID	00010	5	ROOST	00001	1
PRIZE	00028	7	RAIL	00016	6	ROOT	00030	7
PROBE	00006	4	RAIN	00070	9	ROPE	00015	6
PROD	00002	2	RAISE	00052	8	ROT	00008	5
PRONG	00000	1	RAKE	00011	5	ROUGE	00007	4
PROOF	00040	8	RAM	00002	2	ROUGH	00041	8
PROP	00007	4	RANGE	00160	11	ROUND	00081	9
PROVE	00053	8	RANK	00024	7	ROUSE	00002	2
PROWL	00002	2	RANT	00000	1	ROUT	00000	1
PRUNE	00001	1	RAPE	00005	4	ROVE	00001	1
PRY	00006	4	RASP	00002	2	ROW	00035	8
PUFF	00001	1	RAT	00006	4	ROW	00035	8
PUG	00000	1	RATE	00209	11	RUB	00006	4
PUKE	00001	1	RAVE	00000	1	RUCK	00000	1
PULE	00000	1	RAY	00019	6	RUE	00006	4
PULL	00051	8	REACH	00106	10	RULE	00073	9
PULP	00005	4	REAP	00003	3	RUSH	00020	6
PULSE	00009	5	REAR	00051	8	RUST	00010	5
PUMP	00011	5	REEF	00011	5	RUT	00001	1
PUN	00001	1	REEK	00002	2	SACK	00008	5
PUNCH	00005	4	REEL	00002	2	SAG	00004	3
PUNT	00000	1	RENT	00021	7	SAIL	00012	5
PUP	00002	2	REST	00163	11	SALT	00046	8
PURGE	00002	2	RET	00000	1	SALVE	00003	3
PURR	00000	1	RETCH	00001	1	SAND	00028	7
PURSE	00014	6	REV	00033	7	SAP	00001	1
PUSH	00037	8	RHYME	00003	3	SASS	00000	1
PUTT	00007	4	RIB	00001	1	SATE	00000	1
QUACK	00009	5	RICK	00000	1	SAUCE	00020	6
QUAFF	00000	1	RIDGE	00018	6	SAVE	00062	9
QUAIL	00000	1	RIG	00005	4	SCAB	00000	1
QUAKE	00002	2	RILE	00000	1	SCALD	00001	1
QUASH	00000	1	RIM	00005	4	SCALE	00060	9
QUEEN	00041	8	RINSE	00006	4	SCALP	00004	3
QUEER	00006	4	RIP	00006	4	SCAMP	00000	1
QUELL	00003	3	RISK	00054	8	SCAN	00005	4
QUENCH	00001	1	RIVE	00000	1	SCANT	00005	4

Table L.8: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
SCAR	00010	5	SERVE	00107	10	SIP	00002	2
SCARE	00003	3	SEX	00084	9	SIZE	00138	10
SCAT	00000	1	SHADE	00028	7	SKATE	00001	1
SCENT	00006	4	SHAFT	00011	5	SKETCH	00016	6
SCHEME	00033	7	SHAG	00001	1	SKI	00005	4
SCHOOL	00492	13	SHAM	00001	1	SKID	00002	2
SCOFF	00000	1	SHAME	00021	7	SKIM	00000	1
SCOLD	00000	1	SHAPE	00085	9	SKIMP	00000	1
SCOOP	00005	4	SHARE	00098	10	SKIN	00047	8
SCOOT	00000	1	SHAVE	00006	4	SKIP	00005	4
SCORCH	00000	1	SHEAR	00040	8	SKIRT	00021	7
SCORE	00066	9	SHEATHE	00000	1	SKIVE	00000	1
SCORN	00004	3	SHELL	00022	7	SKULK	00001	1
SCOTCH	00005	4	SHELVE	00000	1	SKUNK	00000	1
SCOURGE	00002	2	SHIELD	00008	5	SKY	00058	9
SCOUT	00008	5	SHIFT	00041	8	SLACK	00009	5
SCOWL	00000	1	SHIN	00003	3	SLAKE	00000	1
SCRAG	00000	1	SHIP	00083	9	SLAM	00003	3
SCRAM	00000	1	SHIRK	00000	1	SLANG	00002	2
SCRAP	00008	5	SHIRR	00000	1	SLANT	00003	3
SCRAPE	00003	3	SHOAL	00000	1	SLAP	00002	2
SCRATCH	00009	5	SHOCK	00031	7	SLASH	00003	3
SCRAWL	00000	1	SHOO	00000	1	SLATE	00010	5
SCREAM	00013	6	SHOP	00063	9	SLAVE	00030	7
SCREECH	00001	1	SHORE	00061	9	SLED	00000	1
SCREEN	00048	8	SHORT	00212	11	SLEDGE	00000	1
SCREW	00021	7	SHOUT	00009	5	SLEEK	00002	2
SCRIBE	00004	3	SHOVE	00002	2	SLEET	00001	1
SCRIMP	00000	1	SHOW	00287	12	SLEW	00000	1
SCROUNGE	00000	1	SHRED	00003	3	SLICE	00013	6
SCRUB	00009	5	SHRIEK	00005	4	SLIGHT	00053	8
SCRUNCH	00000	1	SHRILL	00007	4	SLIM	00020	6
SCUD	00000	1	SHRIMP	00002	2	SLIP	00019	6
SCUFF	00001	1	SHRINE	00007	4	SLOG	00000	1
SCULL	00000	1	SHROUD	00000	1	SLOP	00002	2
SCULPT	00000	1	SHRUG	00002	2	SLOPE	00019	6
SCYTHE	00000	1	SHUCK	00000	1	SLOSH	00000	1
SEAL	00017	6	SHUN	00001	1	SLOT	00006	4
SEAR	00002	2	SHUNT	00001	1	SLOUCH	00001	1
SEARCH	00066	9	SHUSH	00000	1	SLOW	00060	9
SEAT	00054	8	SHY	00013	6	SLUG	00010	5
SEED	00041	8	SIDE	00380	12	SLUICE	00002	2
SEEM	00229	11	SIEVE	00001	1	SLUM	00008	5
SEEP	00002	2	SIFT	00000	1	SLUMP	00008	5
SEETHE	00000	1	SIGH	00011	5	SLUR	00000	1
SEINE	00000	1	SIGN	00094	10	SLURP	00000	1
SEIZE	00006	4	SIN	00053	8	SMACK	00004	3
SENSE	00311	12	SINGE	00000	1	SMART	00021	7

Table L.9: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
SMASH	00004	3	SOUND	00204	11	SQUEAL	00001	1
SMEAR	00002	2	SOUSE	00000	1	SQUEEZE	00011	5
SMELT	00003	3	SOW	00003	3	SQUELCH	00000	1
SMILE	00058	9	SPACE	00184	11	SQUINT	00001	1
SMIRCH	00000	1	SPADE	00010	5	SQUIRM	00000	1
SMIRK	00003	3	SPAN	00019	6	SQUIRT	00001	1
SMOKE	00041	8	SPANK	00000	1	STAB	00003	3
SMOOCH	00000	1	SPAR	00000	1	STACK	00009	5
SMOOTH	00042	8	SPARE	00023	7	STAFF	00113	10
SMUDGE	00000	1	SPARK	00012	5	STAGE	00174	11
SMUT	00000	1	SPAWN	00000	1	STAIN	00006	4
SNACK	00006	4	SPAY	00000	1	STAKE	00020	6
SNAG	00003	3	SPEAR	00007	4	STALE	00004	3
SNAKE	00044	8	SPEED	00083	9	STALK	00000	1
SNAP	00012	5	SPEW	00000	1	STALL	00018	6
SNARE	00001	1	SPICE	00004	3	STAMP	00008	5
SNARL	00000	1	SPIEL	00000	1	STANCH	00001	1
SNATCH	00004	3	SPIKE	00002	2	STAR	00025	7
SNEAK	00002	2	SPITE	00056	9	STARCH	00004	3
SNEER	00001	1	SPLASH	00003	3	STARE	00014	6
SNEEZE	00000	1	SPLAY	00000	1	START	00154	11
SNICK	00001	1	SPLICE	00001	1	STARVE	00001	1
SNIFF	00002	2	SPLODGE	00000	1	STASH	00000	1
SNIP	00000	1	SPLOTCH	00000	1	STATE	00808	14
SNIPE	00000	1	SPLURGE	00001	1	STAUNCH	00003	3
SNITCH	00000	1	SPONGE	00007	4	STAY	00113	10
SNOG	00000	1	SPOOF	00001	1	STEAM	00017	6
SNOOP	00001	1	SPOOK	00000	1	STEEP	00013	6
SNOOZE	00000	1	SPOON	00006	4	STEER	00009	5
SNORE	00000	1	SPORT	00017	6	STEM	00029	7
SNORT	00003	3	SPOT	00057	9	STEP	00131	10
SNOW	00059	9	SPOUT	00001	1	STET	00000	1
SNUB	00000	1	SPRAIN	00000	1	STEW	00005	4
SNUFF	00000	1	SPRAWL	00003	3	STILL	00782	14
SNUG	00002	2	SPRAY	00016	6	STINT	00006	4
SOAK	00007	4	SPRINT	00000	1	STIR	00007	4
SOAP	00022	7	SPROUT	00001	1	STITCH	00003	3
SOAR	00000	1	SPRUCE	00005	4	STOCK	00147	10
SOB	00000	1	SPUR	00013	6	STOKE	00000	1
SOCK	00004	3	SPURN	00000	1	STOMP	00000	1
SOIL	00054	8	SPURT	00002	2	STONE	00058	9
SOLE	00018	6	SPY	00009	5	STOOGE	00000	1
SOLVE	00020	6	SQUALL	00007	4	STOOP	00004	3
SOOT	00001	1	SQUARE	00143	10	STOP	00120	10
SOOTHE	00002	2	SQUASH	00002	2	STORE	00074	9
SOP	00001	1	SQUAT	00007	4	STORM	00026	7
SORT	00164	11	SQUAWK	00001	1	STOW	00000	1
SOUGH	00000	1	SQUEAK	00001	1	STRAFE	00002	2

Table L.10: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
STRAIN	00031	7	SWINGE	00000	1	THWACK	00001	1
STRAND	00007	4	SWIPE	00002	2	THWART	00003	3
STRAP	00002	2	SWIRL	00002	2	TICK	00003	3
STRAW	00015	6	SWISH	00000	1	TIE	00023	7
STRAY	00012	5	SWITCH	00043	8	TILE	00016	6
STREAK	00010	5	SWOON	00000	1	TILL	00050	8
STREAM	00051	8	SWOOP	00002	2	TILT	00005	4
STRESS	00107	10	TACK	00004	3	TIME	01599	15
STRETCH	00026	7	TAG	00005	4	TIN	00012	5
STREW	00000	1	TAIL	00024	7	TING	00000	1
STRIP	00030	7	TAINT	00001	1	TINGE	00000	1
STROKE	00019	6	TALK	00154	11	TINT	00001	1
STROLL	00004	3	TAME	00005	4	TIP	00022	7
STROP	00000	1	TAMP	00001	1	TOAST	00019	6
STRUM	00000	1	TAN	00009	5	TOIL	00001	1
STRUT	00003	3	TAP	00018	6	TOLL	00016	6
STUB	00003	3	TAPE	00035	8	TONE	00078	9
STUD	00007	4	TAR	00012	5	TOOL	00040	8
STUFF	00032	7	TASK	00060	9	TOOT	00003	3
STUMP	00002	2	TASTE	00059	9	TOP	00204	11
STUN	00000	1	TAT	00001	1	TOPE	00000	1
STUNT	00001	1	TAUNT	00004	3	TOSS	00009	5
STYLE	00098	10	TAX	00197	11	TOTE	00001	1
SUB	00005	4	TEASE	00006	4	TOUCH	00087	9
SUCK	00005	4	TEE	00005	4	TOUR	00043	8
SUE	00018	6	TEEM	00000	1	TOUT	00001	1
SUIT	00048	8	TEETHE	00000	1	TOW	00001	1
SULK	00000	1	TEMPT	00002	2	TRACE	00023	7
SUN	00112	10	TEND	00043	8	TRACK	00038	8
SUP	00001	1	TENSE	00015	6	TRADE	00143	10
SURF	00001	1	TERM	00079	9	TRAIL	00031	7
SURGE	00009	5	TEST	00119	10	TRAIN	00082	9
SWAB	00000	1	THANK	00036	8	TRAIPSE	00000	1
SWAG	00000	1	THATCH	00000	1	TRAMP	00001	1
SWAMP	00005	4	THAW	00006	4	TRAP	00020	6
SWAN	00003	3	THIEVE	00000	1	TRAWL	00000	1
SWANK	00001	1	THIN	00092	10	TREAT	00026	7
SWAP	00002	2	THIRST	00004	3	TREE	00059	9
SWARM	00003	3	THRASH	00001	1	TREK	00002	2
SWASH	00000	1	THREAD	00015	6	TRENCH	00002	2
SWAT	00000	1	THRESH	00000	1	TREND	00046	8
SWATHE	00000	1	THRILL	00005	4	TRICK	00015	6
SWAY	00005	4	THROB	00000	1	TRILL	00003	3
SWEAT	00023	7	THRONG	00003	3	TRIM	00020	6
SWELL	00007	4	THRUM	00000	1	TRIP	00081	9
SWERVE	00002	2	THUD	00003	3	TROLL	00000	1
SWIG	00002	2	THUMB	00010	5	TROOP	00016	6
SWILL	00000	1	THUMP	00003	3	TROT	00012	5

Table L.11: Regular verbs used in network training

Stem	KFf	l	Stem	KFf	l	Stem	KFf	l
TROUNCE	00000	1	VOW	00002	2	WHIZ	00002	2
TRUDGE	00000	1	WAD	00000	1	WHOOP	00001	1
TRUMP	00001	1	WADE	00002	2	WHOP	00000	1
TRUSS	00000	1	WAFT	00000	1	WIELD	00001	1
TRUST	00052	8	WAG	00000	1	WIG	00001	1
TRY	00140	10	WAGE	00056	9	WILL	02244	16
TUBE	00031	7	WAIL	00003	3	WILT	00003	3
TUCK	00002	2	WAIT	00094	10	WINCE	00000	1
TUG	00003	3	WALK	00100	10	WINCH	00000	1
TUNE	00010	5	WALL	00160	11	WIND	00063	9
TURF	00003	3	WALTZ	00001	1	WING	00018	6
TURN	00233	11	WANE	00000	1	WINK	00007	4
TUT	00000	1	WANK	00000	1	WIPE	00010	5
TWANG	00000	1	WANT	00329	12	WISH	00110	10
TWEAK	00000	1	WAR	00464	13	WOLF	00006	4
TWEET	00000	1	WARM	00067	9	WOO	00003	3
TWIG	00000	1	WARN	00011	5	WORD	00274	12
TWIN	00007	4	WARP	00004	3	WORM	00004	3
TWINE	00000	1	WASH	00037	8	WORST	00034	8
TWIRL	00000	1	WASTE	00035	8	WOUND	00028	7
TWIST	00018	6	WATCH	00081	9	WRAP	00005	4
TWIT	00000	1	WAVE	00046	8	WREATHE	00000	1
TWITCH	00003	3	WAX	00014	6	WRECK	00008	5
TYPE	00200	11	WEAN	00000	1	WRENCH	00000	1
UP	01895	16	WEB	00006	4	WRITHE	00002	2
URGE	00021	7	WEDGE	00004	3	WRONG	00129	10
USE	00589	13	WEE	00005	4	YACHT	00004	3
USE	00589	13	WEED	00001	1	YAK	00000	1
VAMP	00001	1	WEIGH	00004	3	YANK	00007	4
VAULT	00002	2	WELD	00004	3	YAP	00000	1
VAUNT	00000	1	WELL	00897	14	YARN	00014	6
VEER	00002	2	WELSH	00004	3	YAW	00000	1
VEIL	00008	5	WENCH	00000	1	YAWN	00002	2
VEND	00000	1	WHACK	00001	1	YEARN	00001	1
VENGE	00000	1	WHANG	00000	1	YELL	00009	5
VENT	00010	5	WHEEL	00056	9	YELP	00002	2
VEST	00004	3	WHEEZE	00000	1	YEN	00003	3
VET	00001	1	WHELP	00000	1	YIELD	00035	8
VEX	00001	1	WHIFF	00001	1	YOKE	00003	3
VIE	00000	1	WHINE	00004	3	YOWL	00000	1
VIEW	00186	11	WHIP	00019	6	ZAP	00000	1
VOICE	00226	11	WHIRL	00003	3	ZIP	00001	1
VOID	00010	5	WHIRR	00000	1	ZONE	00011	5
VOTE	00075	9	WHISK	00000	1	ZOOM	00000	1

Table L.12: Regular verbs used in network training

Appendix M

PlaNet Network Training Script

```
#!/bin/csh

umask 077

source ./RunNet

nset eta 0.10
nset alpha 0
nset errlimit 0

network net/n.3layernet

foreach t (`count 1 5`)

pattern pat/p.trainingset.full.Z
network reset -s$t

set WTSpref = weights/w.3layernet.trial$t

savefile $WTSpref.0
save
savefile close
grep -v "command pattern" $WTSpref.0 > $WTSpref.0-
mv $WTSpref.0- $WTSpref.0
compress $WTSpref.0

foreach i (`count 1 250`)
nset randomize on
savefile $WTSpref.$i
cycle 1
```

```
save
savefile close
grep -v "command pattern" $WTSpref.$i  > $WTSpref.$i-
mv $WTSpref.$i- $WTSpref.$i
compress $WTSpref.$i
end

nset randomize off
pattern pat/p.testingset.Z
printlist clear
printlist Output
present all | grep -v Pattern > outputs/trial$t.$i.testoutput
printlist clear
cat outputs/trial$t.$i.testoutput|bin/runclose
> results/trial$t.$i.testresults
compress outputs/trial$t.$i.testoutput
concatenate results/desiredtestresults results/trial$t.$i.testresults
> results/trial$t.$i.comparison
bin/desiredandactual results/desiredtestresults
results/trial$t.$i.testresults > results/trial$t.$i.errors

end

quit
netclear
```

Appendix N

Regression Data on Hidden Units

Node	B	Standard Error	β	t	Sig.	Correlations		
						Ø-order	Partial	Part
(Constant)	1.871	3.138		.596	.551			
N1	-3.209E-02	.072	-.041	-.444	.657	-.025	-.011	-.004
N2	.152	.062	.120	2.443	.015	.077	.060	.024
N3	-.208	.046	-.083	-4.529	.000	-.036	-.110	-.045
N4	-.864	4.315	-.002	-.200	.841	-.019	-.005	-.002
N5	-.338	.039	-.154	-8.559	.000	.050	-.205	-.085
N6	-.232	.095	-.400	-2.450	.014	-.010	-.060	-.024
N7	-.121	.053	-.118	-2.305	.021	-.045	-.056	-.023
N8	-2.051E-02	.055	-.016	-.372	.710	-.036	-.009	-.004
N9	.193	.098	.149	1.969	.049	.002	.048	.020
N10	-.683	.490	-.047	-1.394	.163	.070	-.034	-.014
N11	-7.326E-02	.071	-.049	-1.026	.305	-.038	-.025	-.010
N12	-.319	.818	-.006	-.390	.696	-.034	-.010	-.004
N13	.524	.064	.120	8.208	.000	.106	.197	.082
N14	-6.923E-03	.068	-.005	-.102	.919	-.061	-.002	-.001
N15	.475	.371	.024	1.278	.201	.034	.031	.013
N16	-.135	.027	-.165	-4.964	.000	-.097	-.121	-.049
N17	.126	.069	.114	1.821	.069	-.080	.045	.018
N18	.282	.450	.022	.627	.531	.080	.015	.006
N19	.763	.057	.166	13.264	.000	.203	.309	.132
N20	.902	1.005	.015	.897	.370	.014	.022	.009
N21	-9.783E-02	.088	-.057	-1.106	.269	-.022	-.027	-.011
N22	1.354E-02	.043	.014	.314	.754	-.005	.008	.003
N23	-.162	.203	-.018	-.796	.426	-.011	-.019	-.008
N24	-.569	.184	-.053	-3.101	.002	-.142	-.076	-.031
N25	-.322	.027	-.185	-12.029	.000	-.027	-.282	-.120
N26	-5.047E-02	.075	-.033	-.672	.502	.035	-.016	-.007
N27	.361	.075	.058	4.800	.000	.151	.117	.048
N28	-.348	.075	-.458	-4.667	.000	-.008	-.113	-.046
N29	.106	.069	.090	1.551	.121	.014	.038	.015
N30	-.175	.026	-.235	-6.824	.000	.057	-.165	-.068
N31	-.136	.024	-.176	-5.609	.000	-.093	-.136	-.056
N32	-1.008	.708	-.023	-1.424	.155	.014	-.035	-.014
N33	2.125E-02	.058	.008	.368	.713	.114	.009	.004
N34	.192	.021	.264	9.180	.000	-.064	.219	.091
N35	-5.314E-02	.085	-.053	-.626	.531	.007	-.015	-.006
N36	.254	.065	.065	3.929	.000	.237	.096	.039
N37	-6.402E-02	.021	-.090	-3.002	.003	.062	-.073	-.030
N38	-7.972E-02	.067	-.051	-1.188	.235	-.073	-.029	-.012
N39	4.906E-02	.072	.030	.678	.498	.055	.017	.007
N40	-8.110E-02	.070	-.065	-1.161	.246	-.055	-.028	-.012
N41	.232	.016	.380	14.896	.000	.156	.342	.148
N42	.195	.148	.104	1.314	.189	.035	.032	.013
N43	-.938	.081	-.130	-11.531	.000	-.141	-.272	-.115
N44	-.395	.062	-.382	-6.382	.000	-.084	-.154	-.063
N45	-.151	.028	-.186	-5.381	.000	-.074	-.131	-.054
N46	-.104	.076	-.081	-1.368	.171	-.026	-.033	-.014
N47	-.173	.074	-.110	-2.338	.019	-.060	-.057	-.023
N48	.252	.070	.087	3.618	.000	.048	.088	.036
N49	.204	.063	.187	3.257	.001	-.009	.079	.032
N50	.117	.023	.079	5.090	.000	.027	.124	.051

Table N.1: Coefficients and correlations for hidden units

Node	B	Standard Error	β	t	Sig.	Correlations		
						Ø-order	Partial	Part
N51	-.481	2.428	-.003	-.198	.843	.024	-.005	-.002
N52	4.391E-02	.044	.045	1.003	.316	-.047	.025	.010
N53	-.546	.068	-.096	-8.011	.000	-.181	-.192	-.080
N54	-1.682E-02	.092	-.009	-.182	.855	-.054	-.004	-.002
N55	-.136	.026	-.179	-5.280	.000	-.066	-.128	-.053
N56	.266	.023	.341	11.579	.000	-.108	.273	.115
N57	.328	.032	.258	10.225	.000	.010	.243	.102
N58	1.327	3.042	.007	.436	.663	.022	.011	.004
N59	.344	.028	.249	12.265	.000	.113	.287	.122
N60	.105	.076	.090	1.380	.168	-.061	.034	.014
N61	-.363	.059	-.185	-6.117	.000	.015	-.148	-.061
N62	.400	.041	.171	9.735	.000	.040	.232	.097
N63	.272	.021	.469	12.671	.000	-.008	.296	.126
N64	-3.713E-02	.044	-.038	-.840	.401	-.069	-.021	-.008
N65	-.685	.218	-.047	-3.148	.002	-.104	-.077	-.031
N66	.670	.076	.126	8.796	.000	.169	.210	.087
N67	.262	.016	.385	15.941	.000	.099	.363	.159
N68	-.832	.957	-.016	-.870	.384	.036	-.021	-.009
N69	-7.704E-02	.067	-.052	-1.156	.248	.060	-.028	-.011
N70	6.646E-02	.050	.062	1.342	.180	-.033	.033	.013
N71	.488	.046	.191	10.568	.000	.197	.250	.105
N72	.375	.030	.292	12.591	.000	.076	.294	.125
N73	9.950E-02	.040	.034	2.471	.014	.104	.060	.025
N74	.291	.070	.064	4.134	.000	-.027	.101	.041
N75	.447	.122	.054	3.678	.000	.135	.090	.037
N76	-.432	.080	-.071	-5.390	.000	-.149	-.131	-.054
N77	5.423E-02	.074	.041	.731	.465	-.060	.018	.007
N78	-2.478E-02	.097	-.016	-.257	.797	.035	-.006	-.003
N79	1.816E-02	.050	.016	.365	.715	-.146	.009	.004
N80	.126	.046	.124	2.731	.006	-.029	.067	.027
N81	-6.382E-02	.050	-.062	-1.279	.201	-.172	-.031	-.013
N82	-6.693E-03	.073	-.005	-.092	.927	-.089	-.002	-.001
N83	1.623E-02	.063	.014	.258	.797	.072	.006	.003
N84	-7.763E-02	.055	-.070	-1.423	.155	-.021	-.035	-.014
N85	-.117	.072	-.043	-1.620	.105	-.036	-.040	-.016
N86	2.216E-03	.065	.002	.034	.973	.001	.001	.000
N87	.101	.095	.056	1.061	.289	-.001	.026	.011
N88	.167	.020	.220	8.472	.000	-.064	.203	.084
N89	.793	.095	.135	8.373	.000	.181	.201	.083
N90	-5.077E-02	.154	-.010	-.330	.741	.025	-.008	-.003
N91	.227	.099	.109	2.304	.021	.042	.056	.023
N92	.141	.099	.097	1.425	.154	.001	.035	.014
N93	.158	.084	.112	1.892	.059	.010	.046	.019
N94	.393	.190	.036	2.065	.039	.092	.050	.021
N95	.314	.073	.221	4.277	.000	.069	.104	.043
N96	-.233	.144	-.030	-1.623	.105	-.020	-.040	-.016
N97	-.165	.020	-.227	-8.113	.000	.042	-.195	-.081
N98	-.215	.018	-.256	-11.676	.000	-.115	-.275	-.116
N99	.550	.154	.069	3.566	.000	.139	.087	.035
N100	-.239	.054	-.165	-4.426	.000	.000	-.108	-.044

Table N.2: Coefficients and correlations for hidden units

Node	B	Standard Error	β	t	Sig.	Correlations		
						Ø-order	Partial	Part
N101	-3.423E-02	.046	-.035	-.741	.459	.078	-.018	-.007
N102	-1.031	.307	-.059	-3.357	.001	-.063	-.082	-.033
N103	.569	.133	.057	4.288	.000	.058	.104	.043
N104	.173	.018	.260	9.706	.000	-.088	.231	.097
N105	2.118E-02	.047	.009	.454	.650	-.039	.011	.005
N106	.233	.018	.266	12.850	.000	.017	.300	.128
N107	-5.897E-02	.050	-.053	-1.178	.239	-.040	-.029	-.012
N108	-.105	.062	-.086	-1.696	.090	-.071	-.041	-.017
N109	.158	.029	.198	5.527	.000	.024	.134	.055
N110	-4.408E-02	.073	-.031	-.604	.546	-.012	-.015	-.006
N111	-.268	.058	-.224	-4.611	.000	-.004	-.112	-.046
N112	-.414	.777	-.008	-.533	.594	.009	-.013	-.005
N113	-.147	.020	-.208	-7.490	.000	-.059	-.180	-.074
N114	.187	.023	.259	8.127	.000	-.018	.195	.081
N115	-1.935E-02	.071	-.013	-.271	.786	-.048	-.007	-.003
N116	-2.823E-02	.078	-.019	-.360	.719	-.017	-.009	-.004
N117	8.089E-03	.025	.011	.325	.745	.084	.008	.003
N118	-.159	.113	-.030	-1.408	.159	.060	-.034	-.014
N119	1.772	1.252	.023	1.415	.157	-.024	.035	.014
N120	-.267	.019	-.305	-13.827	.000	-.082	-.321	-.138
N121	-.318	1.491	-.003	-.213	.831	-.037	-.005	-.002
N122	-.177	.028	-.224	-6.258	.000	-.138	-.151	-.062
N123	.428	.075	.081	5.709	.000	.144	.138	.057
N124	.267	.044	.283	6.025	.000	.005	.146	.060
N125	.437	.059	.111	7.348	.000	-.050	.177	.073
N126	-.182	.015	-.260	-11.916	.000	-.014	-.280	-.119
N127	.763	.042	.218	18.071	.000	.257	.404	.180
N128	.300	.430	.012	.699	.485	-.009	.017	.007
N129	.106	.081	.096	1.308	.191	-.045	.032	.013
N130	-.149	.016	-.243	-9.357	.000	-.004	-.223	-.093
N131	-.332	.067	-.261	-4.975	.000	.014	-.121	-.049
N132	-.386	.083	-.089	-4.642	.000	.068	-.113	-.046
N133	-.109	.037	-.068	-2.939	.003	-.058	-.072	-.029
N134	-.257	.100	-.053	-2.575	.010	.103	-.063	-.026
N135	1.713E-02	.085	.011	.201	.840	.072	.005	.002
N136	-.265	.029	-.145	-8.999	.000	-.033	-.215	-.090
N137	-3.449E-02	.172	-.003	-.201	.841	-.110	-.005	-.002
N138	7.505E-03	.043	.008	.173	.862	.049	.004	.002
N139	-6.371E-02	.179	-.008	-.355	.723	.024	-.009	-.004
N140	3.061E-02	1.479	.000	.021	.983	.032	.001	.000
N141	.207	.024	.175	8.592	.000	.116	.206	.085
N142	.298	.128	.033	2.325	.020	-.049	.057	.023
N143	-.257	.021	-.333	-12.203	.000	.145	-.286	-.121
N144	.788	.926	.015	.851	.395	.000	.021	.008
N145	-.109	.076	-.089	-1.435	.151	-.093	-.035	-.014
N146	-.420	.141	-.034	-2.979	.003	-.044	-.073	-.030
N147	-7.581	3.294	-.038	-2.301	.021	-.008	-.056	-.023
N148	.295	1.639	.003	.180	.857	-.012	.004	.002
N149	-.117	.033	-.134	-3.517	.000	-.094	-.086	-.035
N150	1.653E-02	.058	.014	.287	.774	-.025	.007	.003

Table N.3: Coefficients and correlations for hidden units

Node	B	Standard Error	β	t	Sig.	Correlations		
						Ø-order	Partial	Part
N151	-.324	.024	-.375	-13.645	.000	-.075	-.317	-.136
N152	-9.672E-02	.071	-.066	-1.356	.175	-.089	-.033	-.013
N153	.141	.094	.094	1.506	.132	-.071	.037	.015
N154	-5.986E-02	.033	-.072	-1.810	.070	-.016	-.044	-.018
N155	.312	.345	.024	.905	.366	.063	.022	.009
N156	-6.504E-02	.073	-.088	-.896	.370	-.035	-.022	-.009
N157	.245	.948	.004	.259	.796	-.072	.006	.003
N158	.196	.075	.125	2.629	.009	-.007	.064	.026
N159	-9.111E-03	.134	-.001	-.068	.946	-.024	-.002	-.001
N160	-.231	.026	-.321	-9.001	.000	-.054	-.215	-.090
N161	.127	.065	.088	1.956	.051	-.066	.048	.019
N162	.282	.016	.422	17.344	.000	.005	.391	.172
N163	-.279	.029	-.168	-9.741	.000	.095	-.232	-.097
N164	-.676	.039	-.281	-17.205	.000	-.328	-.388	-.171
N165	-.213	.082	-.214	-2.607	.009	.006	-.064	-.026
N166	.136	.053	.119	2.546	.011	-.082	.062	.025
N167	.103	.018	.146	5.580	.000	-.045	.135	.055
N168	.158	.017	.229	9.149	.000	.073	.218	.091
N169	.663	.422	.024	1.574	.116	-.024	.038	.016
N170	.314	.049	.135	6.390	.000	-.048	.154	.064
N171	-.495	.192	-.058	-2.581	.010	.044	-.063	-.026
N172	-.272	.028	-.198	-9.694	.000	-.091	-.231	-.096
N173	.204	.038	.221	5.368	.000	-.033	.130	.053
N174	.504	.028	.290	17.845	.000	.046	.400	.177
N175	-.142	.018	-.186	-8.079	.000	-.086	-.194	-.080
N176	-3.239	1.460	-.032	-2.218	.027	-.025	-.054	-.022
N177	.196	.129	.084	1.513	.130	.004	.037	.015
N178	4.505E-02	.021	.061	2.137	.033	-.004	.052	.021
N179	-.153	.022	-.199	-6.892	.000	-.028	-.166	-.069
N180	-.143	.056	-.130	-2.552	.011	-.019	-.062	-.025
N181	-.120	.041	-.129	-2.933	.003	-.092	-.072	-.029
N182	2.906	7.481	.005	.388	.698	-.016	.010	.004
N183	.100	.057	.093	1.772	.077	-.101	.043	.018
N184	.181	.076	.119	2.366	.018	-.065	.058	.024
N185	.268	.072	.211	3.732	.000	.038	.091	.037
N186	-3.482	4.723	-.010	-.737	.461	-.025	-.018	-.007
N187	-.165	.064	-.144	-2.573	.010	.016	-.063	-.026
N188	.107	.035	.051	3.075	.002	-.083	.075	.031
N189	.479	.257	.032	1.865	.062	.129	.046	.019
N190	-.291	.058	-.229	-5.061	.000	-.164	-.123	-.050
N191	7.109E-02	.019	.088	3.818	.000	.033	.093	.038
N192	36.135	17.264	.032	2.093	.036	-.011	.051	.021
N193	-.423	.239	-.054	-1.770	.077	.027	-.043	-.018
N194	.264	.191	.025	1.388	.165	.172	.034	.014
N195	-9.793E-03	.114	-.007	-.086	.932	.069	-.002	-.001
N196	3.538E-02	.064	.029	.555	.579	-.105	.014	.006
N197	-7.964E-02	.024	-.094	-3.274	.001	.079	-.080	-.033
N198	-.125	.073	-.092	-1.712	.087	-.028	-.042	-.017
N199	-.326	.029	-.300	-11.208	.000	-.121	-.264	-.111
N200	.255	.051	.197	4.974	.000	.088	.121	.049

Table N.4: Coefficients and correlations for hidden units

References

Ainsworth-Darnell, K., Shulman, H. & Boland, J. (1998). Dissociating brain response to syntactic and semantic anomalies: Evidence from event-related potentials. *Journal of Memory and Language, 38*, 112–130.

Anderson, J. (1982). Acquisition of cognitive skill. *Psychological Review, 89*, 369–406.

Baayen, R., Piepenbrock, R. & Gulikers, L. (1995). *The CELEX lexical database (CD-ROM).* Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.

Bates, E. & Elman, J. (1993). Connectionism and the study of change. In M. Johnson (Ed.), *Brain development and cognition: A reader.* Oxford: Blackwell.

Beck, M.-L. (1997). Regular verbs, past tense and frequency: Tracking down a potential source of NS/NNS competence differences. *Second Language Research, 13*, 93–115.

Berko, J. (1958). The child's learning of English morphology. *Word, 14*, 150–177.

Breyer, N. (1997). *Neurophysiological studies of semantic organisation in the human brain.* M.Sc. thesis, University of Oxford.

Brown, R. (1973). *A first language: The early stages.* Cambridge, MA: Harvard University Press.

Bybee, J. & Moder, C. (1983). Morphological classes as natural categories. *Language, 59*, 251–270.

Bybee, J. & Slobin, D. (1982). Rules and schemas in the development and use of the English past tense. *Language, 58*, 265–289.

Chomsky, N. & Halle, M. (1968). *The sound pattern of English.* New York: Harper & Row.

Coulson, S., King, J. & Kutas, M. (1998). Expect the unexpected: Event-related brain response to morphosyntactic violations. *Language and Cognitive Processes, 13*, 21–58.

Craik, F. & Lockhart, R. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behaviour, 11*, 671–684.

Daugherty, K. & Seidenberg, M. (1992). Rules or connections? The past tense revisited. In *Proceedings of the 14th annual conference of the Cognitive Science Society* (pp. 259–264). Pittsburgh, PA: Cognitive Science Society.

Daugherty, K. & Seidenberg, M. (1994). Beyond rules and exceptions: A connectionist approach to inflectional morphology. In S. D. Lima, R. L. Corrigan & G. K. Iverson (Edd.), *The reality of linguistic rules.* Amsterdam and Philadelphia, PA: John Benjamins.

Davies, M. (1995). Two notions of implicit rules. *Philosophical Perspectives: AI, Connectionism and Philosophical Psychology, 9*, 153–183.

DeKeyser, R. (1997). Beyond explicit rule learning: Automatizing second language morphosyntax. *Studies in Second Language Acquisition, 19*, 195–222.

Demb, J., Desmond, J., Wagner, A., Stone, M., Vaidya, C., Glover, G. & Gabrieli, J. (1995). Semantic encoding and retrieval in the left inferior prefrontal cortex: A functional MRI study of task difficulty and process specificity. *Neuroscience, 15*, 5870–5878.

Ellis, N. (1996). Sequencing in SLA: Phonological memory, chunking and points of order. *Studies in Second Language Acquisition, 18*, 91–126.

Ellis, N. & Schmidt, R. (1998). Rules or associations in the acquisition of morphology? The frequency by regularity interaction in human and PDP learning of morphosyntax. In K. Plunkett (Ed.), *Language and Cognitive Processes (special issue).* Hove, UK: Psychology Press.

Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D. & Plunkett, K.

(1996). *Rethinking innateness: A connectionist perspective on development.* Cambridge, MA: MIT Press.

Elman, J. & Hare, M. (1997). Single mechanism = single representation? No!! In *Proceedings of the 19th annual conference of the Cognitive Science Society* (p. 907). Pittsburgh, PA: Cognitive Science Society.

Ervin, S. (1964). Imitation and structural change in children's language. In E. Lenneberg (Ed.), *New direction in the study of language.* Cambridge, MA: MIT Press.

Eubank, L. & Gregg, K. (1995). Et in amygdala ego?: UG, (S)LA and neurobiology. *Studies in Second Language Acquisition, 17,* 35–58.

Francis, N. & Kučera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar.* Boston, MA: Houghton-Mifflin.

Friederici, A., Hahne, A. & Von Cramon, D. (1998). First-pass versus second-pass parsing processes in a Wernike's and Broca's aphasic: Electrophysiological evidence for a double dissociation. *Brain and Language, 62,* 311–341.

Friederici, A., Pfeifer, E. & Hahne, A. (1993). Event-related brain potentials during natural speech processing: Effects of semantic, morphological and syntactic violations. *Cognitive Brain Research, 1,* 183–192.

Friederici, A., Wessels, J., Emmorey, K. & Bellugi, U. (1992). Sensitivity to inflec-

tional morphology in aphasia: a real-time processing perspective. *Brain and Language, 43*, 747–763.

Gabrieli, J., Desmond, J., Demb, J., Wagner, A., Stone, M., Vaidya, C. & Glover, G. (1996). Functional magnetic resonance imaging of semantic memory processes in the frontal lobes. *Psychological Science, 7*, 278–283.

Glushko, R. (1979). The organisation and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance, 5*, 674–691.

Goldman-Rakic, P. (1988). Topography of cognition: Parallel distributed networks in primate association cortex. *Annual Review of Neuroscience, 11*, 137–156.

Hagoort, P., Brown, C. & Groothusen, J. (1993). The syntactic positive shift (SPS) as an ERP measure of syntactic processing. *Language and Cognitive Processes, 8*, 439–483.

Halle, M. & Mohanan, K. (1985). Segmental phonology of modern English. *Linguistic Inquiry, 16*, 57–116.

Jaeger, J., Lockwood, A., Kemmerer, D., Van Valin Jr, R., Murphy, B. & Khalak, H. (1996). A Positron Emission Tomographic study of regular and irregular verb morphology in English. *Language, 72*, 451–497.

Jared, D., McRae, K. & Seidenberg, M. (1990). The basis of consistency effects in word naming. *Journal of Memory and Language, 29*, 687–715.

Juola, P. & Zimmerman, P. (1996). Whole-word phonetic distances and the PGPfone alphabet. In *Proceedings of the international conference on spoken language processing.* Philadelphia, PA.

Just, M., Carpenter, P., Keller, T., Eddy, W. & Thulborn, K. (1996). Brain activation modulated by sentence comprehension. *Science, 274*, 114–116.

King, J. & Kutas, M. (1995). Who did what and when: Using word- and clause-level ERPs to monitor working memory usage in reading. *Journal of Cognitive Neuroscience, 7*, 376–395.

Kirsner, K. (1994). Implicit processes in second language learning. In N. Ellis (Ed.), *Implicit and explicit learning of languages.* London: Academic Press.

Kluender, R. & Kutas, M. (1993). Subjacency as a processing phenomenon. *Language and Cognitive Processes, 8*, 573–633.

Kuczaj, S. (1977). The acquisition of regular and irregular past tense forms. *Journal of Verbal Learning and Verbal Behaviour, 16*, 589–600.

Kutas, M. & Hillyard, S. (1983). Event-related brain potentials to grammatical errors and semantic anomalies. *Memory and Cognition, 11*, 539–550.

Kutas, M. & Kluender, R. (1994). What is who violating? A reconsideration of linguistic violations in light of event-related brain potentials. In H. Heinze, T. Münte & G. Magnun (Edd.), *Cognitive electrophysiology.* Boston: Birkhäuser.

Kutas, M. & Van Petten, C. (1994). Psycholinguistics electrified: Event-related brain potential investigations. In M. Gernsbacher (Ed.), *Handbook of psycholinguistics.* San Diego: Academic Press.

Kučera, H. & Francis, N. (1967). *Computational analysis of present-day American English.* Providence: Brown University Press.

Lachter, J. & Bever, T. (1988). The relation between linguistic structure and associative theories of language learning: A constructive critique of some connectionist learning models. *Cognition, 28,* 195–247.

Lahiri, A. & Marslen-Wilson, W. (1991). The mental representation of lexical form: A phonological approach to the recognition lexicon. *Cognition, 38,* 245–294.

Lee, B. (1995). On the processing of regular and irregular inflections: The symbolic-connectionist debate revisited. In C. Koster & F. Wijnen (Edd.), *Proceedings of the Groningen Assembly on Language Acquisition.* The Netherlands: Centre for Language and Cognition, Groningen.

Leiner, H., Leiner, A. & Dow, R. (1993). Cognitive and language functions of the human cerebellum. *Trends in Neurosciences, 16,* 444–447.

MacWhinney, B. (1975). Rules, rote and analogy in morphological formations by Hungarian children. *Journal of Child Language, 2,* 65–77.

MacWhinney, B. (1990). *The CHILDES project: Tools for analyzing talk.* Hillsdale, NJ: Erlbaum.

MacWhinney, B. (1994). The dinosaurs and the ring. In S. D. Lima, R. L. Corrigan & G. K. Iverson (Edd.), *The reality of linguistic rules.* Amsterdam and Philadelphia, PA: John Benjamins.

MacWhinney, B. & Leinbach, J. (1991). Implementations are not conceptualizations: Revising the verb learning model. *Cognition, 40,* 121–157.

Marchman, V. (1997). Children's productivity in the English past tense: The role of frequency, phonology and neighborhood structure. *Cognitive Science, 21,* 283–303.

Marchman, V. & Callan, D. (1995). Multiple determinants of the productive use of the English regular past tense suffix. In *Proceedings of the 17th annual conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

Marchman, V., Wulfeck, B. & Ellis Weismer, C. (1999). Morphological productivity in children with normal language and SLI: A study of the English past tense. *Journal of Speech, Language and Hearing Research, 42,* 206–219.

Marcus, G., Brinkman, U., Clahsen, H., Wiese, R. & Pinker, S. (1995). German inflection: The exception that proves the rule. *Cognitive Psychology, 29,* 189–256.

Marcus, G., Pinker, S., Ullman, M., Hollander, M., Rosen, T. & Xu, F. (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development, 57*(4).

Marslen-Wilson, W. (1998). *Dissociating types of lexical computation.* (Paper Presented at the Department of Brain and Cognitive Sciences, MIT)

Marslen-Wilson, W., Hare, M. & Older, L. (1995). *Priming and blocking in the mental lexicon: The English past tense.* (Paper Presented at the Meeting of the Experimental Society, London)

Marslen-Wilson, W. & Tyler, L. (1997). Dissociating types of mental computation. *Nature, 387*, 592–594.

Marslen-Wilson, W. & Tyler, L. (1998). Rules, representations and the English past tense. *Trends in Cognitive Sciences, 2*, 428–435.

Mazoyer, B., Tzourio, N., Frak, V., Syrota, A., Murayama, N., Levrier, O., Salamon, G., Dehaene, S., Cohen, L. & Mehler, J. (1993). The cortical representation of speech. *Journal of Cognitive Neuroscience, 5*, 467–479.

McKinnon, R. & Osterhout, L. (1997). Constraints on movement phenomena in sentence processing: Evidence from event-related brain potentials. *Language and Cognitive Processes, 11*, 495–523.

McLeod, P., Plunkett, K. & Rolls, E. (1998). *Introduction to connectionist modelling of cognitive processes.* Oxford: OUP.

Minsky, M. & Papert, S. (1969). *Perceptrons.* Cambridge, MA: MIT Press.

Miyata, Y. (1991). *A user's guide to PlaNet version 5.6 : A tool for constructing, running, and looking into a PDP network.*

Morton, J. (1979). Facilitation in word recognition: Experiments causing change in the logogen model. In P. Kolers, M. Wrolstad & M. Bouma (Edd.), *Processing of visible language.* New York, NY: Plenum.

Müller, H., King, J. & Kutas, M. (1997). Event-related potentials elicited by spoken relative clauses. *Brain Research: Cognitive Brain Research, 5*, 193–203.

Münte, T., Heinze, H., Matzke, M., Wieringa, B. & Johannes, S. (1998). Brain potentials and syntactic violations revisited: No evidence for specificity of the syntactic positive shift. *Neuropsychologia, 36*, 217–226.

Neville, H., Nicol, J., Barss, A., Forster, K. & Garrett, M. (1991). Syntactically based sentence processing classes: Evidence from event-related brain potentials. *Journal of Cognitive Neuroscience, 3*, 151–165.

Norman, D. (1986). Reflections on cognition and Parallel Distributed Processing. In D. Rumelhart & J.L.McClelland (Edd.), *Parallel Distributed Processing: Explorations in the microstructure of cognition, Vol. 2: Psychological and biological models.* Cambridge, MA: MIT Press.

Oldfield, R. (1971). The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia, 9*, 97–113.

Orsolini, M. & Marslen-Wilson, W. (1997). Universals in morphological representation: Evidence from Italian. *Language and Cognitive Processes, 12*, 1–47.

Osterhout, L. & Holcomb, P. (1992). Event-related brain potentials elicited by syntactic anomaly. *Journal of Memory and Language, 31*, 785–806.

Paivio, A. (1986). *Mental representations: A dual coding approach.* Oxford: Oxford University Press.

Patel, A., Gibson, E., Ratner, J., Besson, M. & Holcomb, P. (1996). *Examining the specificity of the P600 ERP component to syntactic incongruity in language: A comparative event-related brain potential study of structural incongruity in language and music.* (Paper Presented at the 9th Annual CUNY Conference on Human Sentence Processing, New York)

Pinker, S. (1991). Rules of language. *Science, 253*, 530–534.

Pinker, S. (1994). *The language instinct.* New York: Morrow.

Pinker, S. (1995). Why the child holded the baby rabbits: A case study in language acquisition. In L. Gleitman & M. Liberman (Edd.), *Language: An invitation to cognitive science.* Cambridge, MA: MIT Press.

Pinker, S. (1999). *Words and rules.* London: Weidenfeld & Nicholson.

Pinker, S. & Prince, A. (1988). On language and connectionism. *Cognition, 28*, 73–194.

Plaut, D. (1997). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology, 17*, 291–321.

Plaut, D., McClelland, J., Seidenberg, M. & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasiregular domains. *Psychology Review, 103*, 56–115.

Plunkett, K. & Juola, P. (1999). A connectionist model of English past tense and plural morphology. *Cognitive Science, 23*, 463–490.

Plunkett, K. & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition, 39*, 43–102.

Plunkett, K. & Marchman, V. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition, 48*, 21–69.

Prasada, S. & Pinker, S. (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes, 8*, 1–56.

Prasada, S., Pinker, S. & Snyder, W. (1990). *Some evidence that irregular forms are retrieved from memory but regular forms are rule-governed.* (Paper Presented at the 31st Meeting of the Psychonomic Society, New Orleans)

Quinlan, P. (1992). *Oxford Psycholinguistic Database.* Oxford: Oxford University Press.

Raichle, M., Fiez, J., Videen, T., MacLeod, A.-M., Pardo, J., Fox, T. & Petersen, S. (1994). Practice-related changes in human brain functional anatomy during nonmotor learning. *Cerebral Cortex, 4*, 8–26.

Rayner, K. & Pollatsek, A. (1994). *The psychology of reading.* Hillsdale, NJ: Erlbaum.

Rosenblatt, F. (1959). Two theorems of statistical separability in the perceptron. In M. Selfridge (Ed.), *Mechanisation of thought processes: Proceedings of a symposium held at the National Physical Laboratory.* London: HMSO.

Rösler, F., Pütz, P., Friederici, A. & Hahne, A. (1993). Dissociation of brain activity related to syntactic and semantic aspects of language. *Journal of Cognitive Neuroscience, 5*, 345–362.

Rumelhart, D., Hinton, G. & Williams, R. (1986). Learning internal representations by error propagation. In D.E.Rumelhart & J.L.McClelland (Edd.), *Parallel Distributed Processing: Explorations in the microstructure of cognition, Vol. 1: Foundations.* Cambridge, MA: MIT Press.

Rumelhart, D. & McClelland, J. (1986). On learning the past tenses of English verbs. In D. Rumelhart & J.L.McClelland (Edd.), *Parallel Distributed Processing: Explorations in the microstructure of cognition, Vol. 2: Psychological and biological models.* Cambridge, MA: MIT Press.

Schmahmann, J. (1991). An emerging concept: The cerebellar contribution to higher function. *Archives of Neurology, 48*, 1178–1187.

Schriefers, H., Friederici, A. & Van Kampen, A. (1994). The time course of semantic and syntactic activation during language production in agrammatics. *Brain and Language, 47*, 455–456.

Seidenberg, M. (1989). Visual word recognition and pronunciation: A computational model and its implications. In W. Marslen-Wilson (Ed.), *Lexical representation and process.* Cambridge, MA: MIT Press.

Seidenberg, M. & Bruck, M. (1990). *Consistency effects in the generation of past tense morphology.* (Paper Presented at the 31st Meeting of the Psychonomic Society, New Orleans)

Seidenberg, M. & Hoeffner, J. (1998). Evaluating behavioral and neuroimaging evidence about past tense processing. *Language, 74*, 104–122.

Shao, J. (1995). *Analyzing semantic processing using event-related brain potentials (ERPs).* (unpublished manuscript)

Snodgrass, J. & Vanderwart, M. (1980). A standardised set of 260 pictures: Norms for name agreement, image agreement, familiarity and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory, 6*, 174–215.

Sommer, B. (1980). The shape of Kunjen syllables. In D. Goyvaerts (Ed.), *Phonology in the '80s.* Ghent: Story-Scientia.

Squire, L. & Knowlton, B. (1995). Memory, hippocampus and brain systems. In M. Gazzaniga (Ed.), *The cognitive neurosciences.* Cambridge, MA: MIT Press.

Stanners, R., Neiser, J., Hernon, W. & Hall, R. (1979). Memory representation for morphologically related words. *Journal of Verbal Learning and Verbal Behaviour, 18,* 399–412.

Stemberger, J. (1993). Vowel dominance in overregularizations. *Journal of Child Language, 20,* 503–521.

Ullman, M. (1993). *The computation of inflectional morphology.* Unpublished doctoral dissertation, Department of Brain and Cognitive Sciences, MIT.

Ullman, M., Corkin, S., Coppola, M. & Hickok, G. (1997). A neural dissociation within language: Evidence that the mental dictionary is part of declarative memory, and that grammatical rules are processed by the procedural system. *Journal of Cognitive Neuroscience, 9,* 266–276.

Ullman, M., Corkin, S., Pinker, S., Coppola, M., Growden, J. & Locascio, J. (1994). *Which neural structures subserve language? Evidence from inflectional morphology.* (Paper presented at the First Annual Meeting of the Cognitive Neuroscience Society, San Francisco)

Ullman, M., Corkin, S., Pinker, S., Coppola, M., Locascio, J. & Growden, J. (1993). *Neural modularity in language: Evidence from Alzheimer's and Parkinson's*

diseases. (Paper presented at the 23rd Annual Meeting of the Society for Neuroscience, Washington D.C.)

Ward, G. (1997). *Moby pronunciator.* 3449 Martha Ct., Arcata, CA, USA. (Also available at *http://www.dcs.shef.ac.uk/research/ilash/Moby/index.html*)

Wicha, N. (1998). The structure of language: An integrative exploration of syntactic phenomena. (unpublished manuscript (University of California at San Diego: La Jolla, CA))

Wickelgren, W. (1969). Context-sensitive coding, associative memory and serial order in (speech) behavior. *Psychological Review, 76*, 1–15.

Wise, R., Chollet, F., Hadar, U., Friston, K., Hoffner, E. & Frackowiack, R. (1991). Distribution of cortical networks involved in word comprehension and word retrieval. *Brain, 114*, 1803–1817.

Xu, F. & Pinker, S. (1995). Weird past tense forms. *Journal of Child Language, 22*, 531–556.