Inhibitory Mechanisms in Visual Word Recognition

David R. Harvey

Thesis submitted for the degree of
Doctor of Philosophy in Psychology

Department of Psychology
University of Warwick

June 2021

Contents

Acknowledgements **vii**

Declaration **viii**

Abstract **ix**

1 Introduction **1**

1.1 Background 1

1.1.1 Visual Word Recognition 1

1.1.2 Lexical Decision Task 2

1.1.3 Masked Priming 3

1.1.4 Computational Models 4

1.2 Research Questions 5

1.3 Chapter Summary 6

2 Literature Review **8**

2.1 Lexical Inhibition 8

2.1.1 The Interactive Activation Model 8

2.1.2 Experimental Evidence 10

2.2 N and F 14

2.2.1 Multiple Read-Out Model 19

2.2.2 Shared Neighbours 21

2.2.3 Beyond Coltheart's N 22

2.3 Letter Position Encoding 25

2.4 Individual Differences 30

3 Neighbourly Priming 36
3.1 Introduction 36
3.2 Experiment 1.1 37
3.2.1 Method 37
3.2.2 Results 39
3.2.3 Discussion 40
3.3 Experiment 1.2 41
3.3.1 Method 42
3.3.2 Results 42
3.3.3 Discussion 44
3.4 Experiment 1.3 44
3.4.1 Method 45
3.4.2 Results 45
3.4.3 Discussion 47
3.5 General discussion 48

4 One-Letter Priming 50
4.1 Introduction 50
4.2 Experiment 2.1 50
4.2.1 Method 51
4.2.2 Results 52
4.2.3 Discussion 53
4.3 Experiment 2.2 54
4.3.1 Method 54
4.3.2 Results 55
4.3.3 Discussion 56
4.4 Experiment 2.3 57
4.4.1 Method 57
4.4.2 Results 58
4.4.3 Discussion 59
4.5 Experiment 2.4 59
4.5.1 Method 60
4.5.2 Results 61
4.5.3 Discussion 62

4.6 Experiment 2.5 62
4.6.1 Method 62
4.6.2 Results 63
4.6.3 Discussion 64
4.7 General discussion 65

5 Overcoming Backward Masking 66
5.1 Experiment 3 66
5.1.1 Method 67
5.1.2 Results 70
5.1.3 Discussion 72

6 Sandwich Priming 74
6.1 Experiment 4 74
6.1.1 Method 74
6.1.2 Results 75
6.1.3 Discussion 77

7 Simulations 81
7.1 Introduction 81
7.2 Method 81
7.2.1 Software 82
7.2.2 Model Details 82
7.2.3 Stimuli 84
7.2.4 Trials 84
7.3 Results 84
7.3.1 Neighbourly IA 84
7.3.2 Base IA 85
7.4 Discussion 85

8 Conclusion 89
8.1 Aim of the Thesis 89
8.2 Summary of Results 89
8.2.1 Experiments 1.1 - 1.3 89
8.2.2 Experiments 2.1 - 2.5 90

8.2.3 Experiment 3 90
8.2.4 Experiment 4 90
8.2.5 Simulations 91
8.3 Integrated Discussion 91
8.4 Concluding Remarks 96
8.4.1 Limitations and Future Research 96

References 100

Appendix A Experimental Stimuli 111

Appendix B Simulation Parameters 119

List of Tables

3.1 Experiment 1.1 Mean Response Times 39
3.2 Experiment 1.2 Mean Response Times 43
3.3 Experiment 1.3 Filler Trials 46
3.4 Experiment 1.3 Mean Response Times 47

4.1 Experiment 2.1 Mean Response Times 52
4.2 Experiment 2.2 Mean Response Times 55
4.3 Experiment 2.3 Mean Response Times 58
4.4 Experiment 2.4 Mean Response Times 61
4.5 Experiment 2.5 Mean Response Times 63

5.1 Experiment 3.1 Targets 69
5.2 Experiment 3.1 Mean Response Times 70

6.1 Experiment 3.2 Mean Response Times 76

A.1 Stimuli used in Experiment 1.1 118

B.1 Parameter settings for both the Neighbourly and Base IA Model simulations run in easyNet. 120

List of Figures

7.1 Word representation activation levels over time (cycles) when the Neighbourly IA Model was presented with the trial *sobl-CURE*. 87

Acknowledgements

I am deeply grateful to my supervisor, Dr. James S. Adelman, as the completion of this thesis would not have been possible without his advice and guidance. I would also like to express my sincere gratitude to my PhD Mentor, Dr. Kate Messenger, whose wisdom and support were invaluable. My thanks go too to all of the staff of the University of Warwick Psychology Department, and particularly to the members of the Language and Learning Research Group, for helping to provide an enjoyable and educational studentship.

The writing of this thesis would not have been possible without the endless encouragement of my wife, Ruth, and the support of my friends at King's Church Birmingham; I am so thankful.

Lastly, I would like to express my gratitude to the University of Warwick Psychology Department for funding my PhD studentship.

Declaration

No material contained within this thesis has been published or otherwise used previously.

This thesis is the candidate's own work and has not been submitted for a degree at another university.

Abstract

Many models of visual word recognition rely on a mechanism of lexical inhibition: competitive connections between the representations of visually similar words. This thesis seeks to explore whether effects attributed to lexical inhibition could instead be explained by an alternative mechanism, involving "neighbourly" inhibition based on non-homogeneous inhibitory connections between letter and word representations, based on the presence of visually similar words (orthographic neighbours) in the lexicon. An introduction to visual word recognition is provided, along with some key methodological concepts: the lexical decision task, masked priming and computational modelling, before exploring research literature relevant to visual word recognition and the effects of orthographically similar words. Neighbourly priming is then explored with a series of experiments utilising the masked priming lexical decision methodology, with reaction times and error rates analysed using linear mixed-effects modelling. In the final two experiments, an alphabetic decision task is deployed with a novel collection of non-English foil letters, and sandwich priming is used to explore the effects of introducing a target preprime before neighbourly primes. Computational modelling is used to compare a model that implements the hypothesised neighbourly priming mechanism with one that does not. Finally, the evidence for neighbourly priming is discussed, concluding that neighbourly inhibition is a potential alternative explanation for at least some of the effects attributed to lexical competition.

Chapter 1

Introduction

1.1 Background

1.1.1 Visual Word Recognition

How do we read? What is it about the visual features of a written word that enables a reader to recognise it? Proficient readers can quickly and automatically identify words written in different cases and different fonts, which immediately tells us that the visual system must have some way of extracting abstract identities from these various visual signals. Is a word recognised merely by virtue of having the right letters in the right order? That can't be the case, or we'd never be able to recognise words that contain spelling mistakes, let alone read the infamous 'Cambridge Letter' internet meme[1]:

> Aoccdrnig to a rscheearch at Cmabrigde Uinervtisy, it deosn't mttaer in waht oredr the ltteers in a wrod are, the olny iprmoetnt tihng is taht the frist and lsat ltteer be at the rghit pclae. The rset can be a toatl mses and you can sitll raed it wouthit porbelm. Tihs is bcuseae the huamn mnid deos not raed ervey lteter by istlef, but the wrod as a wlohe.

There is much to consider to get a complete picture of reading, or

[1] Please see https://www.mrc-cbu.cam.ac.uk/people/matt.davis/cmabridge/ for a detailed exploration of the origins and veracity of the meme.

visual word recognition. As well as thinking about how features combine to make letters which combine to make words (orthography), we must think about phonology, spelling-sound consistencies, semantics, orthotactic and phonotactic constraints, context and the effects of other learned languages, to name but a few topics (Adelman, 2012). This thesis will, by necessity, focus on a tiny fraction of this larger whole; specifically, the interactions between letters and words that lead to lexical access. I leave other issues, such as how words combine to form sentences, how recognised words are read aloud, and how meaning is constructed from a recognised word, to other researchers.

Visual word recognition is a long-studied area of cognitive psychology, and the research described by this thesis (by both myself and others) takes two main forms: masked form priming, in tasks such as lexical decision, and computational models of lexical access.

1.1.2 Lexical Decision Task

A common way to determine whether and how quickly the visual form of a word can be recognised is to ask the reader to make a lexical decision, whereby they respond "Yes" if the word is a word they recognise (normally specifically from their native language), and respond "No" if the word is a nonword (a word they do not recognise; every possible unfamiliar letter sequence could technically be a word that the reader does not yet have in their vocabulary). The assumption is that in order for a reader to assert that the collection of letters in front of them matches a word in their vocabulary, they must have recognised and accessed the word. It should be noted however, that this lexical decision could be made based on other factors. If all of the nonword foils used in the experiment violate the language's orthotactics (e.g., the letter combination 'CV' never appears in English words, and 'Q' almost never appears without 'U', except in loan words such as 'NIQAB'), then the 'wordlikeness' of the stimuli may be sufficient for high accuracy on the task. This is discussed further in the literature review in Chapter 2. This is not a new problem; Balota and Chumbley (1984) argued that lexical decision may be a poor measure of lexical access, because not all tasks that would rely on lexical access are as heavily affected by frequency.

Assuming the lexical decision task (LDT) does offer a meaningful insight into word recognition, then it provides two metrics, response times and accuracy rates, which tell us how quickly and how well words are recognised across participant groups. We can then try to investigate how a variety of word factors affect their speed and accuracy of recognition. Word frequency is one of the most studied effects of this kind, with more frequent words being read more quickly and accurately than lower frequency counterparts (e.g. Monsell et al., 1989). Part of the problem with comparing groups of words in this way is ensuring that they are also balanced on all other variables which might affect processing speed and accuracy, such as length, imageability, age of acquisition, or valence (see Adelman, 2012, for other lexical factors that require controlling in such experiments).

1.1.3 Masked Priming

These methodological issues helped to motivate a new way of investigating lexical processes with a lexical decision task: masked priming (Forster, 1987; Forster & Davis, 1984). When a stimulus (known as the prime) is presented for a very short duration (e.g., around 50ms) and masked by other visual stimuli appearing directly before and after the presentation, the viewer is normally not consciously aware of the the masked stimulus. It is, however, processed by the visual system and can influence (or "prime") reaction times and accuracy rates for an immediately following target stimulus. The most robust effect is that of identity priming - if a target word is primed by the same word presented in a different case (e.g., ghost-GHOST), then lexical decisions on the target word are made more quickly and more accurately than if the target word is preceded by an unrelated word (e.g., music-GHOST). In these paradigms, the target acts as the backward mask for the prime, and a forward mask is added. This can either be a specific visual mask or non-alphabetic characters that cover the area of the prime; ##### is a commonly-used forward mask.

Masked priming allows for us to compare the effects of a range of primes on the same targets, instead of separating target words based on a characteristic and attempting to introduce controls so that the separation variable is hopefully the only one that will influence responses. With 50ms primes, participants are normally unaware that there has been a stimulus between the mask and the target, though some researchers use prime durations of around 60-70ms, for which the participant can often report the presence of an extra stimulus before the target, but is usually unable to accurately report on its identity (Forster et al., 2003). This yields the further advantage that priming effects are unlikely to be influenced by strategic choices on the part of the participant, although the proportion of related trials does appear to affect the magnitude of priming (Bodner et al., 2006), suggesting a sensitivity to context that could be argued to reflect strategic choices. Alongside the advantages in making strategic choices more difficult, when conscious perception of the prime is obscured by the forward mask, priming effects appear to rely upon the prime being masked (Forster, 1987). These factors, along with evidence that masked priming appears to be due to almost purely lexical processes, rather than prelexical, makes masked priming a very useful method for investigating visual word processing.

Priming effects that have been researched include: repetition priming effects (with primes commonly referred to as identity primes) where the prime is the same word as the target (keep-KEEP); form priming effects, where the prime has a similar orthographic form to the target, normally one-letter different, and the prime can either be a word (weep-KEEP) or a nonword (teep-KEEP); transposition priming effects, where the prime contains the same letters as the target, but with changes to the order of those letters (ekep - KEEP); morphological priming effects, where the prime and the target share a morphological stem (kept-KEEP); and a semantic priming effects, where the prime and the target share a meaning or association (hold-KEEP) (Forster et al., 2003).

1.1.4 Computational Models

Computational modelling is a compelling tool for helping to describe and understand the patterns of results we see in experimental data, such as from masked priming lexical decision task experiments. For a given input, computations are carried out according to predefined parameters, and some sort of output is generated. In this specific sub-domain of cognitive psychology, those outputs most commonly take the form of a selection of a single word (based on the relative activation of word representations within the model's lexicon) or a lexical decision ("Yes" for a word, and "No" for a nonword). This allows for a direct comparison between human data and the results simulated by the model in a way that is more valuable than mere verbal descriptions of cognition processes, which can often be problematic. Verbal descriptions can often be inherently ambiguous, difficult to falsify, and may contain contradicting assumptions that are not readily apparent. Computational models, however, ensure that the theories can be tested against data.

There is a wealth of available models that each attempt to explain visual word recognition. Many are discussed in this thesis. But the majority of discussion of computational models is given over to the Interactive Activation (IA) model, which is described in detail in Chapter 2. It is arguably the most influential model of lexical access in recent history, with numerous descendants (by which I mean models that have been directly inspired by the IA model, or incorporated some of its mechanisms or features). Complete descriptions of these models are beyond the scope of this thesis. Instead my point is merely that even though more recent models may be able to simulate experimental data more accurately (particularly for prime types such as transposition primes, as I will outline later), there are assumptions and mechanisms contained within the IA model that deserve further probing. This thesis attempts to do that, particularly with regard to lexical competition.

1.2 Research Questions

Many models of visual word recognition, particularly the Interactive Activation model and models it has inspired, rely heavily on lexical competition mechanisms to explain a range of experimental findings, including the apparent inhibitory effect of orthographic neighbours on lexical access. The aim of this thesis was to explore whether non-homogeneous inhibitory connections between letter and word representations exist and so could at least partially explain effects previously attributed to a lexical competition mechanism.

1.3 Chapter Summary

Chapter 2 is a literature review, describing previous research into visual word recognition. It focuses on the Interactive Activation model of lexical access and the predictions it makes with regards to lexical inhibition. I explore the effects of the number and frequency of orthographic neighbours on lexical access, as well as what can be learned from research into letter position effects and individual differences about how letters and words might be represented in the lexicon.

Chapter 3 describes my first three experiments investigating neighbourly priming. I use masked primed lexical decision tasks, and introduce a novel type of prime: the neighbourly prime, in which each letter of the prime is taken from a different neighbour of the target word.

Chapter 4 details five experiments that use single-letter primes in order to examine whether neighbourly letters in isolation can have an inhibitory effect on word access.

Chapter 5 describes an alphabetic decision task experiment, using English and "NonEnglish" letters, manipulating prime identity, prime duration, and the presence of a backward mask, in order to find the conditions necessary for a single letter prime to exert priming influence.

Chapter 6 explains my attempt to use the sandwich priming paradigm to further explore the effects of neighbourly primes.

Chapter 7 introduces a modified version of the Interactive Activation model of lexical access in which non-homogeneous inhibitory connections between the letter and word levels has been implemented. The results of simulating this model on the stimuli used in Experiments 1.1 - 1.3 are compared with those from the original Interactive Activation model.

Finally, **Chapter 8** summarises the aim of the thesis and the results from my experiments and simulations. Support for a neighbourly mechanism in lexical access is explored in an integrated discussion before I present my concluding remarks and make recommendations for future research.

Chapter 2

Literature Review

2.1 Lexical Inhibition

2.1.1 The Interactive Activation Model

Motivated in part by findings that demonstrated that letters are more readily recognised in the context of words than alone or in the context of nonwords (Reicher, 1969), in 1981, McClelland and Rumelhart presented the Interactive Activation model (IA model) of word perception. The model assumes three levels of processing, each containing nodes that represent relevant units: the feature level, the letter level and the word level. Each node is connected to some of the other nodes[1]. Connections between layers are excitatory or inhibitory depending upon consistency. Consistent nodes, such as a word node and the nodes for letters within that word, have excitatory connections between them. Between-layer connections are limited to adjacent levels (i.e., feature-letter and letter-word, but not feature-word).

Connections within the word level are inhibitory, and these inhibitory connections are applied homogeneously and non-specifically (although, as I will discuss later, Davis & Lupker, 2006, have shown that selective inhibitory connections for word nodes with overlapping letters provide a

[1]McClelland and Rumelhart refer to nodes sharing a connection as *neighbours*, though this term is used in this thesis to refer to words that differ by one letter, i.e., orthographic neighbours. In the original implementation of IA model, neighbouring nodes within the word level are not necessarily orthographic neighbours

better fit for experimental data). Nodes at rest have an activation value or level at or below zero, with higher frequency words having higher resting levels than lower frequency words. Nodes are said to be active or inactive if their activation level is positive or not, respectively, and when no longer receiving input, activation levels decay back to their resting values. Each node's activation level is calculated based on a combination of the activation levels of its excitatory and inhibitory neighbours. Each connection has an associated weight value, and it is the product of the weight and a node's activation level that determines its contribution to the activation level of its neighbour, with inactive nodes not influencing other nodes. This homogeneity also applies to between-layer connections; non-consistent between-layer connections (e.g., S(1) →TAKE[2]) are uniformly inhibitory.

Stimuli are presented to the model through the activation of feature inputs. Activation in feature nodes in turn facilitates letters nodes for letters containing those features and homogeneously inhibits all other letters, whilst the active letter nodes attempt to suppress one another. Letter nodes then activate or inhibit word nodes in the same way. Activation can also flow back 'down' to early layers through feedback connections. Excitation and inhibition can feed forward and back, hence *'interactive activation'*. This is also how the model successfully simulates the perceptual advantage of letters in words shown by Reicher (1969). The original model simulations run by McClelland and Rumelhart (1981) were restricted to four-letter words with occurrence rates at or above two times per million according to word counts by Kučera and Francis (1967).

The feature of the IA model that this thesis will focus on is the homogeneity of inhibitory connections. I believe that non-homogeneous letter-word inhibitory weights could provide at least a partial explanation for the apparent lexical competition effect. The model predicts that activation in word nodes will be facilitated by orthographically similar nonwords, but that activation in one word node will suppress activity in the others. The results from empirical research into the effect of ortho-

[2]Here, S(1) symbolises a representation for the letter S in the first letter slot. A further explanation of slot-coding is presented on page 25.

graphically similar words and nonwords has been varied across the years, and much of that research has taken the form of lexical decision tasks with masked primes.

2.1.2 Experimental Evidence

Investigating the repetition effect, whereby the accuracy and speed of responses to previously presented words are increased in lexical decision tasks, Forster and Davis (1984) developed a paradigm in which the first presentation of each word (the prime) was masked. This is the first instance of a masked priming lexical decision task and, although changes to the procedure have been made, it is still widely used today and forms the basis for the majority of the empirical research presented in this thesis. Each trial consisted of a lowercase unrelated word (the mask) presented for 500ms, followed by a 60ms lowercase prime, followed by a 500ms presentation of the target word in uppercase. Having successfully demonstrated a lexical repetition effect, in their second experiment Forster and Davis investigated whether transfer effects could be seen for primes that differed from the target by a single medial letter (e.g., past-POST). They discovered no priming effect for orthographically related word primes, however one-letter-different nonword primes did elicit facilitation. Forster (1987) further demonstrated that this priming effectiveness is affected by the neighbourhood density (N) of the target, and that word neighbour primes to low-N targets can provide a priming effect as strong (and as facilitatory) as nonword neighbours.

Segui and Grainger (1990) carried out further masked neighbour priming research, this time with a forward mask consisting of hash marks (####) instead of an unrelated word. They found, using French stimuli, that a word prime of higher frequency than the target results in an inhibitory priming effect. This, unlike the findings from Forster and Davis (1984) and Forster (1987), supports predictions made by the IA model in which word-level representations have competitive connections (McClelland & Rumelhart, 1981). The presentation of a prime increases this competition, inhibiting recognition of the target. Segui and Grainger argue that when the prime is of lower frequency than the target, preactivation of its representation is insufficient to cause activation levels high enough to inhibit the target representation (which is also receiving positive preactivation from the prime as it shares the majority of its letters).

Grainger and Ferrand (1994) investigated the impact of phonology in lexical retrieval by manipulating the orthographic relatedness of homophone primes in French (e.g., fois-FOIE is a neighbour and homophone pair, sans-CENT is a non-neighbour homophone pair). Whether the prime or the target was the more frequent of the pair was manipulated between conditions. They found that when the prime was higher frequency than the target, there were clear facilitative effects for both orthographically related and unrelated homophones, and slight inhibition when the prime was low frequency relative to the target. The facilitatory effect for higher frequency primes was greater for neighbour pairs than for orthographically dissimilar pairs. However, when Grainger and Ferrand introduced pseudohomophone targets into the lexical decision task, forcing participants to pay more attention to spelling, neighbour primes (homophonic or not) significantly inhibited responses to their targets. This effect was replicated with English stimuli and English-speaking participants. Grainger and Ferrand argue that, without these pseudohomophone targets, participants could be providing "Yes" responses because activation of the representation of the prime reaches a recognition threshold. The task does not discriminate between accurate identification of the target and mistaken identification of the prime.

It is also possible that these effects could be explained with the extension to the IA model proposed by Jacobs and Grainger (1992). They implemented a true lexical decision mechanism by adapting a temporal criterion mechanism originally suggested by Coltheart et al. (1977) for the logogen model, where "Yes" responses in a lexical decision task either arise from a single word representation reach threshold, or due to global lexical activity reaching a separate threshold. (This extension also allowed the generation of a "No" response.) Making the task more difficult by introducing pseudohomophones might force participants to switch strategy and stop relying on global activation levels, instead requiring the individual word representation to reach threshold activation.

Lexical inhibition has also been demonstrated in Dutch (Brysbaert et al., 2000; Drews & Zwitserlood, 1995), across English and French in bilingual participants (Bijeljac-Babic et al., 1997), and for one-letter-addition and one-letter-subtraction primes (again in Dutch; De Moor & Brysbaert, 2000). Lexical inhibition in masked priming paradigms appears to be a robust effect.

In their masked priming experiments, Forster and Veres (1998) found a strong facilitatory effect for nonword neighbour primes, but no inhibition was found for word neighbour primes. This paradigm made use of nonword targets that were also one letter different from legal words. When these items were exchanged for nonword targets that did not bear a resemblance to actual words (or are two letters removed from English words, as in Forster and Veres' Experiment 3), the word neighbour primes produced facilitation equivalent to that from nonword neighbour primes. This parallels the results from Grainger and Ferrand (1994); a difficult lexical decision task showed either inhibitory or insignificant levels or priming, whilst the easier task elicited facilitatory priming. Forster and Veres offer several explanations for this effect, including a checking process that is engaged after the target entry is located, which destroys priming effects from word primes when there are very word-like nonwords present as targets. They also acknowledge the global activation version of the IA model proposed by Jacobs and Grainger (1992), but ultimately conclude that "there is little to choose between the postaccess checking account... and the global versus local activation account of the present results. Each account has merits, but none is entirely satisfactory" (Forster & Veres, 1998, p. 509).

Indeed, there are findings that are hard to reconcile with a global activation account, such as from Evett and Humphreys (1981). They found equivalent facilitatory priming for one-letter-different word and nonword primes in a word identification task where the prime was forward masked and the target was backward masked. Both prime and target had display durations between 36ms and 54ms. Accurate responses in this paradigm cannot rely on global activation levels, so why do the one-letter-different primes cause facilitation relative to controls instead of inhibition? A potential explanation is that the short presentation of primes (mean duration 43ms instead of the 50-60ms priming duration in lexical decision tasks discussed above) does not cause sufficient activation levels in representations of neighbours to the target to result in competititve processes manifesting. In the IA model, nodes do not actively suppress competitors until they reach a certain activity threshold (Davis, 2003; McClelland & Rumelhart, 1981). It seems likely that the prime might not suppress the target, or the target suppress the prime given that the target representation receives evidence over a much shorter time frame (and therefore lower activation levels) than in lexical decision task experiments. The other crucial difference is that word identification can only report accuracy scores, not reaction times (RTs). It could be that target recognition is inhibited in terms of the time, but not accuracy, though it is not clear by what mechanism RTs could be inhibited but accuracy facilitated.

In 2003, Davis described in detail the factors underlying masked priming effects in the IA model. These include frequency difference between prime and target, prime neighbourhood frequency, target neighbours and target-only neighbour frequency. In simple terms, there are various competitors of the target and some are more important than others. For the IA model, the only relevant competitors are direct competitors of the target, i.e., orthographic neighbours, due to its strict coding scheme and parameter settings. Davis' theoretical work and development of simple regression models to capture priming effects opened the way for further research into the effects seen in human participants, particularly around shared neighbours and relative frequencies. Several predictions from the IA model developed by Davis were subsequently put to the test in human participants by Davis and Lupker (2006). For example, for the first time, both facilitatory priming effects of nonwords and inhibitory priming effects of word primes were demonstrated simultaneously in a single experiment. This inhibitory effect was stronger (and only attained significance for) low-frequency targets preceded by high-frequency primes. This interaction with target frequency was not predicted by IA model simulations, and previous research did not reveal such an effect (Sereno, 1991). Davis and Lupker also demonstrated that a better fit for the experimental data was produced by a modified version of the IA model that they refer to as the *selective inhibition model*. Letter-level activation was reset upon presentation of the target, to simulate the masking effect upon the prime, and instead of homogeneous lateral inhibition, word nodes would only receive inhibitory signals from orthographically overlapping word nodes. That is to say that word nodes in this model only compete with other word nodes that share at least one letter in the same position (BEST and BOOT will compete, but DESK and TREE will not). This selective inhibition model also produced the best fit (compared with the standard IA model, and a version implementing only the letter reset modification described above) for Davis and Lupker's experimental data showing inhibitory priming for target and neighbour word prime pairs that share a neighbour (e.g., short-SNORT, where the shared neighbour is SPORT).

Whilst Davis and Lupker attribute these results to lexical inhibition, there is a potential alternative explanation. The inhibitory signals to the target node could come from the letter layer, specifically from the letter that is different in the neighbour prime compared with the target. Henceforth, I will refer to such letters as *neighbourly* letters. An inhibitory signal from the letter node for X(2) to the word node for ABLE would help to suppress the activity received from the letters A(1), L(3), and E(4) and prevent AXLE being confused for ABLE, which has a processing advantage in the IA model, as more frequent words have higher resting activation baselines.

2.2 N and F

So far this exploration of lexical inhibition has focused on direct competition between orthographic neighbours, that is, where the prime is one letter different from the target. There is, however, another important way in which lexical inhibition appears to affect the processing of the target: the size of its neighbourhood. Coltheart et al. (1977) defined *N* (or *Coltheart's N*) as the number of words that can be created by changing a single letter in the target word (originally suggested by Landauer & Streeter, 1973). This measure of neighbourhood size was shown to influence nonword classification; Coltheart et al. (1977) showed that high-N nonwords (nonwords with many neighbours) took longer to be classified as nonwords than low-N nonwords. This makes sense if we consider that many words will receive partial activation due to the presentation of a nonword that shares all but one of their letters, and if we assume that the global activation level in the lexicon inhibits a "No" response (or that a checking process is needed if word activation approaches a threshold). N did not affect the classification speed for words, but as discussed earlier, Forster (1987) found that word neighbour primes to low-N targets can provide a priming effect as strong (and as facilitatory) as nonword neighbours.

In 1989, Grainger et al. again showed no effect of N on word processing with a non-primed paradigm. Slower lexical decision and longer eye-fixation latencies were instead shown for words if one of its neighbours was of a higher frequency, compared with words that had no higher frequency neighbours. In the same year, however, Andrews showed a facilitatory effect of N for low-frequency words in a lexical decision task in which she manipulated N and frequency (F) orthogonally (Andrews, 1989). These findings appear to be contradictory. The IA model predicts the finding from Grainger et al., as the presentation of the target word will also partially activate the representation belonging to a neighbour. If that neighbour is a higher frequency word than the target, it has a higher resting threshold, giving it a boost to activation, helping it to suppress activation of the target. The same situation would be predicted for the experiments run by Andrews, as a low-frequency word with many neighbours is more likely to have neighbours of higher frequencies, which ought to suppress the target identification. Andrews herself points out in a later paper that, due to the dominance of accounts of lexical retrieval like the IA framework, "researchers have been reluctant to accept evidence that is incompatible with their expectation that competing neighbors will interfere with performance, and even more reluctant to believe that similar neighbors may actually facilitate lexical retrieval" (Andrews, 1997, p. 440).

Grainger and Jacobs (1996) attribute the discrepancy between the above findings to task-specific specific processes, whilst Andrews (1997) highlights that facilitatory effects tend to be found with English stimuli and English-speaking participants, whilst inhibitory effects have been found in French (Grainger & Jacobs, 1996) and Spanish (Carreiras et al., 1997). This distinction is important, not only because it cautions against generalising from the results of experiments conducted in a single language. Inhibitory processes may simply be less powerful in English compared with other European languages, or perhaps differences in orthographic density mean that our definition of orthographic neighbours may need to differ from language to language.

Forster and Shen (1996) instead concluded (tentatively) that neighbourhood density does not influence lexical access time, but does influence the lexical decision process via a bias mechanism. Their first experiment showed a facilitatory effect of N in a LDT, with RTs and error rates decreasing as N increased. However, when they manipulated N and the presence of a higher-frequency neighbour across their word stimuli, this did not effect the facilitation from N; higher-frequency neighbours did not interfere. This was true both with nonword stimuli that were very word-like (thus making the task very difficult) and when the task was made easier by using nonwords without word neighbours.

Dutch native speakers (L1) with English as a second language (L2) were faster to respond to English words with more English neighbours, whilst more Dutch neighbours for words in both English and Dutch resulted in slower responses in a language-non-specific lexical decision task carried out by Van Heuven et al. (1998). The presence of a cross-language inhibitory effect suggests that orthographic neighbours across all languages the reader knows are candidates for inhibitory effect. Furthermore, the facilitation for English neighbours appearing alongside inhibition for Dutch neighbours suggests at language differences in the role of neighbours, although another explanation is that the English words are effectively lower frequency words for Dutch L1 speakers, and so less able to exert inhibitory influences.

Overall, studies conducted in English usually find null or facilitatory effects of neighbourhood frequency (e.g., Forster & Shen, 1996; Siakaluk et al., 2002), but Perea and Pollatsek (1998) did find inhibitory effects of neighbourhood frequency in English with their non-primed LDT experiment. The methodology of this study may account for this unusual finding. Sears et al. (2006) point out that Perea and Pollatsek stressed accuracy over speed in their instructions to participants, which is likely to particularly affect the processing of low-frequency words, and that low-frequency words are responsible for driving the significant effect found by Perea and Pollatsek. Sears et al. replicated this experiment, but introduced a condition wherein they manipulated the instructions given to participants and stressed either accuracy (as Perea and Pollatsek had) or accuracy and speed. When accuracy alone was stressed, the replication was successful; there was an inhibitory effect for low-frequency words that had a higher-frequency neighbour compared with those without a higher-frequency neighbour. This effect was only present for low-frequency words, and did not appear when both accuracy and speed were stressed in the participant instructions. Furthermore, when Sears et al. repeated the experiment with new stimuli, they failed to replicate the results, suggesting that the finding of a neighbourhood frequency effect in English was due to a combination of the instructions and the specific stimuli. Perea and Pollatsek used relatively infrequent words in their experiment, many of which resulted in large enough RTs so as to be excluded from the analysis, and many of their nonword foils were very similar to real English words. Perhaps sensitivity to a higher-frequency neighbour is not a function of typical lexical activation, but instead a conscious checking process? Sears et al. argue that English might require weaker lexical inhibition than other languages:

"This neighborhood structure for English words (i.e., larger neighborhoods and many higher frequency neighbors) may necessitate a lexical processor with weaker inhibitory connections than those in other languages. Otherwise, it might be extremely difficult for low-frequency words to accumulate enough activation to reach their identification thresholds" (Sears et al., 2006, p. 1059).

This speaks to a central paradox in the lexical inhibition account. Lateral lexical inhibition is apparently necessary to prevent confusion between similar words, but combined with higher resting activation levels for higher frequency words, it is not clear how low-frequency words are activated without being confused for their higher-frequency neighbours.

Research on how partial-word primes interact with N is also revealing. Perry et al. (2008) further investigated the changes to the IA model proposed by Davis (2003) and made by Davis and Lupker (2006) using partial-word priming. In this version of the masked priming paradigm, rather than complete words or nonwords being used as the prime, a partial word is used, with a non-alphabetic character used in place of the missing letter(s). Grainger and Jacobs (1993) used a % character (e.g., %rown) for their partial-word primes, whereas in these experiments Perry et al. used a # character (e.g., #rown). These partial-word primes can either be ambiguous, where there are 2 or more words that can be made by replacing the missing letter (e.g., #rown is an ambgiuous match for BROWN, CROWN, FROWN, or GROWN), or they can be unambiguous, where there is only one possible word that can be made by replacing the missing letter (e.g., #igar is an unambiguous match for CIGAR and no other words). They manipulated the neighbourhood of their target words, the type of partial-word prime and the difficulty of the task. In Experiment 1A and 1B, high- vs low-N target words were primed with ambiguous partial-word primes, and the difficulty of the task was manipulated between these experiments (1A used nonword foils with low-N, whilst half of the nonwords used in 1B had larger neighbourhoods). There was no effort made to manipulate whether the target word presented after the ambiguous partial-word prime was the highest frequency match. For example, s#eep - SHEEP (KF frequency norm: 24; Kučera and Francis, 1967) has a higher-frequency match of SLEEP (KF: 66), whereas clo#k - CLOCK (KF: 21) has only a lower-frequency other potential match (CLOAK; KF: 3)[3]. Their second pair of experiments used unambiguous partial-word primes, and also added ‘hermit’ words; words with no neighbours (N = 0).

[3] I calculated these frequencies and potential matches using N-Watch (Davis, 2005).

Across these experiments, the high-N and low-N targets were responded to faster than hermit targets, both for related and unrelated primes. This is evidence for a facilitatory effect of N in English, although Perry et al. (2008) argue that this is due to poor item matching; the hermit targets used had lower log CELEX frequencies and higher Age-of-Acquisition than the other targets. As mentioned above, neighbourhood frequency was not accounted for, so no conclusions can be made about this effect. One finding of note was that, for the more difficult experiment (Exp 2B), hermit words showed the largest benefit from unambiguous priming, more so than low- and high-N target words. According to the predictions of the IA model, words with neighbours (particularly higher-frequency neighbours) will have a harder time achieving the activation levels required for lexical access due to the competitive effects of their neighbours. These words should stand to gain the most benefit from an unambiguous partial-word prime, which will boost the activation of the target word with minimal support to any neighbours of the target. Hermit words have no neighbours, so are expected to receive comparably less benefit (or facilitatory priming) from such primes. Overall, these results suggest that the predictions of the version of the IA model implemented by Davis and Lupker do not reflect human performance with partial-word primes, particularly when it comes to the predictions from the lexical competition component of the model.

2.2.1 Multiple Read-Out Model

The lack of consensus over whether neighbourhood frequency effects were inhibitory or facilitatory is (in part) what prompted Grainger and Jacobs (1996) to develop their Multiple Read-Out Model (MROM). Based upon the architecture of the IA model (McClelland & Rumelhart, 1981), the MROM also incorporated three criterion for lexical decisions. The M criterion is sensitive to the activation of single lexical units, and when any one lexical unit reached this activation criterion, lexical selection had occurred and that word had been identified. The Σ criterion is sensitive to the overall lexical activation across all lexical units, and when this criterion is reached then a “word” decision can be made in LDTs, for example. Finally, the T criterion represents a temporal deadline. When sufficient time has passed without the M or Σ criteria being reached, then a "nonword" decision can be made. The parameters underlying the M criterion are assumed to be set, whereas those for Σ and T can be influenced by the nature of the task, the nature of the stimuli or the task instructions provided. For example, when speed is prioritised over accuracy and/or a LDT contains nonwords that are very dissimilar to words, then the Σ criterion is likely to be relied upon and thus lower global activation thresholds are required for a "word" response.

This flexibility in the Σ criterion, Grainger and Jacobs (1996) argued, could account for many of the discrepancies in the literature. When classifying words and nonwords is very difficult (for example, when the nonwords are very word-like) or accuracy is stressed in the task instructions, then lexical decision is made via the M criterion. Individual lexical unit activity is sensitive to lexical competition, so inhibitory effects of neighbourhood frequency can be seen. Whereas when global lexical activity is utilised via the Σ criterion, then activity in neighbouring lexical units contributes towards that threshold, causing facilitation.

Siakaluk et al. (2002) tested this prediction rigorously with MROM simulations and LDT experiments using English stimuli. The word stimuli used were high or low frequency, with small or large neighbourhoods, and had either zero or at least one higher-frequency neighbour. The difficulty of the task was manipulated across these experiments by manipulating the neighbourhood size of the nonwords used. The stimuli contained nonwords that either had no neighbours, small neighbourhoods, the same neighbourhoods as the word stimuli, or large neighbourhoods. For non-existent and small neighbourhoods, the MROM predicts that global lexical activity is sufficient to perform the task, resulting in facilitation from N. This is what was found in the experimental data. However, when the nonwords and words were matched on neighbourhood size, the RTs from participant showed the largest facilitatory effect of N seen across any of the experiments. The MROM predicts a null effect of N in this situation, as the reader cannot rely on the global lexical activity generated by presentation of the stimulus in order to make a word/nonword decision.

The MROM's predictions were also poorly supported by the data in other ways. The effects of N and F are negatively correlated according to the model, that is, larger effects of word frequency should coincide with smaller effects of neighbourhood size. The reverse was seen: larger effects of N were seen with larger effects of F. Siakaluk et al. (2002) found a facilitatory effect of neighbourhood size when the Σ criterion could not be used for accurate decision-making, and they found no evidence for an inhibitory effect of neighbourhood frequency. Instead, large neighbourhoods and the presence of higher-frequency neighbours facilitated responses to low-F words.

Overall, it appears that having many neighbours and having higher-frequency neighbours does not inhibit word recognition in English as the lexical competition component of models such as the IA model would predict. Andrews (1997) makes the case that around 80% of four-letter words in English have a higher-frequency neighbour. Is a mechanism that inhibits so many words really going to be viable? As argued by Siakaluk et al., "a lexical processor that delays the processing of the majority of words and facilitates the processing of the minority is, at best, counterintuitive" (Siakaluk et al., 2002, p. 679).

2.2.2 Shared Neighbours

A lot of the experiments discussed above that investigate N and neighbourhood frequency rely on comparisons between non-overlapping sets of words. The assumption is that the high-N and low-N words used in a LDT differ only in N, and attempts can be made to control other factors, but there's always the possibility that an unknown factor is driving the differences in RT instead. Not only do primed LDT allow for critical comparison with the same target words, but they also are sensitive to smaller effects. They also allow us to investigate another facet of neighbourhood - shared neighbours.

Nakayama et al. (2008) used masked priming LDT experiments to test the lexical competition assumption. Using only words with high accuracy scores in the English Lexicon Project (Balota et al., 2007), they found inhibition from orthographic neighbour primes irrespective of the relative frequencies of the prime and target. They demonstrated that this effect was seen only for words with high N. When the neighbouring prime-target pairs were drawn from words that have fewer than 5 neighbours, inhibition was only found when the prime was of a higher frequency than the target. IA simulations predicted that high-F primes and low-F targets would yield larger latencies than low-F primes and high-F targets, regardless of N. As a lower-frequency orthographic neighbour can cause inhibitory effects for words with large N, Nakayama et al. ran further experiments to investigate whether this was due to shared neighbours.

Orthographic neighbour pairs will have slightly different neighbourhoods. Some words will be a neighbour only of the prime, some of the target, and others will be a *shared neighbour*: a neighbour of both the prime and the target. Davis (2003) explains that shared neighbours can have a particularly inhibitory effect in priming experiments as they receive supporting activation from both the prime and the target. This predicts an inhibitory effect of shared neighbours, dependent upon their number and summed frequency, relative to the frequency of the target (Davis, 2003).

In Nakayama et al.'s (2008) Experiments 4A and 4B, which varied in terms of high or low target N, they manipulated the presence of shared neighbours. The primes were also of a lower frequency than the targets. They found significant inhibitory priming from neighbour primes for targets with many neighbours, but little evidence for inhibition for targets with fewer neighbours. This effect was not modulated by the presence of shared neighbours. Taken together, these results suggest that shared neighbours were not specifically responsible for inhibitory effects from low frequency primes, but that this was instead due to the activation of many neighbours of the target.

2.2.3 Beyond Coltheart's N

So far, the definition of a neighbour used in this thesis has been a substitution neighbour that differs by a single letter, in line with the definition proposed by Coltheart et al. (1977). This is also how the IA model

(McClelland & Rumelhart, 1981) defines neighbours and therefore which words might affect one another via lexical competition.

This interpretation of the findings from Nakayama et al. (2008) relies on this strict definition of orthographic neighbours. They point out that when the definition of a neighbour is broadened to include words that differ by two letters (instead of only one), then many prime-only neighbours become shared neighbours. Also, words that had not been previously considered as neighbours might now become eligible, some of which might be of a higher frequency than the target. In the example of tide-SIDE, TIME is a one-to-two-letter-removed shared neighbour which has a higher frequency than the target SIDE. However, Nakayama et al. argue that even when unconventional neighbours such as these are not a factor, increases in prime-N are correlated with increases in inhibitory priming. For primes with a lower frequency than the target, when the prime has only one neighbour, the priming effect on average is facilitatory. But this priming becomes more inhibitory as prime-N increases. The same correlation is also true of target-N, though to a lesser degree. These results strongly suggest that more neighbours being activated inhibits processing of the target, which is broadly in line with what is predicted by lateral inhibition.

Bowers et al. (2005) favour a wider definition of neighbours. The critical comparison in much of the research surrounding neighbourhood frequency is between words with higher-frequency neighbours and so-called ‘hermit’ words; words with no higher-frequency neighbours. Many of these hermit words, however, do have transposition, deletion and addition neighbours. Bowers et al. investigated the effect of introducing novel words to the lexicon, which would be new neighbours for previous hermit words (according to a broader definition). For example, BANANA has no substitution, transposition, addition or deletion neighbours, and so participants were taught new words such as BANARA. The central question was: would these new neighbours introduce inhibition in lexical activation for their neighbours which were previously hermits? This is indeed what was found; words that now had a neighbour had RTs slow by 17ms in a semantic categorisation task, and this inhibition increased to 48ms over two days of novel words training. Bowers et al. concluded that this inhibition was much stronger evidence for an inhibitory effect of N and, therefore, for lexical competition.

I believe that the paradigm used by Bowers et al. (2005) can hardly be considered an example of natural word learning, however. Participants learnt the orthographic forms of these new words by typing them out upon presentation, but at no point was a referent attached to these new words. It is also likely that participants in this experiment might have used overt strategies to learn these words, and these strategies could be the source of inhibition in subsequent categorisation tasks. With words to learn such as BANARA, VIODIN and TULKEY, it is not implausible that these meaningless words were learnt via associations with BANANA, VIOLIN and TURKEY and memorising the substituted letter instead of the form of the novel word in its entirety. Even if we can count these newly learnt words as if they were any other word in the lexicon, these findings could also support the mechanism proposed by this thesis; non-homogeneous letter-word inhibitory connections. The inhibition of BANANA upon learning BANARA might occur because of a change in the inhibitory connection from R(5) to BANANA.

Whilst not considered neighbours by the IA model as proposed by McClelland and Rumelhart (1981), there is reason to consider words that differ from one another in terms of letter position as neighbours. And so, let us now turn our attention to transposition priming.

2.3 Letter Position Encoding

The IA model uses a slot-coding approach, in which there is one slot for each potential letter position (the original model utilised four slots, allowing it to process four-letter words; McClelland & Rumelhart, 1981). The word CATS might be represented as C(1), A(2), T(3), S(4) whilst CAST as C(1), A(2), S(3), T(4). The important distinction here is that these slots function independently. S(3) bears no relation to S(4), meaning that, as far as the IA model is concerned, CATS is as dissimilar to CAST as it is to CAMP. This is a clear weakness of the model, as the research on transposed letter priming shows.

Chambers (1979; Experiment 2) found an 'interference effect' for anagrams. She showed that, when participants were asked to perform a lexical decision task with tachistoscopic presentation of target words that differed from a more frequent word in the position of two adjacent letters (e.g. *bale* from *able*), participants gave slower and less accurate responses to the transposed-letter targets (842ms and 88%) than to control words (787ms and 96%). This was taken as evidence that the lexical entries of visually similar words are accessed prior to decision making. This would suggest that the way in which letter position is encoded or processed (or both) is not absolutely rigid. There cannot be complete flexibility, however, or a reader would never be able to distinguish STOP from SPOT (or indeed SUPERSONIC from PERCUSSION).

Transposed Letter (TL) primes provide further evidence for this flexibility in letter position encoding. They are most commonly a nonword prime formed by changing the relative positions of two or more letters in the target word (e.g. *salior* - SAILOR) and are extremely effective primes. Forster et al. (1987) found that 60ms TL primes cause no less priming than identity primes of the same duration (approximately 63ms and 62ms priming respectively). These primes were all formed by transposing two adjacent medial letters as in the example above, and so were all visually very similar to their targets.

Guerrera and Forster (2008) failed to find such strong priming for more extreme transformations. In a series of experiments using 8-letter words, they found that whilst transposing the interior six letters (e.g. brihtady-BIRTHDAY) or the exterior four letters (e.g. ibrthdya-BIRTHDAY) produced significant priming (30ms and 23ms respectively) relative to unrelated control primes, the effect size was significantly less than for identity primes (45ms). They also found that primes in which every letter was transposed, either by swapping each letter pair (T-All primes; e.g. dineityf-IDENTIFY) or by reversing each half of the word (Reversed Halves; e.g. nediyfit-IDENTIFY) did not produce significant priming. This suggests a limit on the flexibility of letter position encoding - extreme letter transpositions create primes that no longer resemble the target word sufficiently to activate it, despite no change to letter identities. Lupker and Davis (2009) argued that these T-All primes were only ineffective because many of them activate other words more strongly than the target (e.g. for avacitno-VACATION, the word AVIATION is a stronger match for the prime). They devised a new form of masked form priming, called sandwich priming, in which an identity pre-prime is presented before the prime with a duration of 33ms. This was designed to boost the activation of the target word in order to reduce the inhibition gained from neighbours of the prime. Sandwich priming was shown to produce a significant 40ms priming effect for T-All primes, which provides strong evidence for substantial flexibility in letter position encoding, as very extreme transformations can still provide significant activation to a target word.

Changing letter order is not the only way to alter letter position information. Letters can also be inserted into the prime, which preserves the relative order whilst changing the absolute letter position. Van Assche and Grainger (2006) showed that one or two letter insertions (e.g. *gafrdlen* - GARDEN) lead to a priming effect equivalent to identity priming. This was extended by Welvaert et al. (2008), who showed a graded effect of letter insertion, with the size of the priming effect decreasing as more letters were inserted. They found that significant priming remained even up to 3 inserted letters, which resulted in a 20ms priming effect. They concluded that inserted letters was surprisingly harmless for priming lexical access, which highlights the flexible nature of orthographic coding.

There has been some debate over whether letter identity affects TL priming, specifically with regards vowels and consonants. Perea and Lupker (2004) showed that, with Spanish readers, TL priming is only present for consonant transpositions, and not for vowel transpositions. Lupker et al. (2008) replicated this result in English, but pointed out that vowels have higher frequencies than consonants. To test whether letter frequency was causing these differences, they extended the experiment to compare TL primes where the transposed letters were either high frequency or low frequency consonants. They found that only transposing low frequency consonants could produce a significant priming effect, although they did acknowledge that phonology might also play a role.

Perea and Acha (2009) also found evidence for significant TL priming with consonants but not vowels, which they explored further with a primed Same-Different task in which the participant is asked to judge whether the target is identical or different to a probe shown before the masked prime. This allowed them to more directly compare priming between word and nonword targets; a lexical decision task requires different responses for different lexical statuses, whereas with a same-different task, both words and nonwords can illicit a "same" response. They found significant priming for both transposed consonants and vowels (27ms for each) for "same" word targets. The authors describe the same-different task as "a low-level perceptual task" (Perea & Acha, 2009, p. 135), which is unaffected by phonology (e.g., Norris & Kinoshita, 2008). The finding of equivalent consonant and vowel TL priming in such a task is strong evidence for letter position encoding taking place before consonant/vowel distinctions are made by the perceptual system, so they cannot be the source of the observed differences in TL priming.

Studying evidence from developing and adult readers, Comesaña et al. (2016) concluded that the differences between vowels and consonants in letter position encoding are phonological in nature. There were able to replicate significant TL priming effects for consonants but not vowels when their participants were adults, but 9-10 year old children tested on the same stimuli showed TL priming for both consonants and vowels

(23 and 21ms respectively). These priming effects were not statistically significant, however, which Comesaña et al. attributed to the noisy data that resulted from children participants. When the same experimental material was used in a sandwich priming paradigm, they found significant TL priming for consonants and vowels (47 and 40ms respectively). They also demonstrated that developing readers of this age were sensitive to orthographic primes, but not phonological priming. This means that adults show priming for consonant but not vowel TL primes due to differences in phonological processing, which occurs after early letter position is encoded. Children, who are not sensitive to the phonological characteristics of primes, show equivalent TL priming for both vowels and consonants. Comesaña et al. argue that this is why current models of lexical access cannot predict these differences in TL priming; at time of writing there are no models that incorporate phonological processes alongside flexibility in letter position encoding.

It is clear that the IA model's slot-coding approach is a flawed one, as it cannot capture the apparent similarity between words with transposed letters, or letter addition or subtraction. Davis and Bowers (2006) argue that word recognition models should be sensitive to the relative position and contiguity of letters. Despite this very obvious weakness in the IA model as a model for lexical access, it is still an extremely influential and useful model for examining the way in which letter identities are encoded and relate to one another.

It is important to note that all of the experiments referred to in this section thus far have been carried out in languages that use the Latin alphabet (e.g. English, Spanish). Is flexibility in letter position encoding a universal feature of written language? The most well-studied non-Latin script in this area is Hebrew, which uses semitic morphology; all verbs and most nouns and adjectives are composed of a root of three consonants, and a word pattern of vowels, or vowels and consonants. Morphemes are bound together, and cannot be separated. For example, *LBS* is the root that refers to 'wearing', and *ti___o_et* is the word-pattern that denotes a feminine noun. Together they produce *tiloboset* (written as *tlbwst*; vowels are not transcribed in Hebrew writing) meaning 'a costume'. With 22 letters in the Hebrew alphabet and 3 letters in each root, it is necessary that many roots share letters. For example: *SLX* meaning 'to send', *XLS* meaning 'to dominate', *XSL* meaning 'to toughen' and *LXS* meaning 'to whisper' (Velan & Frost, 2009). Hebrew can be said to have a very dense orthography - roots with very different meanings look very similar. The Hebrew language cannot allow noisy position schemes as a small change in letter position will almost always result in a real word with a very different meaning.

Velan and Frost (2009) found no evidence for facilitatory TL priming in Hebrew, and instead discovered an inhibitory priming effect from TL primes. In a series of lexical decision tasks, they investigated three-letter primes that were: the root of the target word (identity), an anagram of the root of the target word that formed a different legal root (TL-root), an anagram of the target word root that formed a nonsense root (TL-nonsense), or three letters from the target word that were not all taken from the root (control). Identity root primes caused facilitatory priming of 19ms, whereas TL-root primes caused inhibitory priming of 9ms, whilst TL-nonsense primes caused no priming at all. There was no benefit for priming the letters of the root in the wrong order, and an explicit disadvantage when those letters formed a different, existing root. When Velan and Frost extended the experiment to primes with the same number of letters as the target (by creating TL primes that were either nonwords with an existing root contained within it, or nonwords with nonsense roots), they again found that existing root TL primes cause inhibition whilst TL nonsense roots had no priming effect. This shows that the flexibility in letter position coding evident in languages such as English or Spanish does not apply for Hebrew, which has a far denser orthography. Does this mean that readers of different scripts are employing qualitatively different processes for lexical access?

There is strong evidence that these types of differences can also be observed between speakers of English. In the next section, I will consider these individual differences and how our models of lexical access, as well as the neighbourly inhibition mechanism I am positing, can account for them.

2.4 Individual Differences

It would be wrong for us to look at the average amount of priming found in English-speaking experiments and conclude, for example, that all English readers demonstrate flexible letter position encoding. Andrews and Lo (2012) found that English readers show different priming effects depending upon their language proficiency. Participants who scored highly on reading (comprehension and speed), spelling (dictation and recognition) and vocabulary showed greater inhibitory TL priming than their less proficient counterparts. Furthermore, those who scored higher for spelling relative to their own reading and vocabulary scores showed stronger inhibitory form priming on average. Andrews and Lo make the case that many unskilled readers of English do not discriminate between orthographically similar words, likely because this is unnecessary in a language with sparse orthography. However, readers with greater language proficiency, particularly in spelling, do not seem to treat TL primes as being similar to their target words, but instead transposing letters seems to harm lexical access as evidenced by the inhibition found by Andrews and Lo. Perhaps proficient English readers and all Hebrew readers respond in the same way to TL primes because both have very specific and finely-tuned word representations. Such representations would likely be required by all readers of Hebrew, as the dense orthography is intolerant of flexibility in letter position.

Andrews warns against relying on average performance data to make conclusions about human language processing, especially from small samples of university students, and points to the considerably variety in performance, even amongst monolinguals (Andrews, 2012). These individual differences may help to explain several ambiguous situations where average data provide contrasting claims, such as the effects of N, which have been reported as inhibitory, null and facilitatory.

Forster (1987) argued that the match criterion of a target word might be tuned based on the neighbourhood density. With a high N, it is tuned so as to only accept close matches to the target. This is efficient, as it prevents neighbours of the word from being considered and reducing performance. With a low N, on the other hand, there is a broader tuning. These representations can tolerate a mismatching letter, as the pool of candidate words is very low. This is, again, efficient, as the lower match criterion leads to faster acceptance without the risk for error.

This theory was developed further by Forster and Taft (1994), who proposed that this tuning comes in the form of coding words according to their subsyllabic units. Specifically, the body (the rhyming part of the word; e.g., *ace* in *face*) or the ‘antibody’ (the onset and nucleus of the word; e.g., *bla* in *blast*). For example, the word HERD might be encoded, not as H-E-R-D, but as H-ERD. This reduces the neighbourhood density for the word, as it now has only one neighbour (NERD) and is no longer considered to be a neighbour to HEAD, HERO and HERE (Castles et al., 1999). Forster and Taft showed that high-N targets, which do not normally show form priming, could be facilitated by a nonword neighbour prime that shared a low frequency word body (e.g., perd-HERD, where the ERD body is low frequency).

This is a developmental model. Lexical representations are tuned with experience, which increases with age and reading development. This means that we can separate our participants based on their reading proficiency (and hope or assume that we find substantial individual differences) or investigate how readers respond to primes as their proficiency increases. The IA model does not account for this developmental process or differences in individual proficiency, but as we attempt to fit this model, and others like it, to experimental data, the model will be stronger for an ability to explain differences between participants (perhaps represented by changes in parameters) rather than constant parameter tweaking to force it to only represent the ‘average reader’ (if such a person exists).

In work that agreed with the findings of Forster and Taft (1994), Loncke et al. (2009) found that including legal onset-nuclei or rimes in nonwords made them harder to reject for their Dutch-speaking participants. This suggests that lexical tuning might lead to words being encoded not as a series of individual letters, but in terms of their onset, nucleus and coda[4]. Loncke et al. also produced simulations with the IA model; a version of the IA model that encoded input in terms of onset, nucleus and coda (ONC). The standard IA model only provided a good fit for the experimental data if the rime neighbours, for example, were also orthographic neighbours, whereas the ONC model matched the data more strongly. How exactly the lexicon would transition from representations of single letters to representations of larger parts of the word (such as onset-nucleus-coda) is unclear, but there does seem to be strong evidence for changing representations through language experience.

Castles et al. (1999) compared developing readers (7-11 years old) with adult readers and found the children showed a facilitatory effect of form priming from one-letter different nonwords for words from large neighbourhoods. This effect was not seen in the adult readers, nor was it seen for low-N targets. There was also evidence that the form priming effect was attenuated for the most proficient readers amongst the children. These results support the idea of lexical tuning increasing with reading proficiency. As readers become more skilled over time, they are less likely to treat words that are similar in orthography as sufficiently similar to significantly increase activation. Another explanation is that the lexicons of the adult readers contained more representations for the high-N targets' neighbours, which meant that any benefit from overlapping letters was suppressed by lexical competition. This explanation is unlikely, given that Castles et al. were careful to verify that the low/high-N distinction was still true when the children's vocabularies were taken into account.

In 2007, this Lexical Tuning Hypothesis was further supported by the finding that strong form priming seen in the responses of 8-year-olds is no longer present for the same children by the age of 10 (Castles et al., 2007). Transposed-letter priming, on the other hand, was reduced but still present in the responses of the children once they turned 10 years old, but was not seen for the adult (non-longitudinal) sample. Whilst

[4]The nucleus is the vowel within the word, and the coda is the consonant after the vowel. The word's rime is the nucleus and the consonant, whilst the onset is the consonant before the vowel. Splitting the word BEAR into onset-nucleus-coda results in B-EA-R, with the onset-nucleus being BEA- and the rime being -EAR.

there was a weak correlation between RTs and the size of the priming effect in children, this does not appear to be a strong enough relationship to suggest that the differences in priming were simply a function of taking longer to respond. Castles et al. propose that these results can be supported by the IA model, as a substitution prime is likely to also activate neighbours of the target word, particularly any words that are shared neighbours (Van Heuven et al., 2001) of the prime and the target. These shared neighbours will then inhibit the activation of the target, suppressing any beneficial form priming. Developing readers, however, might not yet have any or all of these shared neighbours in their lexicon, allowing form primes to facilitate activation of the target without any such suppression from neighbours.

Whilst plausible, this proposal would be strengthened by demonstrating that items for which form priming was shown in the responses of 8-year-olds but not adults had shared neighbours that were known to the adults but not to the children. My own examination of the stimuli would suggest that this is not the case. Approximately half of the prime-target pairs used by Castles et al. (2007) had no such shared neighbours. Of those that did, many of the shared neighbours had relatively low Age-of-Acquisition (AoA) according to norms calculated using N-Watch (Davis, 2005). For example, gast-FAST has 4 shared neighbours with AoA ratings from Bird et al. (2001): CAST, EAST, LAST, and PAST. These shared neighbours had a mean AoA of 363 (corresponding to an age of approximately 5-6 years) and a minimum AoA of 275 (corresponding to an age of approximately 3-4 years) (Bird et al., 2001). Given these ratings, it seems likely that both the 8-year-old children and the adults did have shared neighbours in their lexicon. As discussed earlier, it has been argued that neighbourhood frequency is the important factor for lexical competition from neighbours, but for this example, shared neighbour LAST has a higher frequency than the target FAST (684 vs 90; Kučera & Francis, 1967).

In order for the IA model to explain the developmental pattern shown by Castles et al. (2007), different parameters would be required for readers of different ages. Whilst it is possible that developing readers demonstrate lexical competition but adult readers do not, this is a poor explanation for a model designed to represent proficient lexical access. Instead, I believe neighbourly inhibition could account for these results. If readers gradually develop inhibitory letter-word connections with weights according to neighbourhoods, this could account for the suppression of facilitatory priming seen in the adult sample. To demonstrate with the example of gast-LAST, when the prime is shown, the letter representations for G, A, S, and T become active, which in turn leads to activation in the representations for all neighbours of GAST, including FAST and LAST. The activation of FAST feeds back to activate F(1), which has developed a stronger inhibitory connection with LAST than it has with words without a neighbour that begins with F. This suppresses the preactivation benefit that LAST receives and adults show no form priming. Developing readers, however, are still tuning the weights of these inhibitory connections with language exposure, and so have a comparatively weaker inhibitory connection between F(1) and LAST, so facilitation is not suppressed.

The IA model has a much harder time explaining the results for transposed letter primes found by Castles et al. (2007), as it uses slot coding, meaning that the prime LPAY is as similar to PLAY as a prime like BNAY. This is a weakness of the IA model, but Castles et al. point out that it could be accounted for by the variant of the SOLAR model (Davis, 1999). This model uses position-independent letter codes, with letter position encoded based on activation levels that decrease across the length of the word; the second letter receives less activation than the first letter. This spatial coding scheme was later adopted by Davis' Spatial Coding Model (Davis, 2010). According to these models, close anagrams have similar patterns of activation allowing the model to account for transposed letter priming.

These individual differences can be thought about both in terms of development over time, but also the differences in proficiency between developed readers - some adults may never reach the same level of reading proficiency as their peers. Andrews and Hersch (2010) found facilitatory form priming from nonword primes for targets from low, but not high, neighbourhoods. When participants were separated based on their spelling ability, measured by dictation and recognition tasks, Andrews and Hersch found that the better spellers showed inhibitory form priming for high-N target words, whilst the worse spellers showed facilitatory effects for both high and low-N targets. This result was based on RT data averaged over prime lexicality. The explanation offered by the authors was that good spellers have more precise lexical tuning, allowing them to activate the prime faster, which leads to the suppression of competing neighbours of the prime, including the target.

Andrews and Hersch (2010) went on to investigate the effects of ambiguous and unambiguous partial word primes, and neighbour primes, compared with unrelated primes, on five-letter targets with high and low neighbourhood densities. Better spelling ability in the participants was associated with stronger inhibitory priming for neighbour primes, and tended towards showing less priming for ambiguous vs unambiguous primes. This supports the idea that skilled readers have more precise lexical representations, which are less vulnerable to being activated by similar stimuli, and that spelling ability is the best indicator of reading skill (Andrews, 2012). Andrews and Hersch argue that this increased lexical precision could take the form of faster accumulation of evidence, or a recoding of the lexicon similar to that proposed by Forster and Taft (1994) discussed above (e.g., HERD coded as H-ERD). Within a framework like the IA model, these changes in lexical precision are likely best reflected by modifications to the connection weights. My research sets out to investigate whether we can find evidence for non-homogeneous weights in inhibitory letter-word connections.

Chapter 3

Neighbourly Priming

3.1 Introduction

This first series of experiments sought to uncover whether a learned pattern of inhibition between letter and word representations in the visual word processing system can account for the range of priming effects found in empirical data on word recognition and lexical access. According to this hypothesis, word representations might have stronger inhibitory connections with the representations of letters that appears in orthographic neighbours of the word. When reading a word with an orthographic neighbour (such as SHOP), processing a neighbourly letter (C in position 1 from CHOP) is more detrimental to accurate reading than processing a non-neighbourly letter, because this could lead to the reading of an incorrect but legal word.

A bottom-up, constraint-based system with a less uniform, more varied and adaptable pattern of letter-level and letter-to-word inhibition could explain many priming effects in empirical data. This hypothesis also predicts that a prime consisting of letters from orthographic neighbours of the target word (neighbourly primes) would inhibit processing of the target relative to an unrelated prime. For example, the target PINS might be primed with the nonword *wagt*, as each letter in the prime would make a word if substituted into the target at the same location (WINS, PANS, PIGS, PINT). Both *wagt* and a nonword like *rlov* share no letters with the target, so would typically be expected to have equal priming effects. However, if letters from orthographic neighbours have stronger inhibitory connections to word representations than unrelated letters do, then neighbourly primes like *wagt* would be expected to cause inhibitory priming relative to unrelated primes like *rlov*. This chapter describes three masked form priming experiments designed in order to test this specific prediction.

3.2 Experiment 1.1

This experiment is the first attempt to investigate the effect of neighbourly primes compared with unrelated primes. Identity primes were also included in the experimental design to confirm that the priming procedure was influencing participants' responses. The hypothesis was that identity primes would show significant facilitation compared with controls, whilst neighbourly primes would show inhibition compared with controls.

3.2.1 Method

Participants 39 participants were recruited at the University of Warwick, and participated in return for either a payment of £3.00 or as a component of an introductory undergraduate research methods class. Data from 15 participants were excluded from the analysis as their mean accuracy was less than 90%[1].

Design The experiment consisted of a lexical decision task, in which I manipulated the prime type (neighbourly, unrelated, identity) within participants. Response times and accuracy were recorded.

Materials The targets were 96 four-letter words, and 96 four-letter nonwords. All were selected because they had an orthographic word neighbour in each position. Each of the word targets had been tested in the English Lexicon Project (Balota et al., 2007) and had a mean lexical decision accuracy of at least .50.

[1] Changing this threshold to 80% did not change the pattern of results observed.

For each target, three primes were generated: neighbourly, unrelated and identity. The neighbourly primes were generated by selecting an orthographic neighbour of the target for each position, and using the letter found in the neighbour for that position. For example, the target CURE has the neighbours SURE, CORE, CUBE and CURL, which generates the neighbourly prime *sobl*. Unrelated primes were randomly generated following certain rules: each letter could only occur once in the prime, no letters found in the target could be used, any letters that could possibly be used in a neighbourly prime could not be used, and each prime could only be used once. Each identity prime was identical to the target. Each target was paired with one other target so that the neighbourly prime of one target in the pair was used as the unrelated prime for the other target, and vice versa, which removes the potential influence of compounding factors such as wordlikeness.

All items were arranged into 3 lists, so that each target appeared once and each prime appeared a maximum of once on each list. Across the three lists, each target appeared with each of its primes once.

In addition to these items, 10 practice items were added to the start of the experiment (5 words and 5 nonwords) with identity primes. These items were not included in the analysis.

The mean accuracy for word targets was 93.6%. Data from 17 word items were excluded from the analysis as their mean accuracy was below 90%. These excluded items are identified in Appendix A.

Procedure Participants were asked to decide whether each target letter string, presented on a CRT screen, was an English word or a nonword. Responses were made by pressing one of two keyboard keys. Accuracy and Response Times (RTs) were recorded. Targets were presented in a randomised order for each participant.

Each trial started with a fixation cross (+) being presented for 300ms, followed by a blank screen for 200ms. Then a forward mask (####) was presented for 500ms, followed by the lowercase prime at five-eighths size for approximately 52ms. The uppercase target was then displayed until a response was given, or until 2000ms had elapsed, at which point a “No

	List 1	List 2	List 3	Mean	Priming
Unrelated	603 (2)	645 (5)	632 (2)	627 (3)	
Identity	605 (1)	589 (2)	612 (1)	603 (2)	23 (1)
Neighbourly	643 (3)	658 (3)	639 (2)	645 (3)	-18 (0)
Mean	629 (2)	642 (3)	640 (2)	637 (2)	

Table 3.1: Experiment 1.1 mean response times (and percentage error rates) for responses to each prime type for each list. Priming is relative to Unrelated condition response times.

response" message was displayed. The only other feedback participants received was "Wrong" being displayed if the participant gave an incorrect response. All stimuli were presented in the Courier New font on a Sony CPD-G200 monitor at 1024 x 768 resolution, with the refresh rate set to 60Hz. The stimulus presentation and data collection were achieved through the use of the DMDX display system, developed by K. I. Forster and J. C. Forster at the University of Arizona (Forster & Forster, 2003).

3.2.2 Results

Data were analysed using linear mixed-effects models in R version 4.0.3 (R Core Team, 2020), using the packages *lme4* (Bates et al., 2015) and *car* (Fox & Weisberg, 2019), with significant factors being further investigated using the *emmeans* package (Lenth, 2020). Models were fitted with full random structure, and simplified if they could not converge or the model fit was singular (Barr et al., 2013).

The analysis was confined to RTs from correct trials to word targets within the range of 250 to 1500ms[2]. Means and error rates are shown in Table 3.1 above. For the latency analysis, a linear mixed-effects model was fitted with prime type and item list and their interactions as fixed factors. By-participant and by-item intercepts were added as random factors (no models including random slopes successfully converged). Type II Wald chi-square tests were used to establish the significance of

[2]6 of the correct trials (0.38%) were removed for having RTs that were too short or too long.

main effects and interactions. This revealed a significant main effect of prime type, $\chi^2(2) = 35.71$, $p < .001$.

Post-hoc comparisons revealed that responses times to targets with identity primes were significantly smaller than to those with neighbourly primes, $t(1501) = 5.97$, $p < .0001$, and unrelated primes, $t(1501) = 3.29$, $p = .001$. Response times to targets with neighbourly primes were significantly larger than to those with unrelated primes, $t(1501) = 2.69$, $p = .007$.

There was no significant effect of item list, $\chi^2(2) = 0.32$, $p = .853$, and no significant interaction between item list and prime type, $\chi^2(4) = 7.58$, $p = .108$[3].

For the error rate analysis, a linear mixed-effects model was fitted with prime type and item list and their interactions as fixed factors. By-participant and by-item intercepts were added as random factors (no models including random slopes successfully converged). There was no main effect of prime type, $\chi^2(2) = 2.13$, $p = .345$, or item list, $\chi^2(2) = 2.87$, $p = .238$, and no significant interaction, $\chi^2(4) = 2.32$, $p = .678$.

3.2.3 Discussion

There was evidence of identity priming taking place; RTs to targets with identity primes were significant shorter than those to neighbourly or unrelated primes. There was also evidence of neighbourly priming, as these RTs were slower than for unrelated primes.

This result supports the suggestion that there could be non-homogeneous inhibitory letter-to-word connections in the visual word recognition system. Each letter that makes up each neighbourly prime is important to distinguish from each letter found in the same position in the target word, as the substitution creates other English words. Therefore, it advantages the system to have a 'vertical' inhibitory connection from C(1) to BATS. The current evidence does not favour a particular solution to position encoding, so it could be that the presentation of C in any position would

[3]Using an 80% (instead of 90%) accuracy criterion still revealed the same pattern of results, with the exception of a significant interaction between item list and prime type, with the significant neighbourly priming only appearing in 1 of the 3 item lists.

cause inhibition. Given that neighbourly priming only causes 18ms of inhibition, however, it seems unlikely that changing the order of letters in the neighbourly primes would retain detectable inhibitory effects.

A potential further explanation is that each letter in the neighbourly prime causes a small amount of activation in the representations of orthographic neighbours of the target, the combined effect of which inhibits activation of the target. Given that direct priming of orthographic neighbours leads to inhibitory effects of a size between 10 and 40ms (e.g., Segui & Grainger, 1990), it seems unlikely that such an indirect effect (letters facilitating words inhibiting words) would produce priming of the same magnitude as direct word-word priming. That these results might have arisen due to currently accepted mechanisms of lateral inhibition, however, cannot be ruled out.

Whilst the items had been randomly assigned to each list, a closer inspection of the frequency of the items in each list revealed that one of the lists had a mean frequency far higher than the other two lists. Whilst this difference was not significant, it could have influenced the results, as items assigned to the same block for the purpose of constructing the lists had the same prime type in each list. Therefore, for participants in one condition, the items with the highest overall frequency all had identity primes.

3.3 Experiment 1.2

To correct for the imbalance in item frequency across lists, I designed another experiment using the same items. To attempt to increase the power of the experiment, I removed the identity primes so that the neighbourly and unrelated primes for each item would be shown to more participants. Having found identity priming already, I was confident that the experimental design was sufficient to elicit priming, meaning that identity primes were no longer required. The items were assigned to new lists, two this time, according to frequency so that the two lists had comparable mean frequencies.

3.3.1 Method

Participants 50 participants were recruited at the University of Warwick, and participated in return for either a payment of £3.00 or as a component of an introductory undergraduate research methods class. Data from 22 participants were excluded from the analysis as their mean accuracy was less than 90%.

Design and Procedure The items and design were identical to Experiment 1.1 with the following exceptions: no identity primes were shown, only neighbourly and unrelated primes. The stimuli were therefore organised into two lists, with each target and each prime being shown once in each list, and each target-prime pair being shown once across both lists.

The mean accuracy for word targets was 90.6%. Data from 25 word items were excluded from the analysis as their mean accuracy was below 90%, leaving 71 word items in the analysis.

Spelling and vocabulary tests were also added to the experiment. The spelling test was a lexical decision task with no priming, using items previous used by Adelman et al. (2014), which were based on a list from Burt and Tate (2002). 40 of the 82 targets were correctly spelt words, the remaining 40 contained a spelling error. These trials did not time-out after 2000ms. In the vocabulary test, participants were asked to choose the one of four displayed words that had the same meaning as a word shown at the top of the screen. There were 40 vocabulary items, taken from the Shipley (1940) vocabulary test, also used previously by Adelman et al. (2014).

3.3.2 Results

The analysis was confined to RTs from correct trials to word targets within the range of 250 to 1500ms[4]. Means and error rates are shown in Table 3.2 on the following page.

[4] 4 of the correct trials (0.21%) were removed for having RTs that were too short or too long.

	List 1	List 2	Mean	Priming
Unrelated	644 (2)	578 (1)	608 (2)	
Neighbourly	648 (2)	589 (3)	616 (3)	-9 (-1)
Mean	646 (2)	583 (2)	612 (2)	

Table 3.2: Experiment 1.2 mean response times (and percentage error rates) for responses to each prime type for each list.

Spelling and vocabulary scores were standardised using a z-score transformation. These factors are henceforth referred to as z-spell and z-vocab.

For the latency analysis, a linear mixed-effects model was fitted with prime type, item list, z-spell and z-vocab and their interactions as fixed factors. A by-participant intercept and a by-item intercept with a slope for item list were added as random factors. Type II Wald chi-square tests were used to establish the significance of main effects and interactions.

This revealed a significant main effect of item list[5], $\chi^2(1) = 9.55$, $p = .002$, with participants assigned to List 1 giving larger response times than those assigned to List 2. There was also a significant main effect of z-spell, $\chi^2(1) = 6.73$, $p = .009$, with participants who scored higher in the spelling task responding providing lower RTs than participants who scored lower. There was also a significant interaction between item list, z-spell and z-vocab, $\chi^2(1) = 4.08$, $p = .043$. This was due to a significant effect of spelling for participants assigned to List 1, $\chi^2(1) = 4.91$, $p = .027$, whereas for those assigned to List 2 there was no effect of spelling, $\chi^2(1) = 0.89$, $p = .346$, but the interaction between spelling and vocabulary approached significance, $\chi^2(1) = 3.63$, $p = .057$.

RTs following neighbourly primes were 9ms longer than those following unrelated primes, but this differences was not significant, $\chi^2(1) = 2.06$, $p = .151$, nor were any of its interactions, $p > .05$.

For the error rate analysis, a linear mixed-effects model was fitted with prime type, item list, z-spell and z-vocab and their interactions as

[5]This main effect was not significant when using an 80% accuracy criterion for participant inclusion. This was the only difference in the overall pattern of results for this experiment.

fixed factors. By-participant and by-item intercepts were added as random factors (no models including random slopes successfully converged). There was no main effect of prime type ($\chi^2(2) = 1.17$, $p = .279$), and none of the other effects or interactions was significant, $p > .05$.

3.3.3 Discussion

There was no evidence of neighbourly inhibition priming in the results; the RTs to targets with neighbourly and unrelated primes were not significantly different from each other. There were significant differences in how participants assigned to each list responded. This could simply be due to individual differences, with the participants randomly assigned to list 2 being faster than those assigned to list 1. This was in spite of careful counterbalancing of the target word frequencies on each list.

These results would seem to suggest that no neighbourly inhibition was taking place. This could be due to the lack of identity primes – without some of the primes providing useful information, the reading system might simply ignore all primes. Bodner et al. (2006) have demonstrated that, in repetition priming paradigms, a lower proportion of task-useful primes reduces priming effects. It is therefore likely that primes were not being processed in the same way as in Experiment 1.1.

There was also no evidence that, for this sample, the difference in RTs between neighbourly and control primes for each participant was correlated with their spelling and/or vocabulary scores. Though, if no neighbourly priming was taking place, this is not surprising.

3.4 Experiment 1.3

This third experiment was designed to help us discover whether the presence of identity primes was influencing how participants processed the neighbourly primes. I kept the items and lists from Experiment 1.2, but added filler items. In one condition, these filler items had identity primes, in the other, the filler items had unrelated primes. Due to time constraints, I was not able to implement the spelling and vocabulary tests again. However, there was very little variation in scores between participants in Experiment 1.2, suggesting that all participants were competent spellers. This is discussed further in the section on future research.

3.4.1 Method

Participants 59 participants were recruited at the University of Warwick and the University of Bristol, and participated in return for either a payment of £3.00 or as a component of an introductory undergraduate research methods class. The mean accuracy of all participants was 87.3%. Data from 25 participants were not included in the analysis as their mean accuracy was below 90%[6].

Design and Procedure The items and design were identical to those used in Experiment 1.2, except that the spelling and vocabulary tests were removed, and 96 filler items were also included. Half of these filler targets were words, half were nonwords. All the filler targets had no orthographic neighbours. Participants were assigned to one of two conditions. In the first condition, these targets were presented with identity primes. In the other condition, these targets were presented with unrelated primes (generated according to the same rules as outlined in the description for Experiment 1.1).

The mean accuracy for word targets was 93.4%. Data from 21 word items and 24 fillers were excluded from the analysis as their mean accuracy was below 90%.

3.4.2 Results

The analysis was confined to RTs from correct trials to word targets within the range of 250 to 1500ms[7].

To ascertain whether the filler trials had performed as intended, these trials were subjected to latency and accuracy analyses. The aggregated

[6] Changing this accuracy criterion to 80% did not change the pattern of results found for this experiment.

[7] 6 of the correct trials (0.30%) were removed for having RTs that were too short or too long.

Filler Type	RT (ms)	Error Rate (%)
Identity Filler	598	2
Unrelated Filler	656	3
Mean	624	2

Table 3.3: Experiment 1.3 mean RTs and error rates for responses to filler trials with Unrelated or Identity primes.

means are show in Table 3.3 above. A linear mixed-effects model was fitted with filler type (unrelated or identity prime) as the fixed factor. By-participant and by-item intercepts were added as random factors. Type II Wald chi-square tests were used to establish the significance of the main effect. This revealed a significant effect of filler type, $\chi^2(1) = 4.03$, $p = .045$, as identity-primed fillers elicited faster responses than unrelated-primed fillers. An error rate linear mixed-effects model was fitted with filler type as the fixed factor and a by-participant intercept (including a by-item intercept resulted in a singular fit), but revealed no effect of filler type, $\chi^2(1) = 0.22$, $p = .639$.

Aggregated means for the whole experiment are shown in 3.4 on the following page. To analyse the effect of prime types on RTs, a linear mixed-effects model was fitted with prime condition (unrelated or neighbourly), filler type (unrelated or identity primes) and their interaction as fixed factors. A by-participant intercept and a by-item intercept with a slope for filler type were added as random factors. Type II Wald chi-square tests were used to establish the significance of the main effects and interaction.

There was significant effect of prime type, with neighbourly inhibition of 12ms, $\chi^2(1) = 4.29$, $p = .038$. There was no main effect of filler type, $\chi^2(1) = 0.49$, $p = .484$, and the interaction was not significant, $\chi^2(1) = 2.68$, $p = .102$.

Planned comparisons revealed that there was a significant neighbourly priming effect when the filler items had identity primes, $t(1817.5) = 2.63$, $p = .009$, but there was no priming effect with unrelated filler items, $t(1809.3) = 0.16$, $p = .870$.

	Identity Fillers	Unrelated Fillers	Mean
Unrelated Primes	604 (1)	633 (2)	617 (2)
Neighbourly Primes	625 (3)	634 (2)	629 (2)
Mean	614 (2)	633 (2)	623 (2)
Priming	-21 (-2)	-1 (0)	-12 (-1)

Table 3.4: Experiment 1.3 mean response times (and percentage error rates) for responses to each prime type for lists containing Identity-Primed and Unrelated-Primed Fillers.

To investigate the error rate, a linear mixed-effects model was fitted with prime condition and filler type and their interaction as fixed factors, with by-participant and by-item intercepts. There was no significant effect of prime condition, $\chi^2(1) = 2.05$, $p = .151$, or filler type, $\chi^2(1) = 0.48$, $p = .488$, and the interaction was not significant, $\chi^2(1) = 1.69$, $p = .194$.

3.4.3 Discussion

There was evidence for identity priming, as the responses to identity-primed fillers were approximately 70ms faster than those to unrelated-primed fillers on average. There was also evidence for a neighbourly priming effect. This was only evident with participants for whom their filler trials were identity-primed. This confirms that the proportion of useful primes affects the strength (or perhaps even presence) or priming effects (e.g., Bodner et al., 2006).

As with Experiment 1.1, neighbourly primes can cause inhibitory priming in the right circumstances (i.e., when identity primes are also used in the paradigm). This helps to explain why Experiment 1.2 failed to find priming effects. The critical prime-target pairs used in Experiment 1.3 were the same as those used in Experiment 1.2, so I am confident that the crucial factor was the presence of identity fillers.

3.5 General discussion

Over the three experiments, the existence of neighbourly priming became evident. In the four conditions containing neighbourly primes, the effect was always numerically in the same inhibitory direction, but it was only significant in the two cases in which identity primes were also used in some of the trials.

The presence of inhibitory neighbourly priming under these conditions would appear to support the assertion that it is possible for single letters to inhibit the activation of word-level representations.

According to the neighbourly priming theory, this inhibition occurs because the lexical system has developed stronger inhibitory connections between letter representations and word representations for which substituting in that letter would form an orthographic neighbour. For an example, let us take the prime-target pair frod-BLEW. Substituting the letter F in at position 1 would form the word FLEW, a one-letter-different orthographic neighbour of the target BLEW. In order to prevent confusing BLEW for FLEW, there is a stronger inhibitory connection between F(1) and BLEW than there is between G(1) and BLEW, as GLEW is not a word in English.

The IA model (McClelland & Rumelhart, 1981) instead implements homogeneous letter-word inhibitory connections for any letters that do no appear in the word. It does not predict that a neighbourly letter would inhibit the target word more than an unrelated letter would. Other models of lexical access that are based upon the IA model would also not predict inhibitory neighbourly letter priming.

However, the mechanism for neighbourly inhibitory priming remains ambiguous. It could be that each of the letters that make up the neighbourly prime lend evidence to and thereby increase the activation of the representations of the target's neighbours. This slightly increased activation from four of the target's neighbours in turn inhibits the target. In this way, letter-to-word activation is still uniformly facilitatory, and word-to-word neighbour inhibition operates according to the established account.

The Spatial Coding Model (SCM Davis, 2010) is one of the models that is based upon the IA model. One of the features the two models have in common is lateral inhibition. Unlike the IA model, the SCM does not have specific inhibitory letter-word connections. Instead, word nodes get inhibited based upon the number of letters not positive contributing to the activity of the letter node. Letter identity is not considered for this mismatch inhibition.

Chapter 4

One-Letter Priming

4.1 Introduction

If four neighbourly letters can activate the target's lexical neighbours and thereby cause sufficient inhibition, it is extremely unlikely that a single letter would be able to do the same. If, however, a single letter can be demonstrated to lead to significant inhibition, this would lend further support for the notion that letter-to-word connections can be inhibitory as well as facilitatory. We designed a series of experiments to find 1) whether a single letter prime can influence the latency of the lexical decision task in the first place, and 2) whether a single neighbourly letter can slow down processing of a word.

4.2 Experiment 2.1

The first experiment in this series was designed to confirm that significant priming can be elicited using a single letter by comparing identity with unrelated single letter primes. Previous examples of partial-word primes have replaced letters in the primes with a non-alphabetic character, such as # (Perry et al., 2008) or % (Grainger & Jacobs, 1993) in order to preserve letter position information. I instead decided to use a plus symbol (+), as the percentage symbol (%) is perceptually 'busy', and the hash symbol # was used for the forward mask. These examples only replaced a single letter. Instead, I replaced all-but-one letters with these symbols. Letters in the initial and final positions were used as primes; interior letter were not used. This is because there is evidence that exterior letters have a special status in lexical access (e.g., White et al., 2008) and that, during masked presentation, they are more easily accessed than interior letters (e.g., Mewhort & Campbell, 1978), so I suspected that they would be more likely to yield significant priming effects.

4.2.1 Method

Participants 81 participants were recruited at the University of Warwick, and participated in return for a payment of £3.00. Of these, 3 were excluded for having an accuracy on word items below 90%[1].

Design The experiment consisted of a lexical decision task with masked form priming, in which we manipulated the type (identity, unrelated) and position (initial, final) of the one-letter primes within participants. Response times and accuracy were recorded.

Materials The targets were 160 four-letter words, and 160 four-letter nonwords. The initial and final position letters of each target was a consonant. Each target contained no repeated letters. All of the word targets had been tested in the English Lexicon Project (Balota et al., 2007) and had a mean lexical decision accuracy of at least .50.

For each target, four primes were generated: initial position identity, initial position unrelated, final position identity and final position unrelated. The initial position primes consisted of a single lower-case letter followed by three plus symbols (e.g. `b+++`); the plus symbol was chosen as a neutral non-alphabetical character. The final position primes consisted of three plus symbols followed by a single lower-case letter (e.g. `+++b`). For the identity primes, the letter was the same as that found in the target in the corresponding position (e.g. `b+++` – BAND). For unrelated primes, the letter was selected from the set of letters that do

[1]Changing this accuracy criterion to 80% did not change the pattern of results found for this experiment.

	Position		
	Initial	Final	Mean
Unrelated Prime	519 (4)	520 (4)	520 (4)
Identity Prime	515 (4)	510 (3)	513 (3)
Mean	517 (4)	515 (3)	516 (4)
Priming	4 (1)	9 (0)	7 (0)

Table 4.1: Experiment 2.1 mean response times and error rates for responses to each prime type (Identity and Unrelated) for each position (Initial and Final).

not appear in the target, nor in any orthographic neighbours of the target (e.g. `j+++` – BAND) In order to attempt to balance prime letter frequency across prime types, letter selection was pseudorandom, with letters given a chance of selection weighted according to their frequency; higher frequency letters were more likely to selected.

The items were arranged into four lists, so that each target appeared with each of its primes only once.

In addition to these items, 10 practice targets were added to the start of the experiment (5 words and 5 nonwords) with four-letter identity primes. These items were not included in the analysis.

The mean accuracy for word targets was 95.5%. Data from 13 word items were excluded from the analysis as their mean accuracy was below 90%.

Procedure The procedure was identical to that used in Experiment 1.1.

4.2.2 Results

Mean response times and error rates are shown in Table 4.1 above. Incorrect trials and trials with response times smaller than 250ms or longer than 1500ms were excluded from the analysis[2]. To analyse RTs, a linear

[2] 20 of the correct trials (0.20%) were removed for having RTs that were too short or too long.

mixed-effects model was fitted with prime condition (identity or unrelated), item list and prime position (initial or final) as fixed factors. By-participant and by-item intercepts with slopes for prime condition were added as random factors (no models with random slopes were able to converge). Type II Wald chi-square tests were used to establish the significance of the main effects and interactions. These revealed a significant main effect of prime condition, $\chi^2(1) = 10.75$, $p = .001$, as response times to identity-primed targets were 7ms shorter than to those with unrelated primes on average. There was no significant effect of item list, $\chi^2(4) = 0.85$, $p = .932$, or prime position, $\chi^2(1) = 0.87$, $p = .351$, and no interactions were significant, $p > .05$.

An error rate analysis was performed by fitting a linear mixed-effects model with prime condition, item list and prime position as fixed factors. By-participant and by-item intercepts with slopes for prime condition were added as random factors (no models with random slopes were able to converge). No significant effects or interactions were found, $p > .05$.

4.2.3 Discussion

I found evidence for identity priming with one-letter primes; responses to targets with identity primes were approximately 8ms faster than those to targets with unrelated primes. This occurred regardless of prime position.

I consider this to be proof that primes consisting of just a single letter are sufficient to affect processing of a target word. As far as I am aware, this is the first time that single-letter priming has been demonstrated.

Depending upon parameter settings, this effect could be predicted by the IA model. In the example of the prime-target pair p+++ – PLAY, the P(1) node will become activated, which will activate all word representations that begin with F. Whether this leads to significant facilitatory priming or not would depend on the extent to which all of these P-words inhibit one another through lexical competition. For example, PLAY and PLAN would both receive activation from P(1), but inhibit each other, according to the IA model. This inhibition might be sufficient to completely suppress any facilitation, in which case these results would further support our suggestion that lexical competition mechanisms do not reflect what is actually happening during lexical activation.

4.3 Experiment 2.2

Having successfully demonstrated facilitatory one-letter priming from both initial and final position identity letter primes, neighbourly letter primes can now be added. In the same fashion as the neighbourly primes used in the previous experiments, this is a letter taken from a neighbour of the target word. This time, only a single neighbourly letter is used rather than four at a time. If a single neighbourly letter can elicit significant inhibitory priming, then this will weaken the argument that the inhibition seen from the all-letter neighbourly primes in previous experiments is due to the combined effect of partial activation of four neighbours of the target word.

4.3.1 Method

Participants 113 participants were recruited at the University of Warwick and participated either in return for a payment of £3.00 or as a component of an introductory undergraduate research methods class. Of these, 22 were excluded for having an accuracy below 90%[3]. All participants began learning English by the age of 5 years.

Design The experiment consisted of a lexical decision task with masked form priming, in which we manipulated the type (identity, unrelated or neighbourly) and position (initial, final) of the one-letter primes within participants. Response times and accuracy were recorded.

Materials The targets were 180 four-letter words, and 180 four-letter nonwords. The target list was a different set of items to those used in Experiment 2.1, though there was some overlap. The initial and final letters of each target were consonants. Neighbourly primes were generated for each target by selecting a letter that could be substituted into the

[3]Changing this accuracy criterion to 80% did not change the pattern of results found for this experiment.

	Position			
	Initial	Final	Mean	Priming
Unrelated Prime	592 (2)	592 (3)	592 (2)	
Identity Prime	589 (2)	590 (2)	589 (2)	2 (0)
Neighbourly Prime	603 (3)	597 (2)	600 (3)	-9 (0)
Mean	594 (2)	593 (2)	594 (2)	

Table 4.2: Experiment 2.2 mean response times and error rates for responses to each prime type (Identity, Unrelated and Neighbourly) for each position (Initial and Final).

target in the initial or final position to form a new word, an orthographic neighbour of the target (e.g. `l+++` – BAND, `+++g` – BAND). Identity primes, unrelated primes and practice items were generated in the same way as for Experiment 2.1. The items were arranged into six lists, so that each target appeared with each of its primes only once.

Data from 33 word targets were rejected for having a mean response accuracy below 90%.

Procedure The procedure was identical to that used in Experiment 1.1.

4.3.2 Results

Mean response times and error rates are shown in Table 4.2 above. Incorrect trials and trials with response times smaller than 250ms or longer than 1500ms were excluded from the analysis[4]. A linear mixed-effects model was fitted with prime condition (identity, unrelated or neighbourly), item list and prime position (initial or final) as fixed factors. By-participant and by-item intercepts were added as random factors (no models with random slopes were able to converge). Type II Wald chi-square tests were used to establish the significance of the main effects and interactions. These revealed a significant main effect of prime con-

[4]26 of the correct trials (0.21%) were removed for having RTs that were too short or too long.

dition, $\chi^2(2) = 14.97$, $p < .001$. Post-hoc comparisons[5] showed that this was due to targets with neighbourly primes eliciting significantly longer response times than those with unrelated primes, $z = 2.94$, $p = .003$, and those with identity primes, $z = 3.75$, $p < .001$. Identity primes elicited slightly shorter response times than unrelated primes, but this difference was not significant, $z = 0.80$, $p = .421$. There was no significant main effect of item list or prime position, and no interactions were significant, $p > .05$.

An error rate analysis was performed by fitting a linear mixed-effects model with prime condition, item list and prime position as fixed factors. By-participant and by-item intercepts with slopes for prime condition were added as random factors (no models with random slopes were able to converge). No significant effects or interactions were found, $p > .05$.

4.3.3 Discussion

The latency analysis showed neighbourly inhibition, but no identity facilitation, whilst the error rate analysis showed identity facilitation, but no neighbourly inhibition. Taken together, it is clear that single letters were able to prime the target words.

Single letters caused significant inhibitory priming in this experiment. Given the size of effects found in Experiments 1.1 and 1.3, it is not feasible that this inhibition is caused by facilitatory activation of a lexical neighbour, which in turn causes inhibitory priming. That is, unless a single letter was causing more activation of a single lexical neighbour than the total activation caused by four neighbourly letters. It seems likely that this would happen as a result of priming, but it could be occurring as a result of perceptual blending.

The prime is presented immediately before target word, with no mask or white space between them. Efforts have been taken to prevent perceptual blending by presenting primes in a smaller font and different case to the targets, but it is still possible that due to the temporal adjacency of the prime and the target, participants are (at least occasionally) per-

[5]These comparisons are presented as z-tests instead of t-tests as I lacked the computational resources required for an appropriate calculation of degrees of freedom.

ceiving a target stimulus in which the prime overlaps, and thus simultaneously activating representations for the target and a lexical neighbour. Similar effects have been seen when participants attend to two words in different locations simultaneously; they occasionally report perceiving an illusory blend of the two words (e.g., Davis & Bowers, 2004).

4.4 Experiment 2.3

In order to rule out that the inhibitory one-letter neighbourly priming was due to perceptual blending of the prime and the target, thereby directly activating a neighbour of the target, we added another mask between the prime and the target, identical to the forward mask (######).

4.4.1 Method

Participants 105 participants were recruited at the University of Warwick and participated either in return for a payment of £3.00 or as a component of an introductory undergraduate research methods class. Only data from 95 participants who had begun learning English before the age of 5 years were considered for analysis. The mean accuracy was 93.8%. Of these, 15 participants were excluded for having an accuracy on word items below 90%[6].

Materials The same target items were used as in Experiment 2.2. Only final position primes were included in this experiment, as the priming effects found in Experiment 2.2 were marginally larger for final over initial letter primes. Data from 30 word targets were rejected for having a mean response accuracy below 90%.

Design and Procedure The design and procedure were identical to that used in Experiment 1.1, except that a 33ms duration backward mask (######) for the prime was inserted after the prime was presented,

[6]Changing this accuracy criterion to 80% did not change the pattern of results found for this experiment.

	RT (ms) and Errors (%)	Priming
Unrelated Prime	583 (3)	
Identity Prime	587 (3)	-4 (0)
Neighbourly Prime	590 (3)	-7 (0)
Mean	586 (3)	

Table 4.3: Experiment 2.3 mean response times and error rates for responses to each prime type (Identity, Unrelated and Neighbourly).

before the target, and that the primed letters only appeared in the final position.

4.4.2 Results

Mean response times and error rates are shown in Table 4.3 above. Incorrect trials and trials with response times smaller than 250ms or longer than 1500ms were excluded from the analysis[7]. For the latency analysis, a linear mixed-effects model was fitted with prime condition (identity, unrelated or neighbourly) and item list as fixed factors. By-participant and by-item intercepts with slopes for prime condition were added as random factors (no models with random slopes were able to converge). Type II Wald chi-square tests were used to establish the significance of the main effects and interactions. The main effect of prime condition approached significance, $\chi^2(2) = 5.47$, $p = .065$, and a planned comparison[8] revealed that the 7ms inhibitory effect of neighbourly priming as compared with unrelated primes was significant, $z = 2.34$, $p = .019$, whilst the other prime condition comparisons were not significant, $p > .05$.

There was no main effect of item list, $\chi^2(2) = 3.70$, $p = .157$, and no interaction between item list and prime condition, $\chi^2(4) = 4.91$, $p = .297$.

An error rate analysis was performed by fitting a linear mixed-effects model with prime condition, item list and prime position as fixed factors.

[7] 11 of the correct trials (0.09%) were removed for having RTs that were too short or too long.

[8] These comparisons are presented as z-tests instead of t-tests as I lacked the computational resources required for an appropriate calculation of degrees of freedom.

By-participant and by-item intercepts with slopes for prime condition were added as random factors (no models with random slopes were able to converge). There was no main effect of prime type, $\chi^2(2) = 2.03$, $p = .363$, or item list, $\chi^2(2) = 1.96$, $p = .375$, and no interaction was found, $\chi^2(4) = 3.93$, $p = .416$.

4.4.3 Discussion

Introducing a mask between the prime and the target appears to have prevented a significant main effect of prime type, along with the significant identity effects that we would expect to see. Had we observed identity priming without neighbourly priming, then this would have suggested that the neighbourly priming previously observed was due to perceptual blending. There was evidence for neighbour inhibition, but the lack of a main effect suggests that visual information from the prime could not always be accessed by our participants.

A possible explanation is that a 500ms ##### mask followed by a 52ms prime followed by a 33ms ##### was perceptually indistinguishable from a single mask without the intervening prime. In a follow-up experiment, we changed the second mask to use a different character in order to make them visually dissimilar.

4.5 Experiment 2.4

In order to rule out a lack of priming due to the prime being made impossible to perceive due to being preceded and followed by an identical mask, in our next experiment we changed the character that comprises the backward mask for the prime. We also changed the targets used in the experiment, to prevent repetition effects if any participants had previously completed a prior experiment due to overlap in participant pools. The percentage symbol (%) was selected as it provides good coverage for the full width and height of a character space, but shares few visual features with the hash symbol (#).

4.5.1 Method

Participants 124 participants were recruited at the University of Warwick and participated either in return for a payment of £3.00 or as a component of an introductory undergraduate research methods class. Only data from 111 participants who had begun learning English before the age of 5 years were considered for analysis. The mean accuracy was 87.3%[9]. Of these, 55 participants were excluded for having an accuracy on word items below 90%[10].

Materials The targets were 180 four-letter words, and 180 four-letter nonwords. None of the targets had been used in a previous experiment. Unlike the stimuli generated for Experiment 2.1, targets with repeated letters and vowels in the final and/or initial position were used. Identity primes, unrelated primes and neighbourly primes were generated in the same way as for Experiment 2.2. The items were arranged into six lists, so that each target appeared with each of its primes only once across the six lists, and each participant saw each target only once.

Data from 80 word targets were rejected for having a mean response accuracy below 90%.

Design and Procedure The design was identical to that used in Experiment 1.1. The procedure was identical to Experiment 2.3 with the following exceptions. The backward mask for the prime was changed to %%%%%%[11]. Primes could appear in either the initial or final position.

All stimuli were presented on an iiyama ProLite B2480HS monitor at 1920 x 1080 resolution, with a refresh rate set to 60Hz. The stimulus presentation and data collection were achieved through the use an executable compiled using 'BlitzMax' (Henderson, 2013).

[9]Accuracy was unusually low in this experiment, with no apparent explanation. This led to a larger-than-usual proportion of the data being excluded from analysis. Lowering the cut-off points for participant and item inclusion to 80% did not change the pattern of results.

[10]Changing this accuracy criterion to 80% did not change the pattern of results found for this experiment.

[11]The backward mask is printed here in the same font used during the experiment.

	Position			
	Initial	Final	Mean	Priming
Unrelated Prime	647 (3)	648 (3)	647 (3)	
Identity Prime	646 (1)	641 (3)	644 (2)	4 (1)
Neighbourly Prime	642 (2)	641 (3)	641 (2)	6 (0)
Mean	645 (2)	644 (3)	644 (2)	

Table 4.4: Experiment 2.4 mean response times and error rates for responses to each prime type (Identity, Unrelated and Neighbourly) for each position (Initial and Final).

4.5.2 Results

Mean response times and error rates are shown in Table 4.4 above. Incorrect trials and trials with response times smaller than 250ms or longer than 1500ms were excluded from the analysis[12]. For the latency analysis, a linear mixed-effects model was fitted with prime condition (identity, unrelated or neighbourly) and item list as fixed factors. A by-participant intercepts and a by-item intercept with position as a slope were added as random factors. Type II Wald chi-square tests were used to establish the significance of the main effects and interactions. There was no main effect of prime condition, $\chi^2(2) = 1.38$, $p = .501$, no main effect of item list, $\chi^2(5) = 4.54$, $p = .475$, and no main effect of primed letter position, $\chi^2(1) = 0.83$, $p = .364$. None of the interactions was significant, $p >.05$.

An error rate analysis was performed by fitting a linear mixed-effects model with prime condition, item list and prime position as fixed factors. By-participant and by-item intercepts with slopes for prime condition were added as random factors (no models with random slopes were able to converge). This also revealed no main effect of prime condition, $\chi^2(2) = 3.33$, $p = .189$, no main effect of item list, $\chi^2(5) = 6.12$, $p = .295$, and no main effect of primed letter position, $\chi^2(1) = 2.13$, $p = .145$. None of the interactions was significant, $p > .05$.

[12] 9 of the correct trials (0.17%) were removed for having RTs that were too short or too long.

4.5.3 Discussion

Not only did we fail to find evidence for neighbourly priming with a backward mask that differed from the forward mask, but there was still no evidence for identity priming taking place. It could be that the backward mask reduces the magnitude of priming effects across the board, although the small differences observed in our results were not always in the expected direction; neighbourly primes had shorter RTs than unrelated primes.

4.6 Experiment 2.5

In order to attempt to elicit significant priming effects with a backward mask, we attempted to increase the power of the experiment by increasing the number of trials. To do this it was necessary to use five-letter words instead of four-letter words as this provided a larger pool of potential targets with orthographic neighbours diverging in both the initial and final position.

4.6.1 Method

Participants 60 participants were recruited at the University of Warwick and participated for a payment of £3.00. The mean accuracy was 92.2%. All participants were native English speakers. Data from 10 participants was excluded from the analysis as they had an accuracy below 90%.

Materials The targets were 200 five-letter words, and 200 five-letter nonwords. Each target had either 1 or 2 orthographic neighbours, with one of those neighbours diverging from the target in either the initial or final letter position. Each target was categorised as Initial Letter or Final Letter based upon the location of this diverging letter, with half of the targets in each category. Initial Letter primes (e.g. `e++++`) were generated for all Initial Letter targets, and Final Letter primes (e.g. `++++e`) for all Final Letter targets. Each target had three primes generated;

	Position			
	Initial	Final	Mean	Priming
Unrelated Prime	622 (2)	613 (2)	618 (2)	
Identity Prime	619 (2)	619 (2)	619 (2)	-1 (0)
Neighbourly Prime	617 (3)	619 (2)	618 (2)	0 (0)
Mean	619 (2)	617 (2)	618 (2)	

Table 4.5: Experiment 2.5 mean response times and error rates for responses to each prime type (Identity, Unrelated and Neighbourly) for each position (Initial and Final).

identity primes (the lowercase letter of the target in the first/final position), neighbourly primes (the lowercase letter from the orthographic neighbour that diverges in that position), and unrelated primes (a random letter selected from a list of all letters that did not occur in either the target or any of its orthographic neighbours). The items were arranged into three lists, so that each target appeared with each of its primes only once across the three lists, and each participant saw each target only once.

Response accuracy to 55 of the word items was below 90%, so data for these items were excluded from the final analysis.

Design and Procedure The design was identical to that used in Experiment 1.1 The procedure was identical to that of Experiment 2.4.

4.6.2 Results

Mean response times and error rates are shown in Table 4.5 above. Incorrect trials and trials with response times smaller than 250ms or longer than 1500ms were excluded from the analysis[13]. For the latency analysis, a linear mixed-effects model was fitted with prime condition (identity, unrelated or neighbourly), prime position (initial or final), and item list as fixed factors. By-participant and by-item intercepts were added as

[13]13 of the correct trials (0.18%) were removed for having RTs that were too short or too long.

random factors (no models with random slopes were able to converge). Type II Wald chi-square tests were used to establish the significance of the main effects and interactions. There was no main effect of prime condition, $\chi^2(2) = 0.34$, $p = .843$, no main effect of item list, $\chi^2(2) = 2.56$, $p = .278$, and no main effect of primed letter position, $\chi^2(1) = 0.19$, $p = .664$. None of the interactions was significant, $p > .05$.

An error rate analysis was performed by fitting a linear mixed-effects model with prime condition, item list and prime position as fixed factors. By-participant and by-item intercepts were added as random factors (no models with random slopes were able to converge). This also revealed no main effect of prime condition, $\chi^2(2) = 1.78$, $p = .411$, no main effect of item list, $\chi^2(2) = 0.81$, $p = .667$, and no main effect of primed letter position, $\chi^2(1) = 0.41$, $p = .522$. None of the interactions was significant, $p > .05$[14].

4.6.3 Discussion

Once again, we failed to find any priming for this paradigm. This suggests that introducing a mask between the prime and the target either prevented the prime from being processed, or any activation from a single-letter prime decays sufficiently over 33ms so as to not affect target processing speed or accuracy.

Priming effects across a backward mask are possible. Grainger and Frenck-Mestre (1998) have demonstrated priming effects with a 14ms backward mask. Duyck and Warlop (2009) had significant priming from 56ms primes with a 56ms backward mask (with whole word primes in a different language to the target). Perhaps more relevant to the present experiment which used single-letter primes, Jacobs and Grainger (1991) elicited significant priming in an alphabetic decision task with 20ms backward masks. Perhaps a 33ms is too long a duration for backward masking a single letter.

[14] Changing the partition accuracy criterion to 80% revealed a significant interaction between prime type and position in the accuracy analysis (p = .035), due to a significantly higher error rate for neighbourly primes as compared with unrelated primes for primes in the final letter position only. The remaining pattern of results was unchanged.

4.7 General discussion

These experiments in single-letter primes have revealed that it is possible to prime a four-letter word target using a single-letter prime with positional information intact. It is unclear as to whether this is due to the activation of representations connected with the representation of the target, as was the intent, or due to perceptual blending. The introduction of a backward mask in order to rule out the latter, however, appears to have stopped all priming. This may mean that we can never resolve whether letter-level-to-word-level inhibition exists as without a backward mask there will always be the possibility that the temporal proximity of prime and target leads to perceptual blending.

It could be that there is a particular combination of prime-length and backward-mask length that would allow priming effects whilst ruling out the perceptual blending explanation. Perhaps a longer prime would lead to larger activation levels that their effects could persist beyond a 33ms mask. This is investigated further in Experiment 3 in Chapter 5.

There have been previous examples of primes that were initially thought to not be capable of eliciting priming effects that were later demonstrated using a redesigned experiment. Guerrera and Forster (2008) found no priming effects for primes with every letter transposed, whilst Lupker and Davis (2009) introduced a sandwich priming paradigm, which showed significant priming for these all-transposed primes (as well as primes with 3 or more letter replacements) when a target identity preprime with 33ms duration was introduced. Perhaps an identity preprime would provide a boost in activation of the target node that would allow a priming effect that had been reduced by the presence of a backward mask to manifest successfully. This option is explored further in Experiment 4 in Chapter 6.

Chapter 5

Overcoming Backward Masking

To address questions raised by previous experiments, Experiment 3 sought to investigate the display duration required for a single letter prime to affect reaction times to a target after a 33ms backward mask.

5.1 Experiment 3

Experiments 2.1 and 2.2 showed that response times to a word target could be significantly affected by a one-letter prime (such as s+++). No significant priming was found in experiments 2.3, 2.4 and 2.5. The consistent difference between these experiments is that a backward mask was introduced between the prime and target from Experiment 2.3. Perhaps the failure to find significant results in these experiments is because any priming caused by one-letter prime cannot survive a 33ms mask.

To investigate this, I designed an Alphabetic Decision Task (ADT), in which participants need to classify single-character targets as either belonging to the Latin (English) Alphabet, or to another source. By varying the prime duration and the presence of a backward mask, I hope to demonstrate that the activation of a single letter requires a display duration longer than 50ms in order to survive across a backward mask.

A previous ADT was conducted by Grainger and Jacobs (1991) in which the primes were words and the targets were letters or non-letters embedded within a string of hash marks (e.g., table-T####). The non-letters foils used were common punctuation symbols: £, %, (, *, +, ?, =,), <, and >. I felt that these are inappropriate foils for an ADT. For one, they are visually dissimilar from Latin alphabet letters. They are also already known to English speakers. To make my ADT task closer to the LDT, in which the foils are unfamiliar to participants and are not able to be disambiguated from the targets based on their features alone, I used characters from Cherokee, Armenian, Cyrillic and Greek scripts selected for their visual similarity in terms of features to uppercase letters from the Latin alphabet used in English.

5.1.1 Method

Participants 82 participants were recruited at the University of Warwick and participated for a payment of £3.00. All participants began learning English by the age of 5 years. The mean accuracy for all responses was 96.2%. One participant was removed from the analysis for having an accuracy less than 90%[1].

Design The experiment consisted of an alphabetic decision task, in which I manipulated the alphabet (Latin or non-Latin), letter identity (one of 10 letters), prime type (identity, unrelated), prime duration (10, 30, 50, 70, 90ms) and the presence of an intermediate mask between the prime and the target (present, not present). This resulted in 400 trials, with each participant seeing each trial once.

Materials The targets were 10 Latin alphabet letters (henceforth referred to as English Letters), and 10 non-English letters or symbols selected to resemble visual features present in English letters (henceforth referred to as NonEnglish Letters). The targets are outlined in Table 5.1. The English Letter targets were selected because their upper and lowercase forms are visuaully distinct. All targets were presented in uppercase. For each target, two primes were generated. An identity prime (the

[1]Changing this accuracy criterion to 80% did not change the pattern of results found for this experiment.

same letter identity as the target) and an unrelated prime (a different letter identity to the target). All English Letter targets were preceded by English Letter primes, and all NonEnglish Letter targets by NonEnglish Letter primes. All primes were presented in lowercase, where available[2].

In addition to these items, 10 practice items were added to the start of the experiment (5 English Letters and 5 NonEnglish Letters) with identity primes. These items were items that were not included in the main experiment, and were not included in the analysis.

Response accuracy to one English Letter target (*Q*) was under 90%, so data for this target were excluded from the analysis.

Procedure Participants were asked to decide whether each target letter-string, presented on a computer screen, was an English Letter or not. Responses were made by pressing one of two keyboard keys. Accuracy and RTs were recorded. Targets were presented in a randomised order for each participant.

Each trial started with a fixation cross (+) being presented for 300ms, followed by a blank screen for 200ms. Then a forward mask (`###`) was presented for 500ms, followed by the prime in 80pt font size for 10, 30, 50, 70 or 90ms. In trials where the intermediate mask was present, the prime was followed by a mask (`%%%`) for 33ms. The target was then presented at 60pt font until a response was given, or until 2000ms had elapsed, at which point a "Too slow!" message was displayed. The only other feedback participants received was "Wrong!" being displayed if the participant gave an incorrect response. All stimuli were presented in the FreeMono font (Peterlin & White, 2016).

All stimuli were presented on an iiyama ProLite B2480HS monitor at1920 x 1080 resolution, with a refresh rate set to 60Hz. The stimulus presentation and data collection were achieved through the use an executable compiled using 'BlitzMax' (Henderson, 2013).

[2]Some of the NonEnglish Letter targets were taken from alphabetic scripts, and so had a corresponding lowercase form. Others were non-alphabetic symbols, and so had no lowercase form. In these cases, the identity prime was identical in form to the target.

Alphabet	Letter	Description	Unicode Number	Identity Prime
English	B	Latin Capital Letter B	U+0042	b
	D	Latin Capital Letter D	U+0044	d
	F	Latin Capital Letter F	U+0046	f
	G	Latin Capital Letter G	U+0047	g
	H	Latin Capital Letter H	U+0048	h
	L	Latin Capital Letter L	U+004C	l
	N	Latin Capital Letter N	U+004E	n
	Q	Latin Capital Letter Q	U+0051	q
	R	Latin Capital Letter R	U+0052	r
	T	Latin Capital Letter T	U+0054	t
NonEnglish	ɮ	Latin Small Letter Lezh	U+026E	ɮ
	Δ	Greek Capital Letter Delta	U+0394	δ
	Ц	Cyrillic Capital Letter Tse	U+0426	ц
	Э	Cyrillic Capital Letter E	U+042D	э
	б	Cyrillic Small Letter Be	U+0431	б
	ծ	Armenian Small Letter Ca	U+056E	ծ
	Ꭾ	Cherokee Letter He	U+13AE	ꭾ
	Ꮌ	Cherokee Letter Mo	U+13BC	ꮌ
	Ꮘ	Cherokee Letter Qui	U+13C8	ꮘ
	Ꮨ	Cherokee Letter Ti	U+13D8	ꮨ

Table 5.1: A list of target letters used in Experiment 3.1, along with the descriptive name and Unicode number, as classified by The Unicode Standard 10.0 (Unicode, Inc., 2018). Identity Prime shows the letter used as the identity prime for the target (see footnote on the previous page).

	Prime Duration	Identity Prime	Unrelated Prime	Priming
Backward Mask Absent	10	517 (3)	522 (4)	5 (1)
	30	511 (2)	519 (2)	8 (0)
	50	491 (3)	506 (3)	15 (0)
	70	489 (2)	499 (2)	10 (0)
	90	490 (2)	497 (2)	7 (0)
	Mean	500 (2)	509 (3)	9 (0)
Backward Mask Present	10	509 (2)	514 (2)	5 (0)
	30	514 (2)	509 (3)	-5 (1)
	50	506 (2)	501 (2)	-5 (0)
	70	494 (3)	507 (2)	13 (-1)
	90	492 (2)	509 (4)	17 (2)
	Mean	503 (2)	508 (3)	5 (0)
	Mean	501 (2)	508 (3)	7 (0)

Table 5.2: Experiment 3.1 mean response times and error rates for responses to each prime type (Identity and Unrelated) for each duration (10, 30, 50, 70 and 90ms) with a backward mask present or absent.

5.1.2 Results

Data were analysed using linear mixed-effects models in R version 3.4.0 (R Core Team, 2020), using the packages *lme4* (Bates et al., 2015) and *car* (Fox & Weisberg, 2019), with significant factors being further investigated using the *lsmeans* package (Lenth, 2016). Models were fitted with full random structure, and simplified if they could not converge.

Mean response times and error rates are shown in Table 5.2 above. Incorrect trials and trials with response times longer than 1500ms and shorter than 250ms were excluded from the latency analysis[3]. A linear mixed-effects model was fitted with prime type, mask presence and prime duration and their interactions as fixed factors. By-participant and by-target intercepts were added as random factors. Type II Wald chi-square

[3]14 of the correct trials (0.10%) were removed for having RTs that were too short or too long.

tests were used to establish the significance of main effects and interactions. This revealed a significant main effect of prime type, $\chi^2(1) = 17.95$, $p < .001$, as participants responded faster to identity-primed items than to those with unrelated primes. A significant main effect of duration was found, $\chi^2(4) = 95.12$, $p < .001$. Post-hoc comparisons revealed that this was caused response times to items with 10 and 30ms primes being longer than to items with 50, 70 or 90ms primes. There was no significant main effect of mask presence, $\chi^2(1) = 0.68$, $p = .408$. There was a significant interaction between mask presence and duration, $\chi^2(4) = 18.13$, $p = .001$; this was due to 10ms primes eliciting faster response times when masked, $t(14115) = 2.26$, $p = .012$. There was no such effect for other prime durations, $p > .05$. The were no significant interactions between prime type and mask presence, $\chi^2(1) = 1.41$, $p = .235$, or prime type and prime duration, $\chi^2(4) = 6.09$, $p = .192$.

Finally, there was a significant three-way interaction between prime type, mask presence and prime duration, $\chi^2(4) = 11.27$, $p = .024$. Investigating only the trials where the backward mask was absent, there was a significant effect of prime type, $\chi^2(1) = 15.08$, $p = .259$ due to the facilitatory effect of identity primes. There was a significant effect of prime duration, $\chi^2(4) = 93.01$, $p < .001$, as response times to items with 10 and 30ms primes were longer than to items with 50, 70 and 90ms primes. There was no significant interaction. Turning to trials with the backward mask present, there was significant facilitatory identity priming effect, $\chi^2(1) = 4.54$, $p = .033$ and a significant effect of prime duration, $\chi^2(4) = 17.97$, $p = .001$, as response times to items with 10 and 30ms primes were longer than to items with 50, 70 and 90ms primes. Finally, there was a significant interaction between prime type and prime duration, $\chi^2(4) = 14.57$, $p = .006$. This was due to significant identity priming only being evidenced with 90ms primes, $t(7010) = 3.31$, $p < .001$.

A planned comparison was conducted to investigate the effects of prime type and mask presence for items with a prime duration of 50ms, as this was the prime duration used for Experiments 2.1 - 2.5. A linear mixed-effects model was fitted with prime type and mask presence and their interactions as fixed factors. By-participant and by-target intercepts and slopes for prime type were added as random factors. Type II Wald chi-square tests were used to establish the significance of main effects and interactions. There was no significant effect of prime type, $\chi^2(1) = 2.01$, $p = .156$; or mask presence, $\chi^2(1) = 1.53$, $p = .216$; but the interaction between the two was significant, $\chi^2(1) = 8.01$, $p = .005$. When there was no backward mask present, there was significant facilitatory identity priming, $t(31.9) = 2.96$, $p = .006$, whereas when there was a mask present, there was no difference between response times to identity and unrelated primes, $t(31.3) = 0.92$, $p = .437$.

An error rate analysis was conducted by fitting a linear mixed-effects model with prime type, mask presence and prime duration and their interactions as fixed factors. By-participant and by-target intercepts with slopes for prime type were added as random factors. Type II Wald chi-square tests were used to establish the significance of main effects and interactions. There was no significant effect of prime type, $\chi^2(1) = 0.19$, $p = .660$, mask presence, $\chi^2(1) = 0.00$, $p = .958$, or prime duration, $\chi^2(4) = 3.36$, $p = .499$. There was no significant prime type by mask presence interaction, $\chi^2(1) = 0.00$, $p = .958$, nor for prime type by duration, $\chi^2(1) = 3.42$, $p = .490$. There was a significant interaction between mask presence and prime duration, $\chi^2(4) = 13.34$, $p = .010$. This was due to a significant effect of prime duration only for unmasked primes, $\chi^2(4) = 12.76$, $p = .013$, in which 10ms primes evoked higher error rates than 30ms, $t(7200) = 2.90$, $p = .004$, and 90ms primes, $t(7200) = 3.01$, $p = .003$, whilst this effect was not present for masked primes, $\chi^2(4) = 3.83$, $p = .430$. The three-way interaction between prime type, mask presence and prime duration was not significant, $\chi^2(4) = 3.26$, $p = .516$.

5.1.3 Discussion

This experiment showed evidence for identity priming in an ADT, as well as demonstrating that for 50ms prime durations, introducing a 33ms backward mask will indeed prevent identity priming. In order for single letter primes to survive a backward mask of this duration, the minimum prime duration appears to be greater than 70ms (here being successfully demonstrated with 90ms primes). Anecdotally, primes with a duration of 90ms can be visible, and at least carry an increased risk of being consciously recognised by participants, even with masking, which might make them unsuitable for use in alphabetic and lexical decision tasks.

The other results of interest from this experiment relate to effects of mask presence and prime duration irrespective of prime type. 10ms primes led to faster RTs when a mask was introduced, and responses to primes of 10 and 30ms were slower on the whole than to primes of 50ms and above. My belief is that this represents an effect of the total duration between the trial start and the target presentation. For shorter, non-backward-masked primes, the participants had less time to prepare to respond compared with longer, backward-masked primes. This led to response delays to trials with a shorter start-to-target duration, particularly in an experiment where this duration varied significantly between trials.

These results help to explain why priming effects were no longer elicited when a backward mask was introduced in Experiments 2.3, 2.4 and 2.5. 50ms single letter primes can only elicit identity priming when unmasked. It is also not possible that the priming demonstrated in this experiment was due to perceptual blending. This must be a lexical effect due to abstract letter representations, as the letter targets used had different lowercase and uppercase forms. This does not rule out the perceptual blending possibility for the results demonstrating neighbourly inhibition from single letter primes, but it does indicate that single letters can elicit priming in situations were blending cannot be blamed.

Chapter 6

Sandwich Priming

6.1 Experiment 4

The sandwich priming paradigm was developed by Lupker and Davis (2009) in order to overcome lexical competitor effects in masked priming. It does so by raising the activation level of the target through use of an identity preprime.

This experiment seeks to explore whether this paradigm can help to overcome the issues posed by introducing a backward mask to Experiments 2.3 - 2.5. A higher starting activation level in the target at the point of one-letter prime presentation may allow the manifestation of smaller changes in activation levels that can arise in a standard lexical decision task using a backward mask. I have kept the 33ms backward mask between the prime and the target, and also included one between the preprime and the prime, in order to prevent the perception of a blended preprime-prime, which, if perceived, would cause the paradigm to act as a standard lexical decision with an orthographic neighbour prime.

6.1.1 Method

Participants 136 participants were recruited as undergraduate students at the University of Warwick and participated as a component of an introductory methods class. 118 of these participants began learning English by the age of 5 years, so only data from these participants was considered for analysis. The mean accuracy across these participants was 90.4%. Data from 41 of these participants were excluded from the analysis for having a word accuracy below 90%.

Materials The items used were identical to those used in Experiment 2.1. The mean accuracy of responses to 62 of the 200 word items was less than 90%, so data from these items were not used in the final analysis.

Design and Procedure The experiment used a masked sandwich priming paradigm. Each trial started with a fixation cross (+) being presented for 500ms, followed by a blank screen for 200ms. Then a forward mask (#####) was presented for 500ms, followed by the lowercase preprime in 50pt font size, displayed for 33ms. A backward mask (%%%%%) was presented after the preprime for 33ms. This was followed by the lowercase prime, which was displayed for at 60pt font size for 70ms. This was masked (#####) for 33ms. Finally, the uppercase target was presented at 50pt font until a response was given, or until 2000ms had elapsed, at which point a "Too slow!" message was displayed. The only other feedback participants received was "Wrong!" being displayed if the participant gave an incorrect response.

All stimuli were presented on an iiyama ProLite B2480HS monitor at 1920 x 1080 resolution, with a refresh rate set to 60Hz. The stimulus presentation and data collection were achieved through the use an executable compiled using 'BlitzMax' (Henderson, 2013).

6.1.2 Results

Data were analysed using linear mixed-effects models in R version 4.0.3 (R Core Team, 2020), using the packages *lme4* (Bates et al., 2015) and *car* (Fox & Weisberg, 2019), with significant factors being further investigated using the *emmeans* package (Lenth, 2020). Models were fitted with full random structure, and simplified if they could not converge.

Mean response times and error rates are shown in Table 6.1 on the next page. Incorrect trials and trials with response times smaller than

	Position			
	Initial	Final	Mean	Priming
Unrelated Prime	653 (2)	641 (2)	647 (2)	
Identity Prime	639 (2)	640 (1)	639 (2)	7 (0)
Neighbourly Prime	643 (2)	641 (2)	642 (2)	5 (0)
Mean	645 (2)	640 (2)	643 (2)	

Table 6.1: Experiment 3.2 mean response times and error rates for responses to each prime type (Identity, Unrelated and Neighbourly) for each position (Initial and Final).

250ms or longer than 1500ms were excluded from the analysis[1]. For the latency analysis, a linear mixed-effects model was fitted with prime condition (identity, unrelated or neighbourly), prime position (initial or final) and item list as fixed factors. By-participant and by-item intercepts were added as random factors (no models with random slopes were able to converge). Type II Wald chi-square tests were used to establish the significance of the main effects and interactions. These revealed a significant main effect of prime condition, $\chi^2(2) = 7.08$, $p = .029$[2]. Post-hoc tests revealed that identity primes elicited shorter response times than unrelated primes, $t(10179) = 2.49$, p =.013. Neighbourly primes also elicited shorter response times than unrelated primes, but this did not reach significance $t(10179) = 1.74$, $p = .081$. Response times for targets with identity and neighbourly primes did not significantly differ, $t(10179) = 0.74$, $p = .456$.

There was no significant main effect of item list, $\chi^2(2) = 3.49$, $p = .174$, or prime position, $\chi^2(1) = 0.64$, $p = .422$. The interaction between prime condition and prime position approached significance, $\chi^2(2) = 5.82$, $p = .054$[3]. Post-hoc tests revealed this to be to due to a significant effect of prime condition occurring with the initial position primes, $\chi^2(2) = 12.83$, $p = .002$, but not the final position primes, $\chi^2(2)$

[1] 36 of the correct trials (0.35%) were removed for having RTs that were too short or too long.

[2] This main effect was no longer significant (p = .094) when the participant accuracy criterion was changed to 80%

[3] This interaction was significant with a participant accuracy criterion of 80%

= 0.09, p = .955. For initial position primes, identity primes elicited faster responses than unrelated primes, $t(5123) = 3.43$, $p < .001$, as did neighbourly primes, $t(5123) = 2.61$, $p = .009$, which did not differ from identity primes, $t(5123) = 0.82$, $p = .411$ The interaction between item list and prime condition was not significant, $\chi^2(4) = 1.06$, $p = .900$, nor was the interaction between item list and prime position, $\chi^2(2) = 1.53$, p = .466. The three-way interaction was also not significant, $\chi^2(4) = 3.09$, $p = .542$.

An error rate analysis was performed by fitting a linear mixed-effects model with prime condition, item list and prime position as fixed factors. By-participant and by-item intercepts were added as random factors (no models with random slopes were able to converge). This also revealed no main effect of prime condition, $\chi^2(2) = 1.52$, $p = .467$, no main effect of item list, $\chi^2(2) = 1.09$, $p = .581$, and no main effect of primed letter position, $\chi^2(1) = 0.65$, $p = .421$. None of the interactions was significant, $p > .05$.

6.1.3 Discussion

This backward-masked sandwich priming experiment with single letter primes showed a significant identity priming effect, which was predicted. There was also some evidence for a facilitatory neighbourly priming effect, which is the opposite of the hypothesised neighbourly inhibition. The results were in this direction overall, with a significant effect for initial letter primes. Neither the main prime type effect nor the overall interaction between prime type and position reached significance so, whilst the implications for a facilitatory neighbourly effect will be discussed, it should be remembered that this may not reflect a reproducible effect.

An explanation for identity letter facilitation in this sandwich priming experiment is no different from the explanation in standard masked priming experiments; activation of the letter feeds forward to the representation of the target word, giving it a performance boost relative to an unrelated letter.

I hypothesised that neighbourly priming would show significant inhibition in a sandwich priming paradigm, as the identity preprime would raise the activation levels of the target and so make small inhibitory effects easier to detect. The original sandwich priming experiments by Lupker and Davis (2009) found facilitatory priming for all-letter-transposed (T-All) primes, which had previously shown no facilitation. Lupker and Davis attributed this to the boosted activation from the preprime helping to overcome lexical competitor effects. If the T-All primes were activating neighbours of their targets more strongly than the targets themselves, then the lexical competition from those neighbours would suppress any facilitation of the target. However, if the target were already activated from a preprime, then the target suppresses neighbour activation and prevents them from causing inhibition.

This account would at least explain the lack of inhibitory priming from neighbourly letter primes - target activation from the preprime suppresses the neighbour word representations preventing them from showing inhibitory effects. This does not explain neighbourly facilitation, however.

When the target preprime is presented, it will activate the target representation, but this preprime is also a partial match for neighbours of the target. The design of my stimuli ensures that each word target has at least one neighbour with which it shares 3 of its 4 letters. This neighbour should receive some activation from the preprime, though models implementing lexical competition would predict that this activation will be suppressed by the activation of the target word. Non-homogeneous letter-word inhibition could also account for suppression directly from the letter level, from the letter in the preprime that differs from the letters in the neighbour. When the next stimulus, the one letter prime, is a neighbourly letter, this neighbour of the target will receive more activation. Models such as the IA implement top-down feedback (McClelland & Rumelhart, 1981), so this neighbour activation would then feed activation back down to the letter layer, raising the activation of its letter, 3 of which are shared with the target, which would in turn contribute to the target representation's already raised activation from the preprime. This would account for the facilitation seen relative to unrelated letter primes, which do not activate neighbour words with substantial letter overlap with the target.

The difficulty with this explanation is accepting that the feedback from the neighbour word representation to the letter level would be substantial enough to significantly boost activation of the target if lexical competition (or non-homogeneous letter-word inhibition) from the target's preprime activation is suppressing its neighbours. Perhaps this suppression is sufficient to prevent inhibitory activation from the neighbour or neighbourly letters, but not so strong as to completely suppress any facilitatory activation.

In research conducted since this I carried out this experiment, Trifonova and Adelman (2018) found evidence that, contrary to the claims of Lupker and Davis (2009), the sandwich priming paradigm does not eliminate lexical competition effects. They found facilitatory form and T-All priming for both identity preprimes and neighbour preprimes (e.g. around - bauodn - ABOUND). This means that the reason that facilitatory effects were found for T-All primes in the sandwich priming experiment was not because the identity preprime removed lexical competition; a neighbour preprime, which would be expected to increase lexical competition, also showed facilitation. Instead, it appears that sandwich priming effects are a prelexical effect affected by the joint similarity between the prime and the target.

This account may help to explain the results from my experiment. Whilst the unrelated and neighbourly letter primes are equally dissimilar to the target, it could be that a prelexical combination of the preprime and the prime leads to identity and neighbourly primes contribute towards a 'Yes' response in the LDT moreso than an unrelated prime.

A possibility that I have not yet discussed is that the identity and neighbourly primes are not facilitatory, but that the unrelated primes alone are inhibitory. Even in the unrelated prime condition, an identity preprime is used. It could be the case that representations activated by an unrelated prime interfere with lexical selection in some way. Any such interference would not be offset by the feedback activation of overlapping letters as it could be for the neighbourly letters, as I've outlined above.

The previous experiments that found significant neighbourly priming showed an inhibitory effect of neighbourly priming, and previous experiments from other researchers either found no priming effect for orthographic neighbours in sandwich priming (Burt & Duncum, 2017) or an inhibitory effect (for primes of a higher frequency than the target; Nakayama et al., 2008). I will discuss how both results might be understood together in my integrated discussion in Chapter 8.

The results indicated that this neighbourly priming effect was confined to initial letter primes, but not final letter primes, which further raises questions around non-equivalence of these exterior letters. It could be that the relationships between initial letter nodes and their word representations are different to those of interior and/or final letters. Forster and Taft (1994) argued that skilled readers might recode orthographic representations as head and rhyme body (e.g. H-ERD). This would mean that a final letter prime primes part of a unit, but the inital letter primes prime the entire unit.

This could also be a purely perception-based phenomenon. Whilst the fixation cross provided at the start of each trial was centrally aligned relative to the target word, perhaps the preprime draws fixations to the Preferred Viewing Location slightly to the left of the centre (e.g. Rayner, 1979). It should be noted once again, however, that the interaction between prime type and letter position only approached significance, so this is unlikely to be a reliable effect, especially as my previous experiments that manipulated letter position revealed no effect.

Chapter 7

Simulations

7.1 Introduction

In order to directly test that the neighbourly inhibition mechanisms described in this thesis would produce a viable model of lexical access and lexical decision that would produce the key masked priming effect, I ran a number of computational model simulations.

7.2 Method

Simulations were run using a modified version of the IA model (McClelland & Rumelhart, 1981), in which the connections between the letter and word representations had non-homogeneous connection weights. Whilst some amount of trial and error was necessary to create a model that successfully captures the hypotheses of this thesis, only one model simulation will be described in this chapter. Apart from small changes made to the model parameters in order to produce a model that could discriminate between words and nonwords, no search of the parameter space was conducted in order to find a model that best fit human data. The purpose of these simulations was to demonstrate what is possible with small changes to the IA model, not to demonstrate a simulation with cherry-picked parameters that performs well on the trial items described below, but would require significant changes to the parameters in order to generalise well on other trial items.

To compare this Neighbourly IA model, simulations were also conducted using a non-modified IA with analogous parameter changes to parameters common to both versions of the model

7.2.1 Software

Simulations were conducted using the *easyNet* simulation software developed by Adelman et al. (2017).

7.2.2 Model Details

The Neighbourly IA model was a version of the IA model, modified from the model as described by Davis (2003). The main modification was that connections from the letter layer to the word layer had their connection strengths adjusted based on the presence of neighbours in the model's vocabulary. The model's vocabulary was all 4-letter words given frequency norms by Kučera and Francis (1967), with resting activation levels set according to the log frequency of the word, as generated using N-Watch (Davis, 2005); the same method used to set resting activation levels as Davis (2003). Connection weights from letter to word representations were set to +0.07 if that letter appears in that position in that word (e.g. C(1) →CAST), to -0.5 if the letter appears in an orthographic neighbour of the word in that position (e.g. F(1) →CAST, with the neighbour being FAST), and to -0.01 if the letter does not appear in that position in that word (e.g. D(1) →CAST).

As described by Davis (2003), the rate of the model, along with other time-based parameters, was reduced by a factor of 10 to produce response times more akin to the millisecond range of response times captured with human participants.

In order to simulate lexical decision, the Multiple Read-Out Model (MROM; Grainger & Jacobs, 1996) module was added to the model as made available by the *easyNet* software. This implements a three-decision criteria, whereby as "Yes" response can be reached if any word unit reaches a set activation level threshold, or if the sum of all word unit activation (the total lexical activity) reaches a different threshold, and a "No" response can be reached if a "Yes" decision is not made after a certain number of cycles. The module referred to as MROM in *easyNet* actually implements the three-decision mechanism used by Coltheart et al. (2001) in the Dual Route Cascaded (DRC) model, which uses the Grainger and Jacobs (1996) criteria with the following modification; the total lexical activity "Yes" threshold and the "No" deadline value can be updated during processing. The DRC reduces the total lexical activity threshold and/or increases the "No" deadline value at cycle 20 depending upon the total lexical activity at that point.

In the Neighbourly IA model, this was implemented as follows: if the total lexical activity reached 0.65 by cycle 70[1], the total lexical activity "Yes" threshold was reduced from 0.90 to 0.72, and if the total lexical activity reached 0.22 by cycle 120, the "No" deadline value was increased from 250 to 300.

The specific word threshold parameter (known as the *M* criterion in the MROM, Grainger and Jacobs, 1996; the *A* threshold in the DRC, Coltheart et al., 2001; and the 'unit decider' in *easyNet*, Adelman et al., 2017) was brought down to 0.4 after initial simulations revealed that the model accuracy was extremely poor, particularly on nonwords with identity primes. The lexical activity began increasing from the point that the prime was presented, and could quickly reach levels sufficient to elicit a "Yes" response based on total lexical activity within the same time period that unrelated- and neighbourly-primed targets would not yet have reached sufficient activity levels for a "Yes" response based upon individual lexical activity.

The full list of parameters used in *easyNet* can be found in Appendix B.

Simulations were also conducted with a Base IA model, using all of the same parameters. However, instead of implementing neighbourly letter to word connection weights, excitatory connections (when the letter is present in the word in that position) were set to +0.07, and inhibitory weights (when the letter is not present in the word in that position) were

[1]The cycle count in *easyNet* begins with presentation of the prime, so for a 50 cycle prime duration, cycle 70 corresponds to 20 cycles after the presentation of the target.

homogeneously set to -0.04.

7.2.3 Stimuli

The stimuli used in this simulation were the same as those used in Experiment 1.1. Neighbourly primes (consisting of letters from orthographic neighbours of the target word), unrelated primes and identity primes were presented to the model for each four-letter target word. Not all of the original stimuli were used in the simulations; 30 of the 288 words were not used in the simulation as they did not appear in the model's vocabulary, whilst 15 of the 288 nonwords were not used as they did appear in the model's vocabulary. The discrepancy is due to the original stimuli selection classifying words or nonwords using the CELEX Lexical Database (Baayen et al., 1996), whilst the model vocabulary was generated using words appearing in the Kučera-Francis frequency norms (Kučera & Francis, 1967).

7.2.4 Trials

To simulate each trial, the prime was presented for 50 cycles and then the target was presented until a lexical decision had been made. This took a maximum of 300 cycles, corresponding with the late timeout parameter. The response of the model was recorded, as was the number of cycles required to reach this decision.

7.3 Results

The model results are reported in terms of reaction time in cycles (although, as described above, the rate of the model was adjusted so that cycles can be taken as a proxy for milliseconds) and the accuracy rate.

7.3.1 Neighbourly IA

The accuracy rate to all words with unrelated primes was 89.58%, with identity primes was 91.67% and with neighbourly primes was 88.54%.

This was 91.67% for nonwords with unrelated primes, 72.92% for nonwords with identity primes, and 93.75% for nonwords with neighbourly primes.

Amongst the correct responses, the average response to words with unrelated primes was 190 cycles. There was clear evidence of 77 cycles of identity facilitation, with an average identity prime response time of 113 cycles. Neighbourly primes has a mean response time of 199, showing 9 cycles of inhibition relative to unrelated primes.

7.3.2 Base IA

The accuracy rate to all targets was 100%, except for nonwords with identity primes, which had an accuracy rate of 91.21%.

Amongst the correct responses, the average response to words with unrelated primes was 179 cycles. There was clear evidence of 78 cycles of identity facilitation, with an average identity prime response time of 101 cycles. Neighbourly primes has a mean response time of 179, exactly the same as for unrelated primes.

7.4 Discussion

Whilst both models demonstrated identity priming effects, the Neighbourly IA model showed a small, but consistent, inhibitory neighbourly priming effect. The Base IA model showed no effect at all for neighbourly primes relative to the unrelated primes.

These results confirm that the neighbourly mechanism proposed in this thesis can lead to a viable model capable of performing lexical decision, but also one that demonstrates inhibitory effects of neighbourly primes. The only difference between the two models was the implementation of this mechanism.

Whilst these results show that neighbourly inhibition can be elicited via non-homogeneous inhibitory connections between the letter and word layers, it does not prove that this is the only way that these results can arise. The IA model available with the *easyNet* software does not include the changes made by Davis and Lupker (2006) that implemented specific lateral inhibition for only lexical neighbours in the word layer. As this model has not yet been tested, it cannot be stated as fact that this other form of lexical inhibition would not predict neighbourly inhibition in primed lexical decision.

There are also weaknesses with the Neighbourly IA model described in this chapter. A 27% error rate for nonwords with identity primes is very high, and does not match human performance on similar tasks. It is conceivable that further adjustments to the model parameters would have improved the model accuracy. It is also conceivable, however, that human behaviour on a specific, individual task could be approximated with a wide variety of lexical access models through the tweaking of such parameters. Instead, I intended to restrict my parameter changes to ones that were necessary to reflect the modelling conducted by Davis (2003), and ones that the MROM (Grainger & Jacobs, 1996) specifically describe as being changeable, task-dependent parameters. Once a set of parameters that produced passable results across all conditions (accuracy rates of 70% or higher), I ceased tweaking.

An important consideration is that performing well on a lexical decision task is no guarantee that the model is accurately identifying the target word. It could be that a word other than the target reaches the identification threshold, leading to a "Yes" response, or that the decision is based upon total lexical activity instead. Whilst these possibilities are also true for human participants, the Neighbourly IA model could hardly be said to be a model of lexical access if it frequently misidentified words. This does not appear to be the case, however. The *easyNet* software does not include an option for primed lexical identification tasks, but the activity levels of the nodes in the model can be investigated for individual trials. Doing so indicated that the target word was the most active node in the word layer for the subset of trials I investigated, as demonstrated in Figure 7.4, which plots the activation levels for word representations for the neighbourly trial sobl-CURE. CURE achieves an activation level of 0.4, which is the threshold for a "Yes" response in this simulation.

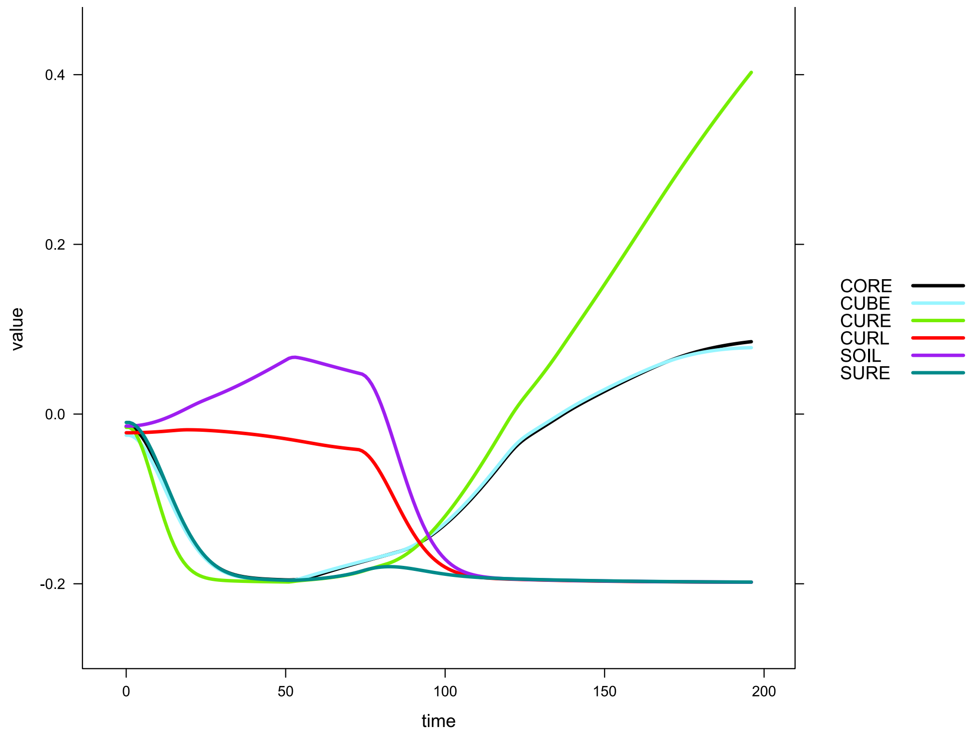

Figure 7.1: Word representation activation levels over time (cycles) when the Neighbourly IA Model was presented with the trial *sobl-CURE*.

Figure 7.4 includes the activation levels for the target, CURE, along with each of the neighbours used to construct its neighbourly prime (sobl; SURE, CORE, CUBE, CURL). Also included is SOIL, which is a neighbour of the prime. Each of these neighbours will receive some activation from the prime, as they each have one letter in common with it. However, as the plot shows, these primes respond in different ways to the presentation of the target in the 0 to 50 cycles time period. CURL virtually unaffected by the appearance of the prime and maintains a near constant activation. SURE, on the other hand, is also inhibited by the prime. This is because SURE has the neighbour SORE, so the model has an inhibitory weight between O(2) and SURE. For target words in dense orthographies, it is likely that a neighbourly prime will also contribute facilitatory activation to neighbours of the target's neighbours. This makes it harder for neighbourly inhibition effects to be explained with lexical inhibition. Even if we could accept that the partial activation of four neighbours of the target would lead to significant inhibition of the target word, there were often be trials in which these neighbours are themselves suppressed by its neighbourly prime, which would in turn prevent them from suppressing the target effectively.

In conclusion, these simulations show that a model based upon the IA model with a neighbourly mechanism of non-homogeneous letter-word inhibition does predict inhibitory effects in lexical decision RTs. Whilst there is clearly room for improvement of the Neighbourly IA model in terms of improving its accuracy, these simulations show that a small inhibitory effect from all-letter neighbourly primes can be predicted by a model implementing non-homogeneous letter-word inhibitory connections. If the effect really is as small as was found in this simulations (199 vs 190 cycles), then this could help to explain why small changes to the paradigm prevented the effect from appearing in behavioural experiments with human participants.

Chapter 8

Conclusion

8.1 Aim of the Thesis

The aim of this thesis was to explore whether non-homogeneous inhibitory connections between letter and word representations could at least partially explain effects previously attributed to a lexical competition mechanism.

8.2 Summary of Results

8.2.1 Experiments 1.1 - 1.3

My first set of experiments introduced neighbourly primes: nonwords constructed by taking each letter of the target and replacing it with a letter that, if substituted in at that position, would form a one-letter-different substitution neighbour of the word. Through masked priming lexical decision tasks, I found that these neighbourly primes could significantly inhibit response times, so long as identity primes were also present during the experiment (the differences were still numerically in the same inhibitory direction without identity primes in the same experiment).

Motivated by previous research showing that spelling proficiency could be a good indicator of lexical precision (Andrews, 2012), I also investigated whether spelling and vocabulary scores might reveal differences between participants. I found very little variation between scores for my participants, suggesting that they were all extremely proficient, so spelling and vocabulary were not measured in subsequent experiments.

8.2.2 Experiments 2.1 - 2.5

The next set of experiments was designed to test whether the neighbourly inhibition effect could also be elicited with a single letter. In these experiments, the neighbourly primes were a single letter, taken from a neighbour of the target, presented in the differing letter location (restricted to initial and final letters).

Again, I found significant neighbourly priming. To rule out this being due to perceptual blending between the prime and the target, 33ms backward masks were introduced. Both masks (##### and %%%%%) appeared to prevent any priming effects, including identity priming, and increasing the power of the experiment by using more items also revealed no significant priming effects.

8.2.3 Experiment 3

Experiment 3 was an Alphabetic Decision Task with a forward mask (###) in all trials and a backward mask (%%%) for the prime in 50% of trials. The prime was either an identity or unrelated prime presented for 10, 30, 50, 70 or 90ms. I selected the NonEnglish foil letters in this experiment to be closely matched to the English letters used based on their features. With no backward mask, there was significant identity priming for primes with a duration of at least 50ms. When the backward mask was present, only the 90ms primes caused significant priming effects.

8.2.4 Experiment 4

My final experiment was a sandwich priming experiment (Lupker & Davis, 2009) using single-letter neighbourly, identity or unrelated primes. Masks were interleaved between each stimulus presentation: ##### - preprime - %%%%% - prime - ##### - TARGET. Contrary to expectations, and previous experiments, both the identity and neighbourly primes yielded significant facilitatory priming relative to the unrelated primes. The identity primes caused slightly larger priming effects, but this difference was not significant.

8.2.5 Simulations

Simulations using the stimuli from Experiment 1.1 conducted using the *easyNet* software showed that a modified version of the IA model, whereby neighbourly letters had stronger inhibitory connections with word representations than unrelated letters, could perform lexical decision and demonstrated facilitatory identity priming and inhibitory neighbourly priming. The Base IA model did not show inhibitory neighbourly priming.

8.3 Integrated Discussion

Across the experiments utilising a standard masked priming procedure, neighbourly primes were shown to inhibit lexical decision, so long as identity primes were also present amongst the stimuli. This was the case for neighbourly primes created by combining neighbourly letters and for one-letter neighbourly primes. When a backward mask was introduced after the prime, neither neighbourly nor identity priming was found. An Alphabetic Decision Task revealed that priming from a single letter does not manifest with a 33ms backward mask, unless the prime is over 70ms in duration, at which point it ceases to be a reliably subconscious prime. This would expose the paradigm to participant strategies, limiting its utility.

Sandwich priming, however, did not yield the inhibitory priming that was predicted. Instead, when one-letter neighbourly primes were used in a sandwich priming lexical decision task, I found facilitatory effects that were not statistically different from identity priming effects. This is difficult to square away with the theories of neighbourly priming that motivated this research. First, I want to address whether this result is possible to understand in a way that is compatible with the lexical competition mechanism that the IA model and its descendants implement.

Situations in which an orthographic neighbour receives significant activation should, according to models that implement lexical competition, inhibit recognition of that target. Perhaps for situations in which the target is already activated and there is slight activation of an orthographic neighbour, the overall activation level is raised for nodes connected to the target, but not significantly enough that they reach a competition threshold and inhibition is exhibited. This would suggest that orthographic neighbours do not always cause inhibition - perhaps this is true when activation of the neighbour reaches a certain threshold, but for activation levels below this point the general increase in nodes connected to a target raise its activation level leading to faster recognition. Competitive processes only become relevant if activation between prime and target are similar, otherwise any amount of lexical activation raises the overall activation of word representations, yielding faster "Yes" responses in a LDT.

This is not quite the same as predicted by the Multiple Read-Out Model (MROM; Grainger & Jacobs, 1996), in which global lexical activation can yield faster "Yes" responses. This mechanism in the MROM is based upon global lexical activation across all lexical units. Even unrelated single letter primes would contribute to the global lexical activation, so I would expect to see no difference between unrelated and neighbourly primes. Instead, I am arguing that the connections between lexical units can be facilitatory at low activation levels, a mechanism which is not currently implemented by any of the models based upon the IA model.

Another possibility is that the activation of the target from the prepprime inhibits the neighbour to the extent that, when activated by the neighbourly prime, the neighbour representation cannot inhibit the target, but not to the extent that the partially activated neighbour cannot raise the activation level of the letters in that neighbour word. These letters, with the exception of the neighbourly letter, are also present in the target word, hence the facilitation. Unrelated primes will not activate any orthographic neighbours of the target word, so feedback activation to the letter level is less likely to activate letters that appear in the target word. I say "less likely" because unrelated letters were selected so as to not come from neighbours of the target as defined by Coltheart et al. (1977). The unrelated letter might still partially activate words that are 2- or 3-letter substitution neighbours of the target, which would then lead to feedback activation of some letters that appear in the target word. The net effect across all stimuli, however, is that identity and neighbourly primes are still more able to activate the target word representation before presentation of the target stimulus, causing facilitation relative to the unrelated primes.

This explanation is also plausible if we consider that the facilitatory effect comes from the raised activation at the word level caused by the feedback facilitation. If the representations of letters in the target word more quickly reach high activation levels, then this could reduce the time needed for lexical access upon target presentation. After all, Experiment 3 in this thesis has shown that single letter priming has a demonstrable effect.

This verbal explanation of the effects according to the IA model relies on quite strong feedback facilitation in the face of inhibition of the neighbourly word representation due to the preprime. Neighbourly mechanisms could also account for this result in a more plausible way.

According to the neighbourly mechanism hypothesised in this thesis, the connection between a word representation and the representations of any of its neighbourly letters are more inhibitory than those between when the word and non-neighbourly letters. So when the preprime activates the target representation, the neighbourly letters are more strongly inhibited than unrelated letters. When a neighbourly one-letter prime is presented, its activity will be impeded by the feedback inhibition from the word level due to the preprime[1]. An unrelated letter will not be inhibited to the same degree, and so is able to become activated, contributing an inhibitory influence on the target. In this way, unrelated primes inhibit

[1]The neighbourly inhibition implemented in my simulations was unidirectional, not bidirectional; only from the letter level to the word level. Both versions of the IA model used had only excitatory feedback connections. A model with non-homogeneous feedback inhibitions is conceivable, but has not yet been tested.

processing of the target more than identity or neighbourly primes, which the sandwich priming results indicated, and identity primes will facilitate relative to neighbourly primes, which was indicated in the direction of the results, though the difference was not statistically significant.

This explanation relies on acknowledging lateral inhibition in the letter layer, whilst this thesis is trying to provide an alternative to lateral inhibition in the word level. This might appear contradictory, but the point is not to suggest that there is zero lexical competition, but to argue that the the role of this mechanism has been exaggerated. So it is important to consider to what extent the neighbourly effects discussed in this thesis are compatible with the research regarding lexical inhibition.

The Interactive Activation model (McClelland & Rumelhart, 1981) accounts for inhibitory effects between word neighbours on the basis of direct inhibitory connections between word representations. This was supported by findings such as those from Segui and Grainger (1990) where higher frequency neighbour words caused inhibition. This could be caused by lateral inhibition from one word representation to another, or it could be caused by bottom-up inhibition from the neighbourly letters themselves, with the frequency effects potentially explained by stronger feedback to that inhibitory neighbourly letter representation if the neighbour word is of higher frequency, in turn causing stronger inhibition from the letter-to-word connection.

We also need to account for the many experiments that have revealed a facilitatory effect of neighbour priming (Forster & Veres, 1998, e.g.,). This appears to be, at least in part, affected by task difficulty; when the lexical decision task is easier (for instance, because there are obvious orthotactic differences between the words and nonwords used) then it is not necessary to identify specific words to perform well on the task. The global activation account from Jacobs and Grainger (1992) explained this with the possibility of producing a “Yes” response in a lexical decision task using summed word activation levels across all word nodes instead of a specific word node reaching a threshold activation. These findings are compatible with both a neighbourly inhibition mechanism and a lateral inhibition mechanism, as the source of the inhibition to a specific word node (from a neighbouring word or a neighbourly letter representation) has no bearing on the overall amount of lexical activation.

Davis and Lupker (2006) made a change to the IA model that implemented non-homogeneous, selective inhibition at the word level with inhibitory signals only sent between words that share at least one letter in the same position. Davis and Lupker argued that reducing the inhibitory effect of unrelated primes on target recognition provided simulated results that better fit human data. This model successfully predicted the inhibitory priming elicited by one-letter-different orthographic neighbours. However, this selective inhibition model does not account for the null or facilitatory effects of neighbours in English (e.g., Forster & Shen, 1996). Siakaluk et al. (2002) found that backward- and forward-masked word primes facilitated identification of a target word, which poses a particular problem for the lateral inhibition account as global lexical activation would not allow accurate identification of specific words.

Does this finding pose a problem for neighbourly inhibition? A short-duration, backward masked prime might not produce sufficient neighbourly inhibitory priming. Instead, the prime provides feedback activation for the letters shared between the prime and the target, whose activations contribute towards a facilitatory effect, in a similar way to my explanation for the sandwich priming effects described above. It is possible that lateral inhibition also only provides 'enough' inhibition at certain thresholds, but this is not how it is presently implemented by the IA model (McClelland & Rumelhart, 1981), nor by Davis and Lupker (2006).

The simulations presented in Chapter 7 have shown that a neighbourly mechanism can result in reasonably accurate lexical decision. Simulations of the sandwich priming paradigm described in Chapter 6 would help clarify whether facilitatory effects from neighbourly primes are consistent with the non-homogeneous inhibitory connections between letters and words[2].

[2] I have not been able to find the correct parameters to enable the model to perform accurate lexical decision on single-letter primes in *easyNet*. A partial prime like t+++ partially activates many word nodes (any words beginning with T), and so the total lexical activity increases significantly as soon as the prime is presented, leading to

There is a clear argument that a lateral inhibition mechanism that would suppress the activation of the majority of four-letter words in the English language is a counter-intuitive one (Andrews, 1997; Siakaluk et al., 2002). This thesis provides evidence that neighbourly inhibition is possible and viable, whilst also providing a potential way that tuning could occur to explain the effects of readers at different developmental stages or levels of ability (e.g., Andrews, 2012; Andrews & Hersch, 2010).

8.4 Concluding Remarks

This thesis sought to explore a potential mechanism for models of lexical access via investigating the effects of neighbourly primes on lexical decision. Lexical competition mechanisms appear to be insufficient to explain the full range of effects seen in English lexical decision tasks, whilst non-homogeneous letter-word inhibition does appear to be a mechanism that could help to model lexical access processes more accurately, and warrants further investigation. This mechanism may also explain individual differences and developmental effects, as changing inhibitory connections between the letter and word levels offers a clear method of tuning the lexical access system over time with language exposure.

This research has also demonstrated that position-specific single letter primes (e.g., `+++d`) can elicit significant priming effects, and so could be used in future investigations into the relationships between letters and words, and the connections between them. Furthermore, I believe that the NonEnglish stimuli used in the Alphabetic Decision Task in Experiment 3 (see Table 5.1 on page 69) provide a suitable foil for English letters, and so might have further utility in investigations into visual letter identification.

8.4.1 Limitations and Future Research

Whilst I am confident that the thesis shows neighbourly inhibitory processes to be a viable candidate for models of lexical access in proficient

"Yes" responses to virtually all nonwords.

readers, there are plenty of reasons to be cautious about applying these conclusions too widely. One of the major limitations with the research described in this thesis is a lack of statistical power. Each experiment was carried out using as many participants as opportunity and budgets allowed at the time, but an experiment with a much larger participant pool would have greater power, allowing us to be more confident that if neighbourly inhibition is a real effect, it will be detected.

These experiments were limited to the English language, but also to educated, proficient readers. It is not clear how these results would generalise to other languages or developing readers. Languages other than English have been largely ignored, both in terms of whether, for example, Chinese would work the same way, but also in terms of assuming readers to be monolingual. Nonwords in my experiments were labelled so based only on an English lexicon - nonwords in other languages would likely be more difficult to reject for those who know the (non)word. The majority of my participants were university students, many of whom were bi- or poly-lingual. It is increasingly accepted that bilinguals access words based on letter overlap rather than language membership. The Bilingual Interactive Activation model (BIA), for example, implements language membership as a layer above the word level, influencing selection via top-down processes (Van Heuven et al., 1998). I only ever considered the impact of orthographic neighbours within English, whilst orthographic neighbours from L1 can exert an inhibitory effect on words in L2 (Mulder et al., 2018). My assumption is that the neighbourly mechanisms would remain unchanged when considering bilingual readers and that cross-language neighbours would have the same effect as within-language neighbours in tuning the letter-word inhibitory connections, but more research would be needed to investigate this further.

More research is also needed to further clarify the role of spelling ability in this tuning process. Experiment 1.2 attempted to investigate the effect of spelling ability (as a measure of language proficiency) on the presence of neighbourly inhibition. The lack of an effect led to this line of inquiry being dropped for future experiments, partly because only those extremely proficient in English were wanted as participants in order to maximise the likelihood finding a neighbourly effect. Now that I have presented stronger evidence that such an effect does manifest, future research could investigate whether this effect is only apparent for the most proficient readers. Better yet, a longitudinal study that shows no neighbourly effect for developing readers followed by a clear inhibitory effect of neighbourly primes would provide strong evidence for lexical tuning via a neighbourly mechanism. It may also be possible to show the development of this neighbourly mechanism artificially by training participants on novel neighbours, in a paradigm similar to that used by Bowers et al. (2005). This would likely require a lengthy training process in order to replicate the effects of years of language exposure, however.

The case for neighbourly inhibition would also be further strengthened were the modelling work reported in the thesis more complete. Whilst the modified version of the IA described in Chapter 7 achieved accuracy for nonword lexical decision above 70%, this is still quite low compared with accuracy rates seen in experimental data, suggesting that this model is not yet a good simulation of human cognitive processes during a lexical decision task. This model is a prime candidate for improvements in future research. It would also be advantageous to run simulations with the same model for the other paradigms described in this thesis; one-letter priming, sandwich priming and an alphabetic decision task. The central claim made by this thesis, that non-homogeneous inhibitory connections between letter and word representations provide a better explanation for empirical data than strong lexical competition mechanisms, would be greatly supported by experiments demonstrating robust neighbourly inhibition and simulations demonstrating that a model implementing the neighbourly mechanism is capable of producing this inhibitory effect alongside other robust effects found in the visual word recognition literature.

Overall, the greatest weakness of this thesis is that is had not provided substantive proof for a process of neighbourly inhibition in visual lexical access. I believe that I have made a strong case for the *possibility* of effects previously attributed to lexical competition being driven by letter-to-word neighbourly inhibition, but the possibility remains that the pattern of results found in these experiments are caused by inhibitory effects exerted by lexical neighbours of the target word. Perhaps future experiments and simulations will more decisively support (or disprove) the central argument of this thesis.

References

Adelman, J. S. (2012). Methodological issues with words. In J. S. Adelman (Ed.), *Visual word recognition volume 1: Models and methods, orthography and phonology* (pp. 116–138). Psychology Press.

Adelman, J. S., Davis, C. J. & Gubian, M. (2017). *easyNet*. http://adelmanlab.org/easyNet/

Adelman, J. S., Johnson, R. L., McCormick, S. F., McKague, M., Kinoshita, S., Bowers, J. S., Perry, J. R., Lupker, S. J., Forster, K. I., Cortese, M. J. et al. (2014). A behavioral database for masked form priming. *Behavior research methods*, *46*(4), 1052–1067. https://doi.org/10.3758/s13428-013-0442-y

Andrews, S. (1989). Frequency and neighborhood size effects on lexical access: Activation or search? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 802–814.

Andrews, S. (1997). The effect of orthographic similarity on lexical retrieval: Resolving neighborhood conflicts. *Psychonomic Bulletin & Review*, *4*(4), 439–461. https://doi.org/10.3758/bf03214334

Andrews, S. (2012). Individual differences in skilled visual word recognition and reading: The role of lexical quality. In J. S. Adelman (Ed.), *Visual word recognition volume 2: Meaning and context, individuals and development* (pp. 163–184). Psychology Press.

Andrews, S. & Hersch, J. (2010). Lexical precision in skilled readers: Individual differences in masked neighbor priming. *Journal of*

Experimental Psychology: General, *139*(2), 299–318. https://doi.org/10.1037/a0018366

Andrews, S. & Lo, S. (2012). Not all skilled readers have cracked the code: Individual differences in masked form priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(1), 152–163. https://doi.org/10.1037/a0024953

Baayen, R. H., Piepenbrock, R. & Gulikers, L. (1996). *The CELEX lexical database*. https://doi.org/10.35111/gs6s-gm48

Balota, D. & Chumbley, J. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(3), 340–357. https://doi.org/10.1037/0096-1523.10.3.340

Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B. & Treiman, R. (2007). The English lexicon project. *Behavior Research Methods*, *39*(3), 445–459. https://doi.org/10.3758/bf03193014

Barr, D. J., Levy, R., Scheepers, C. & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Bates, D., Mächler, M., Bolker, B. & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Bijeljac-Babic, R., Biardeau, A. & Grainger, J. (1997). Masked orthographic priming in bilingual word recognition. *Memory and Cognition*, *25*(4), 447–457. https://doi.org/10.3758/bf03201121

Bird, H., Franklin, S. & Howard, D. (2001). Age of acquisition and imageability ratings for a large set of words, including verbs and function words. *Behavior Research Methods, Instruments, & Computers*, *33*(1), 73–79. https://doi.org/10.3758/BF03195349

Bodner, G. E., Masson, M. E. J. & Richard, N. T. (2006). Repetition proportion biases masked priming of lexical decisions. *Memory* *and Cognition*, *34*, 1298–1311. https://doi.org/10.3758/bf03193273

Bowers, J. S., Davisa, C. J. & Hanley, D. A. (2005). Interfering neighbours: The impact of novel word learning on the identification of visually similar words. *Cognition*, *97*, B45–B54. https://doi.org/10.1016/j.cognition.2005.02.002

Brysbaert, M., Lange, M. & Wijnendaele, I. V. (2000). The effects of age-of-acquisition and frequency-of-occurrence in visual word recognition: Further evidence from the dutch language. *European Journal of Cognitive Psychology*, *12*(1), 65–85. https://doi.org/10.1080/095414400382208

Burt, J. S. & Duncum, S. (2017). Masked form priming is moderated by the size of the letter-order-free orthographic neighbourhood. *The Quarterly Journal of Experimental Psychology*, *70*(1), 127–141. https://doi.org/10.1080/17470218.2015.1126289

Burt, J. S. & Tate, H. (2002). Does a reading lexicon provide orthographic representations for spelling? *Journal of Memory and Language*, *46*(3), 518–543. https://doi.org/10.1006/jmla.2001.2818

Carreiras, M., Perea, M. & Grainger, J. (1997). Effects of orthographic neighborhood in visual word recognition: Cross-task comparisons. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(4), 857–871. https://doi.org/10.1037/0278-7393.23.4.857

Castles, A., Davis, C., Cavalot, P. & Forster, K. I. (2007). Tracking the acquisition of orthographic skills in developing readers: Masked priming effects. *Journal of Experimental Child Psychology*, *97*(3), 165–182. https://doi.org/10.1016/j.jecp.2007.01.006

Castles, A., Davis, C. & Letcher, T. (1999). Neighbourhood effects on masked form priming in developing readers. *Language and Cognitive Processes*, *14*(2), 201–224. https://doi.org/10.1080/016909699386347

Chambers, S. M. (1979). Letter and order information in lexical access. *Journal of Verbal Learning and Verbal Behaviour*, *18*, 225–241. https://doi.org/10.1016/s0022-5371(79)90136-1

Coltheart, M., Davelaar, E., Jonasson, J. T. & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and performance VI* (pp. 535–555). Academic Press.

Coltheart, M., Rastle, K., Perry, C., Langdon, R. & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, *108*(1), 204–256. https://doi.org/10.1037/0033-295x.108.1.204

Comesaña, M., Soares, A. P., Marcet, A. & Perea, M. (2016). On the nature of consonant/vowel differences in letter position coding: Evidence from developing and adult readers. *British Journal of Psychology*, *107*(4), 651–674. https://doi.org/10.1111/bjop.12179

Davis, C. (2005). N-Watch: A program for deriving neighborhood size and other psycholinguistic statistics. *Behavior Research Methods*, *37*(1), 65–70. https://doi.org/10.3758/BF03206399

Davis, C. & Lupker, S. (2006). Masked inhibitory priming in English: Evidence for lexical inhibition. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(3), 668–687. https://doi.org/10.1037/0096-1523.32.3.668

Davis, C. J. (1999). *The self-organising lexical acquisition and recognition (SOLAR) model of visual word recognition* (Doctoral dissertation). University of New South Wales, Syndey, Australia.

Davis, C. J. (2003). Factors underlying masked priming effects in competitive network models of visual word recognition. In S. Kinoshita & S. J. Lupker (Eds.), *Masked priming: The state of the art* (pp. 121–170). Psychology Press.

Davis, C. J. (2010). The spatial coding model of visual word identification. *Psychological Review*, *117*(3), 713–758. https://doi.org/10.1037/a0019738

Davis, C. J. & Bowers, J. S. (2004). What do letter migration errors reveal about letter position coding in visual word recognition? *Journal of Experimental Psychology*, *30*(5), 923–941. https://doi.org/10.1037/0096-1523.30.5.923

Davis, C. J. & Bowers, J. S. (2006). Contrasting five different theories of letter position coding: Evidence from orthographic similarity effects. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(3), 535–557. https://doi.org/10.1037/0096-1523.32.3.535

De Moor, W. & Brysbaert, M. (2000). Neighborhood-frequency effects when primes and targets are of different lengths. *Psychological Research*, *63*(2), 159–162. https://doi.org/doi.org/10.1007/PL00008174

Drews, E. & Zwitserlood, P. (1995). Morphological and orthographic similarity in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(5), 1098–1116. https://doi.org/10.1037/0096-1523.21.5.1098

Duyck, W. & Warlop, N. (2009). Translation priming between the native language and a second language: New evidence from Dutch-French bilinguals. *Experimental Psychology*, *56*(3), 173–179. https://doi.org/10.1027/1618-3169.56.3.173

Evett, L. J. & Humphreys, G. W. (1981). The use of abstract graphemic information in lexical access. *The Quarterly Journal of Experimental Psychology*, *33*(4), 325–350. https://doi.org/10.1080/14640748108400797

Forster, K. & Veres, C. (1998). The prime lexicality effect: Form-priming as a function of prime awareness, lexical status, and discrimination difficulty. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(2), 498–514. https://doi.org/10.1037/0278-7393.24.2.498

Forster, K. I. (1987). Form-priming with masked primes: The best match hypothesis. In M. Coltheart (Ed.), *Attention and performance XII* (pp. 127–146). Erlbaum.

Forster, K. I., Davis, C., Schoknecht, C. & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *The Quarterly Journal of Experimental Psychology Section A*, *39*(2), 211–251. https://doi.org/10.1080/14640748708401785

Forster, K. I. & Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(4), 680–698. https://doi.org/10.1037/0278-7393.10.4.680

Forster, K. I. & Forster, J. C. (2003). DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, & Computers*, *35*(1), 116–124. https://doi.org/10.3758/bf03195503

Forster, K. I., Mohan, K. & Hector, J. (2003). The mechanics of masked priming. In S. Kinoshita & S. J. Lupker (Eds.), *Masked priming: The state of the art* (pp. 2–20). Psychology Press.

Forster, K. I. & Shen, D. (1996). No enemies in the neighborhood: Absence of inhibitory neighborhood effects in lexical decision and semantic categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(3), 696–713. https://doi.org/10.1037/0278-7393.22.3.696

Forster, K. I. & Taft, M. (1994). Bodies, antibodies, and neighborhood-density effects in masked form priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(4), 844–863. https://doi.org/10.1037/0278-7393.20.4.844

Fox, J. & Weisberg, S. (2019). *An R companion to applied regression* (Third). Sage. https://socialsciences.mcmaster.ca/jfox/Books/Companion/

Grainger, J. & Ferrand, L. (1994). Phonology and orthography in visual word recognition: Effects of masked homophone primes. *Journal of Memory and Language*, *33*(2), 218–233. https://doi.org/10.1006/jmla.1994.1011

Grainger, J. & Jacobs, A. M. (1993). Masked partial-word priming in visual word recognition: Effects of positional letter frequency.

Journal of Experimental Psychology: Human Perception and Performance, *19*(5), 951–964. https://doi.org/10.1037/0096-1523.19.5.951

Grainger, J. & Frenck-Mestre, C. (1998). Masked priming by translation equivalents in proficient bilinguals. *Language and Cognitive Processes*, *13*(6), 601–623. https://doi.org/10.1080/016909698386393

Grainger, J. & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, *103*(3), 518–565. https://doi.org/10.1037/0033-295x.103.3.518

Grainger, J. & Jacobs, A. M. (1991). Masked constituent letter priming in an alphabetic decision task. *European Journal of Cognitive Psychology*, *3*(4), 413–434. https://doi.org/10.1080/09541449108406237

Grainger, J., O'Regan, J. K., Jacobs, A. M. & Segui, J. (1989). On the role of competing word units in visual word recognition: The neighborhood frequency effect. *Perception & Psychophysics*, *45*(3), 189–195. https://doi.org/10.3758/bf03210696

Guerrera, C. & Forster, K. I. (2008). Masked form priming with extreme transposition. *Language and Cognitive Processes*, *23*(1), 117–142. https://doi.org/10.1080/01690960701579722

Henderson, B. A. (2013). *BlitzMax*. https://blitzmax.org

Jacobs, A. & Grainger, J. (1992). Testing a semistochastic variant of the interactive activation model in different word recognition experiments. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(4), 1174–1188. https://doi.org/10.1037/0096-1523.18.4.1174

Jacobs, A. M. & Grainger, J. (1991). Automatic letter priming in an alphabetic decision task. *Perception & Psychophysics*, *49*(1), 43–52. https://doi.org/10.3758/BF03211615

Kučera, H. & Francis, W. N. (1967). *Computational analysis of present-day American English*. Brown University Press.

Landauer, T. K. & Streeter, L. A. (1973). Structural differences between common and rare words: Failure of equivalence assumptions for theories of word recognition. *Journal of Verbal Learning and Verbal Behavior*, *12*(2), 119–131. https://doi.org/10.1016/s0022-5371(73)80001-5

Lenth, R. V. (2016). Least-squares means: The R package lsmeans. *Journal of Statistical Software*, *69*(1), 1–33. https://doi.org/10.18637/jss.v069.i01

Lenth, R. V. (2020). *Emmeans: Estimated marginal means, aka least-squares means* [R package version 1.5.2-1]. https://CRAN.R-project.org/package=emmeans

Loncke, M., Martensen, H., Van Heuven, W. J. B. & Sandra, D. (2009). Who is dominating the dutch neighbourhood? On the role of subsyllabic units in dutch nonword reading. *Quarterly Journal of Experimental Psychology*, *62*(1), 140–154. https://doi.org/10.1080/17470210701851206

Lupker, S. J. & Davis, C. J. (2009). Sandwich priming: A method for overcoming the limitations of masked priming by reducing lexical competitor effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(3), 618–639. https://doi.org/10.1037/a0015278

Lupker, S. J., Perea, M. & Davis, C. J. (2008). Transposed-letter effects: Consonants, vowels and letter frequency. *Language and Cognitive Processes*, *23*(1), 93–116. https://doi.org/10.1080/01690960701579714

McClelland, J. L. & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*(5), 375–405. https://doi.org/10.1037/0033-295x.88.5.375

Mewhort, D. J. K. & Campbell, A. J. (1978). Processing spatial information and the selective-masking effect. *Perception & Psychophysics*, *24*, 93–101. https://doi.org/10.3758/BF03202978

Monsell, S., Doyle, M. C. & Haggard, P. N. (1989). Effects of frequency on visual word recognition tasks: Where are they? *Journal of*

Experimental Psychology: General, *118*(1), 43–71. https://doi.org/10.1037/0096-3445.118.1.43

Mulder, K., Van Heuven, W. J. B. & Dijkstra, T. (2018). Revisiting the neighborhood: How L2 proficiency and neighborhood manipulation affect bilingual processing. *Frontiers in Psychology*, *9*, 1860. https://doi.org/10.3389/fpsyg.2018.01860

Nakayama, M., Sears, C. & Lupker, S. (2008). Masked priming with orthographic neighbors: A test of the lexical competition assumption. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 1236–1260. https://doi.org/10.1037/0096-1523.34.5.1236

Norris, D. & Kinoshita, S. (2008). Perception as evidence accumulation and Bayesian inference: Insights from masked priming. *Journal of Experimental Psychology: General*, *137*(3), 434–455. https://doi.org/10.1037/a0012799

Perea, M. & Pollatsek, A. (1998). The effects of neighborhood frequency in reading and lexical decision. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(3), 767–779. https://doi.org/10.1037/0096-1523.24.3.767

Perea, M. & Acha, J. (2009). Does letter position coding depend on consonant/vowel status? Evidence with the masked priming technique. *Acta Psychologica*, *130*(2), 127–137. https://doi.org/10.1016/j.actpsy.2008.11.001

Perea, M. & Lupker, S. J. (2004). Can CANISO activate CASINO? Transposed-letter similarity effects with nonadjacent letter positions. *Journal of Memory and Language*, *51*, 231–246. https://doi.org/10.1016/j.jml.2004.05.005

Perry, J. R., Lupker, S. J. & Davis, C. J. (2008). An evaluation of the interactive-activation model using masked partial-word priming. *Language and Cognitive Processes*, *23*(1), 36–68. https://doi.org/10.1080/01690960701578112

Peterlin, P. & White, S. (2016). *GNU FreeFont*. https://www.gnu.org/software/freefont

R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. https://www.R-project.org/

Rayner, K. (1979). Eye guidance in reading: Fixation locations within words. *Perception*, *8*(1), 21–30. https://doi.org/10.1068/p080021

Reicher, G. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, *81*(2), 275–280. https://doi.org/10.1037/h0027768

Sears, C. R., Campbell, C. R. & Lupker, S. J. (2006). Is there a neighborhood frequency effect in English?: Evidence from reading and lexical decision. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(4), 1040–1062. https://doi.org/10.1037/0096-1523.32.4.1040

Segui, J. & Grainger, J. (1990). Priming word recognition with orthographic neighbors: Effects of relative prime-target frequency. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(1), 65–76. https://doi.org/10.1037/0096-1523.16.1.65

Sereno, J. A. (1991). Graphemic, associative, and syntactic priming effects at a brief stimulus onset asynchrony in lexical decision and naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(3), 459–477. https://doi.org/10.1037/0278-7393.17.3.459

Shipley, W. C. (1940). A self-administering scale for measuring intellectual impairment and deterioration. *The Journal of Psychology*, *9*(2), 371–377. https://doi.org/10.1080/00223980.1940.9917704

Siakaluk, P. D., Sears, C. R. & Lupker, S. J. (2002). Orthographic neighborhood effects in lexical decision: The effects of nonword orthographic neighborhood size. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(3), 661–681. https://doi.org/10.1037/0096-1523.28.3.661

Trifonova, I. V. & Adelman, J. S. (2018). The sandwich priming paradigm does not reduce lexical competitor effects. *Journal of*

Experimental Psychology: Learning, Memory, and Cognition, *44*(11), 1743–1764. https://doi.org/10.1037/xlm0000542

Unicode, Inc. (2018). *The Unicode Standard, version 10.0*. https://www.unicode.org

Van Assche, E. & Grainger, J. (2006). A study of relative-position priming with superset primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 399–415. https://doi.org/10.1037/0278-7393.32.2.399

Van Heuven, W. J. B., Dijkstra, T. & Grainger, J. (1998). Orthographic neighborhood effects in bilingual word recognition. *Journal of Memory and Language*, *39*(3), 458–483. https://doi.org/10.1006/jmla.1998.2584

Van Heuven, W. J. B., Dijkstra, T., Grainger, J. & Schriefers, H. (2001). Shared neighborhood effects in masked orthographic priming. *Psychonomic Bulletin & Review*, *8*(1), 96–101. https://doi.org/10.3758/BF03196144

Velan, H. & Frost, R. (2009). Letter-transposition effects are not universal: The impact of transposing letters in Hebrew. *Journal of Memory and Language*, *61*(3), 285–302. https://doi.org/10.1016/j.jml.2009.05.003

Welvaert, M., Farioli, F. & Grainger, J. (2008). Graded effects of number of inserted letters in superset priming. *Experimental Psychology*, *55*(1), 54–63. https://doi.org/10.1027/1618-3169.55.1.54

White, S. J., Johnson, R. L., Liversedge, S. P. & Rayner, K. (2008). Eye movements when reading transposed text: The importance of word-beginning letters. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(5), 1261–1276. https://doi.org/10.1037/0096-1523.34.5.1261

Appendix A

Experimental Stimuli

The following stimuli were used in Experiment 1.1.

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
CLAP	W	frim	gobe	clap
COPS	W	huwy	filn	cops
CURE	W	sobl	wifk	cure
COME	W	hakb	lygt	come
FARE[1]	W	bicm	nlus	fare
BLEW[1]	W	frod	gatp	blew
HARE[1]	W	divm	josg	hare
CAST	W	vorh	dikb	cast
LACE	W	fink	bodh	lace
LARK[1]	W	bucd	gose	lark
BANK	W	tucg	meif	bank
HUNT	W	cirg	bose	hunt
COAT	W	ghsx	pukd	coat
LONE	W	divg	fukh	lone
NODE[1]	W	muts	jiak	node
COIN	W	jhrl	bate	coin
BEAT	W	mrld	wugy	beat

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
CROP	W	dhaw	bukn	crop
PALE	W	wufk	divm	pale
DUNE[1]	W	tokg	vash	dune
FOND[1]	W	bult	wacs	fond
PUSS[1]	W	fabh	wock	puss
CORE	W	gamk	hily	core
BITS	W	hade	cogy	bits
LADS	W	piwy	menh	lads
HUGS	W	bote	pknl	hugs
MOSS	W	bapt	fhnk	moss
DEED[1]	W	fyar	ckot	deed
PORT[1]	W	fauk	cliw	port
BOAT	W	celr	wupk	boat
LANE	W	cikd	fhub	lane
SALE	W	dokt	chib	sale
BASH	W	dute	fniy	bash
RAGS	W	wuye	ckod	rags
HARD	W	cenm	ftib	hard
FIRE	W	davm	cpub	fire
ROBS[1]	W	jide	pnat	robs
LEAD	W	ronf	gcit	lead
BATS	W	ceyh	kluw	bats
PINK	W	wuct	hamd	pink
COPE	W	hany	bult	cope
BOLT	W	jead	gimk	bolt
PLAY	W	cron	dimk	play
HALL	W	biuf	ryks	hall

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
POST	W	mauh	bixd	post
BINS	W	guad	hely	bins
MUCK[1]	W	bosh	geip	muck
DARE	W	fizk	yonp	dare
CUTS	W	gobe	frim	cuts
DAME	W	filn	huwy	dame
DART	W	wifk	sobl	dart
DIES	W	lygt	hakb	dies
FEED	W	nlus	bicm	feed
HUMS[1]	W	gatp	frod	hums
HUNK[1]	W	josg	divm	hunk
LAME	W	dikb	vorh	lame
LASS	W	bodh	fink	lass
LIFT	W	gose	bucd	lift
LOAN	W	meif	tucg	loan
MATH[1]	W	bose	cirg	math
MILE	W	pukd	ghsx	mile
OATS	W	fukh	divg	oats
PEER	W	jiak	muts	peer
PINS	W	bate	jhrl	pins
PITS[1]	W	wugy	mrld	pits
PORE[1]	W	bukn	dhaw	pore
RANT	W	divm	wufk	rant
RICE	W	vash	tokg	rice
RIPE	W	wacs	bult	ripe
RISE	W	wock	fabh	rise
ROSE	W	hily	gamk	rose

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
RUBS[1]	W	cogy	hade	rubs
RUST[1]	W	menh	piwy	rust
SAID[1]	W	pknl	bote	said
SEED	W	fhnk	bapt	seed
SHIP[1]	W	ckot	fyar	ship
SHOP	W	cliw	fauk	shop
SINS[1]	W	wupk	celr	sins
SLAG	W	fhub	cikd	slag
SLAM	W	chib	dokt	slam
SLAP	W	fniy	dute	slap
SLIP	W	ckod	wuye	slip
SLOP[1]	W	ftib	cenm	slop
SLOT[1]	W	cpub	davm	slot
SLUG	W	pnat	jide	slug
SLUM	W	gcit	ronf	slum
SNOB	W	kluw	ceyh	snob
SOLE[1]	W	hamd	wuct	sole
SORE	W	bult	hany	sore
SUNS[1]	W	gimk	jead	suns
TALE	W	dimk	cron	tale
TAPE	W	ryks	biuf	tape
TENT	W	bixd	mauh	tent
TINT	W	hely	guad	tint
TRAM[1]	W	geip	bosh	tram
WARD[1]	W	yonp	fizk	ward
NIPE	N	wocs	dlab	nipe
MOSE	N	ruvt	cwig	mose

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
CALD	N	borf	guti	cald
DADE	N	wuzs	vinh	dade
WARK	N	moln	cigh	wark
HOAL	N	cewx	pusn	hoal
LONK	N	micg	pudt	lonk
LOID	N	varn	kesm	loid
NUME	N	fadb	khop	nume
HUST	N	gork	weiy	hust
DUTE	N	caky	fovl	dute
PALT	N	hecm	fogk	palt
DISE	N	roch	lugt	dise
JURS	N	oagy	hice	jurs
HOLK	N	funy	wirt	holk
LUSE	N	forh	pacy	luse
MUME	N	fils	bhop	mume
SOUT	N	phfr	mand	sout
PRAS	N	beoy	hunf	pras
DATS	N	bome	wilg	dats
SOND	N	belg	crim	sond
FOLL	N	dawk	cybn	foll
SONE	N	bamg	wuvh	sone
SULF	N	gerk	htan	sulf
DEEL	N	huar	wovs	deel
SELT	N	manf	cody	selt
HAGE	N	cuvs	boik	hage
RINT	N	hufg	wyds	rint
FASE	N	vudt	giln	fase

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
SLUB	N	cnag	peid	slub
BALT	N	meid	wucf	balt
RUDS	N	boge	clat	ruds
FARN	N	yewm	hicg	farn
HARL	N	euid	wocs	harl
LINS	N	bedk	cuah	lins
MISK	N	ralt	fogp	misk
POSS	N	lidh	cwfk	poss
NOPS	N	tade	rugk	nops
PONE	N	biky	mufh	pone
LUME	N	firp	hont	lume
DUNT	N	hesg	mlix	dunt
GOLL	N	diaf	cesk	goll
SALD	N	boie	wrup	sald
SULT	N	caik	honf	sult
BOAN	N	lert	hicg	boan
TOAR	N	beud	gimh	toar
FOUD	N	lenr	bagp	foud
BUMB	N	nolp	terh	bumb
GRUM	N	dlab	wocs	grum
SHAP	N	cwig	ruvt	shap
SEMS	N	guti	borf	sems
BOTE	N	vinh	wuzs	bote
WASE	N	cigh	moln	wase
HORT	N	pusn	cewx	hort
RENS	N	pudt	micg	rens
FILT	N	kesm	varn	filt

Target	**Target Lexicality**	**Neighbourly Prime**	**Unrelated Prime**	**Identity Prime**
SNIT	N	khop	fadb	snit
MAND	N	weiy	gork	mand
GIRE	N	fovl	caky	gire
TUSS	N	fogk	hecm	tuss
PESS	N	lugt	roch	pess
MULK	N	hice	oagy	mulk
MALK	N	wirt	funy	malk
TINK	N	pacy	forh	tink
TRAT	N	bhop	fils	trat
WILK	N	mand	phfr	wilk
CALT	N	hunf	beoy	calt
HANT	N	wilg	bome	hant
WHAP	N	crim	belg	whap
TURE	N	cybn	dawk	ture
FISE	N	wuvh	bamg	fise
SOOP	N	htan	gerk	soop
GINE	N	wovs	huar	gine
RUBE	N	cody	manf	rube
PARN	N	boik	cuvs	parn
TIPE	N	wyds	hufg	tipe
HORE	N	giln	vudt	hore
SOLL	N	peid	cnag	soll
HALK	N	wucf	meid	halk
SHUM	N	clat	boge	shum
LUNK	N	hicg	yewm	lunk
FANT	N	wocs	euid	fant
GOST	N	cuah	bedk	gost

Target	Target Lexicality	Neighbourly Prime	Unrelated Prime	Identity Prime
HUME	N	fogp	ralt	hume
SOAT	N	cwfk	lidh	soat
SICH	N	rugk	tade	sich
DAST	N	mufh	biky	dast
BELD	N	hont	firp	beld
COAN	N	mlix	hesg	coan
BULT	N	cesk	diaf	bult
CHAM	N	wrup	boie	cham
SELD	N	honf	caik	seld
SONK	N	hicg	lert	sonk
LASP	N	gimh	beud	lasp
JUMS	N	bagp	lenr	jums
DASK	N	terh	nolp	dask

Table A.1: Stimuli used in Experiment 1.1

[1]Data from these target words were excluded from the final analysis as the mean accuracy of responses was below 90%.

Appendix B

Simulation Parameters

Parameter Name	easyNet Parameter Name	Value
Rate	rate	0.1
Feature-Letter Excitation	flc::excitation	0.005
Feature-Letter Inhibition	flc::inhibition	-0.5
Letter-Word Excitation[1]	loc::excitation	0.07
Letter-Word Inhibition[1]	loc::inhibition	-0.04
Word-Letter Excitation	olc::excitation	0.3
Word-Letter Inhibition	olc::inhibition	0
Lateral Word Inhibition	lato::inhibition	-0.21
Lateral Letter Inhibition	leto::inhibition	0
Word Resting Level Multiplier	words::resting_multiplier	0.05
Letter Activation Decay	letters::decay	0.007
Word Activation Decay	words::decay	0.007
Minimum Activation Level	min_act	-0.2
Maximum Activation Level	max_act	1
Total Lexical Activation Multiplier	sum::excitation	1
Specific Word "Yes" Threshold	unit_decider::threshold	0.4

Parameter Name	easyNet Parameter Name	Value
Total Lexical Activation “Yes” Check Time	tla_yes::check_time	70
Total Lexical Activation “Yes” Check Threshold	tla_yes::check_threshold	0.65
Total Lexical Activation “Yes” Threshold (Low)	tla_yes::threshold_low	0.72
Total Lexical Activation “Yes” Threshold (High)	tla_yes::threshold_high	0.9
“No” Decision Check Time	mrom_no::check_time	120
“No” Decision Check Threshold	mrom_no::check_threshold	0.22
Early Timeout for “No” Decision	mrom_no::early_timeout	250
Late Timeout for “No” Decision	mrom_no::lat_timeout	300

Table B.1: Parameter settings for both the Neighbourly and Base IA Model simulations run in easyNet.

[1]These parameters are only relevant for the Base IA. The Neighbourly excitation and inhibition weights are described on page 82.