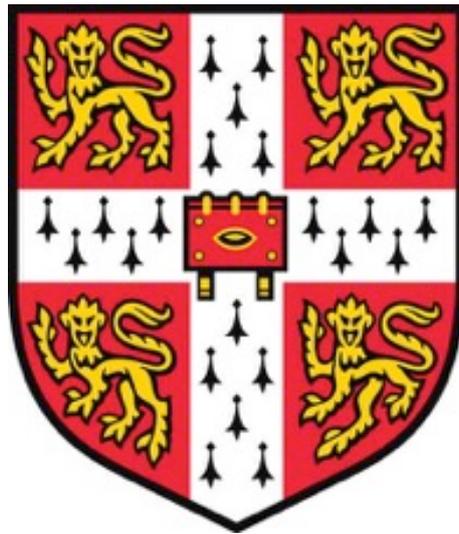


Learning, Consolidating, and Generalising Novel Morphology



Lydia Viñals-Castonguay
Lucy Cavendish College

Department of Theoretical and Applied Linguistics
University of Cambridge

September 2017

This dissertation is submitted to the University of Cambridge
for the degree of Doctor of Philosophy

Learning, Consolidating, and Generalising Novel Morphology

Lydia Viñals-Castonguay

Despite a central role for morphological knowledge in supporting linguistic generalisation, the neural representations supporting its learning remain largely unexplored. This thesis addressed this gap by exploring the role of memory consolidation in morphological learning and generalisation. In three experiments, adult participants learned an artificial language in which stems (e.g. *gleet*, *shiln*) combined with plural affixes (e.g. *-aff*, *-opp*; *gleetaff*, *shilnopp*) to refer to the occupation of multiple male and female characters. Mimicking properties of morphological systems in natural languages, the plurals varied in their phonological consistency/ambiguity and type/token frequency. Two sets of plurals, distinguished by gender, were trained on two successive days. Experiment 1 revealed that generalisation to novel phonologically ambiguous forms measured on the second day showed a greater influence of token frequency for plurals trained on the previous day, suggesting overnight changes in their underlying representations. Experiment 2 examined this effect further by using fMRI to compare the neural representations underlying plurals learned on the day of scanning or on the previous day. Representational Similarity Analysis revealed increased similarity structure among high type frequency plurals and reduced similarity structure among high token frequency plurals following overnight consolidation in the left superior temporal gyrus (STG). These results are consistent with a Complementary Learning Systems (CLS) model in which overnight consolidation supports the development of overlapping representations among several items sharing the same feature (here, an affix; type frequency) and strengthens item-specific representations for frequently occurring items (token frequency). Additionally, connectivity analyses showed that the functional coupling between the left STG and the left dorsolateral prefrontal cortex was weaker for high type frequency plurals and stronger for high token frequency plurals following overnight consolidation. These results suggest that the engagement of prefrontal control processes in retrieving the newly-learned plurals is subject to overnight consolidation and sensitive to the similarity structure underlying the plurals to be retrieved. However, the overnight changes in similarity structure and

functional networks observed in Experiment 2 were not mirrored by changes in generalisation to novel forms as were observed in Experiment 1. Experiment 3 aimed to address the discrepancy in consolidation-related changes in generalisation behaviour between the first two experiments. Type/token frequencies were manipulated to bias learning, consolidation, and generalisation towards high token frequency plurals. Despite this manipulation, no consolidation-related changes in generalisation were observed. Findings from all three experiments are interpreted in the context of the CLS model and a role for overnight consolidation in morphological learning and generalisation is discussed.

Table of contents

Abstract	3
Table of contents	5
List of figures	8
List of tables	10
Acknowledgements	11
Declaration	12
Chapter 1: Learning morphology	14
1.1. Learning rules or strengthening connections	14
1.2. Influences on learning inflections	18
1.2.1. Frequency	18
1.2.2. Phonological consistency.....	22
1.3. Neural bases of inflectional processing	24
1.4. Cross-linguistic inflectional processing	30
1.5. Summary	32
Chapter 2: Consolidating morphology	33
2.1. Origins of consolidation theory	33
2.2. Consolidation and Complementary Learning Systems	36
2.2.1. Sleep and CLS	41
2.3. Consolidation in language learning	45
2.3.1. Consolidation of speech sounds and words	45
2.3.2. Consolidation of morphology	49
2.4. Summary	51
Chapter 3: Generalising morphology	52
3.1. Generalisation	52
3.2. Generalisation and Complementary Learning Systems	55
3.3. Morphological generalisation	62
3.3.1. Frequency	63
3.3.2. Phonological consistency.....	65
3.3.3. Generalisation of morphology and CLS	67
3.4. Summary	69
Chapter 4: Learning, consolidating, and generalising morphology	71
4.1. Building bridges	71

4.2. Thesis outline	73
Chapter 5: Experiment 1.....	74
5.1. Introduction	74
5.2. Materials and methods.....	79
5.2.1. Participants	79
5.2.2. Experimental stimuli	80
5.2.3. Experimental design and procedure.....	82
5.2.4. Data analysis	87
5.3. Results.....	89
5.3.1. Training.....	89
5.3.2. Testing	92
5.3.3. Generalisation.....	95
5.4. Discussion.....	98
Chapter 6: The representation of morphology	104
6.1. Decoding mental representations.....	104
6.2. Representational Similarity Analysis	106
6.2. RSA and morphology	109
6.4. RSA and consolidation.....	110
6.5. Summary	112
Chapter 7: The connectivity of morphology	113
7.1. Functional connectivity.....	113
7.2. Psychophysiological interaction analysis	113
7.3. PPI and morphology	114
7.4. PPI and consolidation.....	116
7.5. Summary	118
Chapter 8: Experiment 2.....	119
8.1. Introduction	119
8.2. Materials and methods.....	122
8.2.1. Participants	122
8.2.2. Experimental stimuli	122
8.2.3. Experimental design and procedure.....	124
8.2.3.1. Day 9	124
8.2.4. fMRI scanning	125
8.2.5. Data analysis	128
8.3. Results.....	134

8.3.1. Behavioural results.....	134
8.3.2. Univariate results	144
8.3.3. RSA results	146
8.3.4. PPI results.....	154
8.4. Discussion.....	158
8.4.1. Learning and generalisation	158
8.4.2. Neural representations and functional networks	161
8.4.3. Summary	167
Chapter 9: Experiment 3.....	168
9.1. Introduction	168
9.2. Materials and methods.....	175
9.2.1. Participants	175
9.2.2. Experimental stimuli	175
9.2.3. Experimental design and procedure.....	177
9.3. Results.....	177
9.3.1. Training.....	177
9.3.2. Testing	179
9.3.3. Generalisation.....	182
9.4. Discussion.....	186
Chapter 10: Conclusions and future directions	190
10.1. Thesis summary.....	190
10.2. Main contributions	192
10.2.1. Overnight changes in generalisation	193
10.2.2. Overnight changes in neural representations	195
10.2.3. Overnight changes in functional networks	201
10.3. Limitations	203
10.4. Future directions	205
Appendix 1	209
Appendix 2	210
Appendix 3	211
References	212

List of figures

Figure 1. Complementary Learning Systems and their interactions	38
Figure 2. Sleep architecture and sleep stages	43
Figure 3. Hippocampal replay and neocortical consolidation	62
Figure 4. Mirković and Gaskell (in prep.): training stimuli set.....	75
Figure 5. Mirković and Gaskell (in prep.): experimental design.....	76
Figure 6. Mirković and Gaskell (in prep.): generalisation results	77
Figure 7. Experiment 1: training stimuli set	81
Figure 8. Experiment 1: phonological distribution of novel affixes.....	82
Figure 9. Experiment 1: experimental design	84
Figure 10. Experiment 1: experimental procedure and tasks.....	88
Figure 11. Experiment 1: training accuracy	92
Figure 12. Experiment 1: test accuracy (picture naming).....	94
Figure 13. Experiment 1: test accuracy (2AFC)	94
Figure 14. Experiment 1: generalisation results	96
Figure 15. Experiment 1: generalisation results for equalised correct responses...	97
Figure 16. Experiment 1: generalisation errors	98
Figure 17. Representational Similarity Analysis (RSA)	107
Figure 18. PPI functional connectivity analysis.....	114
Figure 19. Experiment 2: training stimuli set	123
Figure 20. Experiment 2: experimental design	124
Figure 21. Experiment 2: fMRI scanning procedure	127
Figure 22. Experiment 2: training accuracy	136
Figure 23. Experiment 2: test accuracy (picture naming).....	137
Figure 24. Experiment 2: test accuracy (2AFC)	138
Figure 25. Experiment 2: generalisation results	140
Figure 26. Experiment 2: generalisation results for equalised correct responses.	140
Figure 27. Experiment 2: generalisation errors	141
Figure 28. Experiment 2: generalisation 3AFC results.....	143
Figure 29. Experiment 2: fMRI test trials accuracy	144
Figure 30. Experiment 2: univariate whole-brain analysis results	145
Figure 31. Experiment 2: univariate ROI analysis results	146

Figure 32. Experiment 2: RSA models RDMS (factorial design)	148
Figure 33. Experiment 2: RSA ROI analysis results (factorial design)	150
Figure 34. Experiment 2: RSA model RDMS (type frequency)	151
Figure 35. Experiment 2: RSA ROI analysis results (type frequency).....	152
Figure 36. Experiment 2: RSA model RDMS (phonological ambiguity)	153
Figure 37. Experiment 2: PPI seed-to-voxel analysis affix regularity x affix training day interaction	156
Figure 38. Experiment 2: PPI seed-to-voxel analysis main effect of phonological consistency	157
Figure 39. Experiment 2: PPI seed-to-voxel analysis main effect of affix regularity	158
Figure 40. Experiment 3: generalisation results Experiments 1 vs. 2.....	170
Table 1. Experiments 1 vs. 2 analysis summary	171
Figure 41. Experiment 3: training stimuli set	176
Figure 42. Experiment 3: training accuracy	179
Figure 43. Experiment 3: test accuracy (picture naming).....	180
Figure 44. Experiment 3: test accuracy (2AFC)	181
Figure 45. Experiment 3: generalisation results	183
Figure 46. Experiment 3: generalisation results for equalised correct responses.	184
Figure 47. Experiment 3: generalisation errors	185
Figure 48. Experiment 3: generalisation 3AFC results.....	186
Figure 49. Overnight consolidation of type and token frequency.....	200

List of tables

Table 1. Experiments 1 vs. 2 analysis summary 169

Acknowledgements

I would like to thank my supervisors, Matt Davis and John Williams, for their help and support over the past few years. I must express my special thanks to Matt for his patient guidance, encouragement, and advice. I have developed into a confident and independent researcher thanks to his mentorship. My sincerest thanks and appreciation go to Helen Blank for her invaluable help and friendship. I am also indebted to Jelena Mirković and Gareth Gaskell for introducing me to their exciting research, which has formed the basis of the work reported in this thesis, and for the many stimulating conversations. I have also benefited from the thoughtful insights of past and present members of the Davis Group.

I am thankful to Denise Klein, Xiaoqian Chai, and all other members of the Klein Lab for generously hosting me during my overseas study visit at the Montreal Neurological Institute and teaching me about the exciting possibilities afforded by functional connectivity analysis methods.

Many friends have supported and helped me along the way. Among these wonderful people, I am particularly grateful to Jean-Michel and Fernando for their love, care, and hospitality. A special thanks to Sophie, Noémie, Claudel, Sarah, and Daphnée for the time spent together in Montreal when I needed it most.

Most importantly, though, I am thankful to my family for encouraging me through all my endeavours. To my mother, Édith, for her kindness and strength. To my father, José, for his candour and resolve. I owe my grit to you both. To my brother, Thomas, for his sensitivity and compassion. To my brother, Jordi, for his determination and dependability. To my tieta, for her selflessness and generosity. To my iaia, for teaching me more than words can say. You have each played a role in making this journey possible, and, for that, I am filled with gratitude.

The greatest debt of thanks is owed to my partner, Mike, for his constant love and unwavering support. I know it was a fortunate stroke of serendipity that brought us together, but it feels like fate had its way. You nurture me and, every day, you bring me one step closer to everything.

Declaration

This dissertation is the result of my own work and includes nothing, which is the outcome of work done in collaboration except as specified in the text and Acknowledgements.

It is not substantially the same as any work that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution.

It does not exceed the 80 000 word limit (including footnotes, references, and appendices) for the Modern and Medieval Languages Degree Committee.

*“Il y a plus de courage que de talent
dans la plupart des réussites”*

Félix Leclerc

Chapter 1: Learning morphology

1.1. Learning rules or strengthening connections

Understanding the nature of language learning has been a central issue in linguistics and cognitive neuroscience for decades. Generative grammar (Chomsky, 1956, 1959) views language as the manipulation of symbols (e.g. words and phrases) according to rules (e.g. syntax). These rules are said to arise from a domain-specific language acquisition device programmed with innate knowledge of core linguistic principles. Generative grammar has thus tended to focus on formalising rules and establishing how they are learned, represented, and processed. As such, it allies with the broader symbolic approach to studying cognitive processes (Fodor, 1968, 1975; Newell, 1980; Newell & Simon, 1981; Pylyshyn, 1984). An alternative view comes from connectionism, or Parallel Distributed Processing (PDP), which eschews the need for symbols and rules and instead views language as best understood by considering the properties of artificial neural networks (Rosenblatt, 1958, 1962; Rumelhart, Hinton, & McClelland, 1986). Artificial neural networks, in which processing units are akin to neurons and their connections to synapses, are used to model different cognitive processes in a way that resembles the functioning of the brain. Learning involves the emergence and strengthening of connections between processing units over time on the basis of statistical contingencies in the environment. Knowledge comes to be represented as patterns of activity distributed across large numbers of units. Processing consists of using previously learned knowledge. On this view, language relies on the same domain-general learning mechanisms as other aspects of cognition.

English past tense inflection has spearheaded much of the debate between these opposing theoretical views. It has attracted contentious consideration due to its sharp distinction, within the same cognitive domain, between a highly regular pattern and a set of exceptions (Marslen-Wilson & Tyler, 1998, 2005; Pinker, 2001; McClelland & Patterson, 2002a, 2002b; Pinker & Ullman, 2002a, 2002b). The majority of English verbs form their past tense by adding the affix '*ed*' to an otherwise

unchanging stem (e.g. ‘walk – walked’)¹. Because it applies to almost all verbs and is typically extended to new verbs that enter the language (e.g. ‘Skype – Skyped’), the regular past tense has been considered a paradigmatic example of a rule-based mechanism. By contrast, there are about 180 irregular English verbs that do not follow the regular affixation pattern (e.g. ‘take – took’, ‘go – went’). Symbolic and connectionist accounts have mostly disagreed on whether the learning, representation, and processing of regulars is qualitatively different to that of irregulars.

According to symbolic accounts of the English past tense, regulars are composed by a rule that concatenates symbols for the verb stems and the past tense affix (e.g. [walk] + [ed]). Storage of past tenses is not required as they can be assembled online by rule. By contrast, irregular stems and past tenses are learned and stored as independent lexical entries. This account has yielded *dual-mechanism* theories according to which regulars and irregulars are subserved by distinct mechanisms (Chomsky & Halle, 1968; Halle & Mohanan, 1985; Lachter & Bever, 1988; Pinker & Prince, 1988, 1994; Pinker, 1991, 1999, 2001; Prasada & Pinker, 1993; Ullman et al., 1997; Ullman, 1999, 2001; Marcus, 2001; Pinker & Ullman, 2002b; Marslen-Wilson & Tyler, 2007). Early dual-mechanism theories considered irregulars to be learned by rote and stored in distinct slots of memory (Chomsky & Halle, 1968; Halle & Mohanan, 1985). Rote memorisation of this kind could account for the learning of suppletive forms (e.g. ‘go – went’, ‘be – was’) but was inconsistent with important characteristics of other irregulars. First, like regulars, most irregulars show some degree of overlap between their stems and past tenses. For example, the regular ‘bake – baked’ and the irregular ‘take – took’ both retain the onset and coda of their stem in their past tense form (Halle & Mohanan, 1985). Second, there are sub-regularities in the mapping between stems and past tenses among pools of irregulars such as the vowel change in ‘cling – clung’, ‘fling – flung’, and ‘sling – slung’ (Chomsky & Halle, 1968). Third, some irregulars are similar to regulars in that their past tense

¹ One of three allomorphs of the affix ‘-ed’ is applied depending on the final sound of the stem to which it attaches: /t/ as in ‘walk – walked’ if it is a voiceless consonant, /d/ as in ‘play – played’ if it is a voiced consonant or a vowel, and /əd/ as in ‘twist – twisted’ if it is an alveolar stop.

ends in /t/, one of the three regular ‘*ed*’ allomorphs (e.g. ‘*creep – crept*’, ‘*sleep – slept*’ like in ‘*crop – cropped*’, ‘*step – stepped*’). These have been referred to as weak irregulars or pseudo-regular irregulars (Pinker, 1999; Joanisse & Seidenberg, 2005). Thus, irregulars are not as idiosyncratic as they might first appear. Similarities in the structure of regulars and irregulars are problematic for a model that assumes rote learning and independent representation of irregular stems and past tenses.

To account for these sub-regularities among irregulars, connectionist accounts – or *single-mechanism* theories - of the English past tense do not posit the same categorical distinction between regulars and irregulars as dual-mechanism theories. Both regulars and irregulars are learned, represented, and processed in a single network of connections that computes stem – past tense transformations from distributed phonological representations² (Rumelhart & McClelland, 1986; Seidenberg & Bruck, 1990; Cottrell & Plunkett, 1991; Plunkett & Marchman, 1991, 1993, 1996; MacWhinney & Leinbach, 1991; Seidenberg, 1992; Daugherty & Seidenberg, 1992; Hare & Elman, 1992, 1995; Hare, Elman, & Daugherty, 1995; Marchman, Plunkett, & Goodman, 1997). For example, a network receives distributed phonological representations of regular and irregular verb stems as input (e.g. ‘*bake*’, ‘*sing*’) and computes distributed phonological representations of their past tenses as output (e.g. ‘*baked*’, ‘*sang*’) by spreading activation over the connections between the input and output units. Through repeated exposure to the stem – past tense pairs³, the network gradually learns stem – past tense mappings by adjusting the weights between its connections using learning algorithms. Thus, single-mechanism accounts make no assumptions about a structural difference between regular and irregular mappings. Later instantiations of dual-mechanism theories, which are epitomised by Pinker’s ‘Words and Rules’ theory (Pinker, 1991, 1999), maintain that regulars are composed by rule in a symbol-processing system, but concede that irregulars are learned, represented, and processed in an associative memory network. The rule is said to be

² Distributed phonological representations refer to the fact that a word like ‘*bake*’ is not represented as a single unit but is rather encoded by separate units that represent its phonological features.

³ All connectionist models of the English past tense have used supervised learning where stems and their associated past tense are presented together. Unsupervised learning, in which stems and their associated past tense would not need to be presented together, have not yet been investigated.

preferentially applied whilst the associative memory network can block its application if an appropriate irregular is identified.

Support for dual-mechanism theories dates back to the cognitive revolution of the 1960s during which a number of studies focusing on children's acquisition of the English past tense played an important role in establishing the view that regulars are rule-governed (Berko, 1958; Ervin & Miller, 1963; Ervin, 1964; Brown, 1973; Kuczaj, 1977). Native English speaking children typically begin by producing the correct past tense for a limited number of regulars (e.g. '*played*') and irregulars (e.g. '*went*') before moving on to a period where they produce overregularisation errors (e.g. '*goed*'). These overregularisation errors involve overextending the regular affix '*ed*' to irregulars (e.g. '*goed*') that had previously been correctly used (e.g. '*went*'). This break from correct use to overregularisation errors has been argued to imply the deduction of a rule governing the regulars. Once discovered, this rule is temporarily misapplied to irregulars. Children then learn to apply the rule to regulars only and memorise irregulars as whole forms. However, how children learn to distinguish between regulars and irregulars or what mechanism supports the deduction of a rule for regulars is not explained by this account.

Rumelhart and McClelland (1986) challenged dual-mechanism theories by showing that a network could successfully simulate key characteristics of children's acquisition of the English past tense without invoking the deduction of a rule for regulars. The network initially correctly generated the past tense of regular and irregular verb stems before moving on to producing overregularisation errors similar to those produced by children. Thus, a simple network could learn both regular and irregular mappings using a common set of weighted connections. It also displayed rule-like behaviour and produced overregularisation errors in the absence of an explicit representation of the regular rule or a separate mechanism for irregulars. Rumelhart and McClelland (1986) argued that these findings greatly undermine dual-mechanism accounts of children's acquisition of the English past tense. Rumelhart and McClelland's (1986) network spawned a great deal of criticism from proponents of dual-mechanism theories focusing mainly on the inappropriateness of its training regime, its poor generalisation performance, and its inadequacies of detail,

particularly in terms of the phonological representations of its inputs (Pinker & Prince, 1988; Lachter & Bever, 1988; Prasada & Pinker, 1993; Ling & Marinov, 1993). Several subsequent networks have, however, managed to address most of these criticisms (e.g. Plunkett & Marchman, 1991, 1993 but see Marcus, 1995 for a dissenting view).

As the debate between single- and dual-mechanism theories unfolded, it broadened its scope to consider some of the factors influencing the learning of regulars and irregulars. The influence of the frequency and phonological consistency of verbs has received the bulk of investigative attention. Because they embody a categorical distinction between regulars and irregulars, dual-mechanism theories predict that the two should almost always be dissociable. Since only irregulars are memorised, they should be sensitive to properties of associative memory including frequency and phonological consistency effects whilst regulars should not. By contrast, single-mechanism theories predict that both regulars and irregulars should be sensitive to frequency and phonological consistency effects. The next section focuses on network simulations, behavioural experiments, and diachronic data that have been used to pit these predictions against each other.

1.2. Influences on learning inflections

1.2.1. Frequency

Frequency effects are well-established at all levels of language processing, typically yielding advantages for more frequent relative to less frequent forms (Cattell, 1886; Howes & Solomon, 1951; Solomon & Postman, 1952; Broadbent, 1967; Morton, 1969; Forster & Chambers, 1973; Bybee & Hopper, 2001; Ellis, 2002). However, frequencies can be measured in different ways. In the context of inflectional morphology, *type frequency* refers to the number of words that adhere to an inflectional pattern (Ellis, 2009). For example, many more English verbs follow the regular ‘*ed*’ affixation pattern than the vowel change in ‘*sing-sang*’ to form their past tense. That is, regulars have higher type frequency than irregulars. *Token frequency*, in contrast, refers to how often individual inflected forms occur (Ellis, 2009). English irregular past tense forms ‘*had*’ and ‘*did*’ have higher token frequency than ‘*crept*’ or ‘*forbade*’ because they occur more often in the language. Although English irregulars have much lower type

frequency than regulars, they have amongst the highest token frequencies (Bybee & Slobin, 1982; Bybee, 1995). Thus, in English, the regular – irregular dichotomy typically implies a type – token frequency distinction such that regular verbs have high type frequency whereas irregular verbs have high token frequency. According to dual-mechanism theories, only irregulars should be sensitive to type and token frequency effects. By contrast, according to single-mechanism theories, both regulars and irregulars should be affected by these factors.

In two separate experiments, Prasada, Pinker, and Snyder (1990) and Seidenberg and Bruck (1990) presented native English speakers with the stems of regular and irregular verbs and asked them to produce their past tenses as quickly as possible. Participants were faster at naming the past tense of high token frequency irregulars (e.g. *'took'*) than low token frequency irregulars (e.g. *'rang'*), even when stem frequencies were equated. By contrast, participants produced the past tense of high token frequency regulars (e.g. *'liked'*) and low token frequency regulars (e.g. *'biked'*) with similar latencies. For proponents of dual-mechanism theories, this frequency by regularity interaction is consistent with a rule governing the regulars. Efficiency of retrieval from the mental lexicon is modulated by token frequency such that high token frequency irregulars are retrieved faster than low token frequency irregulars. However, since regulars are computed online by rule, they do not display token frequency effects. Once the stem is retrieved, the rule takes the same amount of time to be applied to all regulars.

The frequency by regularity interaction is not uniquely compatible with dual-mechanism theories. Daugherty and Seidenberg (1992) trained a network on regular and irregular mappings. To examine the frequency by regularity interaction, Daugherty and Seidenberg (1992) constructed sets of the 10 highest token frequency regulars, 10 lowest token frequency regulars, 10 highest token frequency irregulars, and 10 lowest token frequency irregulars from an original training set. The model was more accurate on high token frequency irregulars than on low token frequency irregulars. By contrast, it was as accurate on high token frequency regulars as it was on low token frequency regulars. These results show that the frequency by regularity interaction reported by Prasada et al. (1990) and Seidenberg and Bruck (1990) can be

replicated within a single network undermining the claim that the interaction is evidence of a separate rule for regulars. Daugherty and Seidenberg (1992) further speculate that the frequency by regularity interaction could reflect not only token frequency effects but also type frequency effects. Learning the regular affixation pattern may benefit from several verbs sharing the same affix (high type frequency) such that token frequency effects are superseded.

Ellis and Schmidt (1998) considered the frequency by regularity interaction in adults learning regular and irregular inflections in an artificial language. Half of the words had a high type frequency regular prefix (e.g. *bugarth* but also *bupid*, *bulant*) and the other half had low type frequency irregular prefixes (e.g. *zecharp* but also *niwoop*, *rekag*). Within the regulars and the irregulars, half of the words were trained with high token frequency and half were trained with low token frequency. Participants were faster and more accurate at naming high token frequency regulars and irregulars at early stages in training. However, as training progressed, the token frequency effect disappeared for regulars but remained for irregulars giving rise to the frequency by regularity interaction reported by Prasada et al. (1990) and Seidenberg and Bruck (1990). Ellis and Schmidt (1998) further showed that a network could simulate the participant data. Like Daugherty and Seidenberg (1992), Ellis and Schmidt (1998) show that the frequency by regularity interaction can be explained by a single mechanism. The authors further argue that it is not the case that there are no token frequency effects for regulars but rather that these effects are difficult to observe due to asymptotic performance potentially driven by high type frequency. In addition, token frequency effects have been reported for the processing of regulars (Stemberger & MacWhinney, 1986; Alegre & Gordon, 1999; Hare, Ford, & Marslen-Wilson, 2001). For example, Alegre and Gordon (1999) reported the results of a lexical decision task for regular English past tenses. Higher token frequency regulars were responded to faster than low token frequency regulars. These results cannot be accounted for if regular past tenses are derived by rule rather than being stored in an associative memory network.

Proponents of dual-mechanism theories have argued that token frequency effects on irregulars are consistent with patterns of diachronic change (e.g. Pinker,

1991, 1999). For example, low token frequency irregulars are more likely to be regularised over time (e.g. ‘weep – weeped’) than high token frequency irregulars (e.g. ‘keep – kept’). In addition, all surviving suppletive forms (e.g. ‘go – went’) have high token frequencies (Bybee, 1985). However, these patterns of diachronic change do not necessitate postulating separate processing systems for regulars and irregulars. In the usage-based model (Bybee, 1985, 1988, 1995, 2001), regular and irregular inflectional patterns emerge from a single network of connections among lexical representations of past tenses. A high number of specific usage events with a word (i.e. high token frequency) strengthens its lexical representation. A stronger lexical representation makes a word more autonomous, easier to access whole, and less likely to be regularised over time (Bybee & Slobin, 1982; Bybee, 1985; Hooper, 1976). The regularisation and loss of low token frequency forms over time are argued to be consequences of insufficient usage. Speakers are unable to retrieve and use weakly represented variants. Furthermore, Hare and Elman (1995) report the results of a network simulating diachronic change of past tense inflection from Old to Modern English, which largely agrees with the usage-based model. The output of one network was used as the input of another to simulate change over generations of language users. High token frequency mappings were learned more quickly and accurately and were more likely to be preserved over generations. Low token frequency mappings, by contrast, were harder to learn and more prone to regularisation over generations.

As with token frequency, dual-mechanism theories predict that type frequency should have little influence on the learning and processing of regulars. Single-mechanism theories, by contrast, consider type frequency to be a determining factor of the productivity of the regular affixation pattern. Productivity refers to the extent to which an affixation pattern is used to inflect novel forms. As such, it is closely related to the notions of *default* inflection and *generalisation*, which are discussed in more detail in Chapter 3. Nonetheless, it is worth noting here that connectionist networks are able to simultaneously learn specific mappings and regularities that are shared across several mappings. Plunkett and Marchman (1993) report the results of a network simulation of past tense inflection in which the size of the training set was increased incrementally. The ability of the network to correctly acquire new regular verbs was poor as the size of the training set increased. However, after the training set

reached a critical mass (around 50 verbs), the network's performance started to improve rapidly. Plunkett and Marchman (1993) argue that these findings suggest that the network shifted from memorising specific mappings to discovering regularities shared across mappings allowing it to treat new verbs entering the training set in a systematic way. Plunkett & Juola (1999) found a similar critical mass effect in a connectionist model trained to produce the plurals of English nouns. Thus, the ability of a single network to extract regularities shared by several items (much like the deduction of a rule) can be simulated without invoking a symbol-processing system for regulars. The extent to which these results are an artefact of the idiosyncratic frequency statistics of English, however, has been a matter of debate, which is considered in Chapter 3.

1.2.2. Phonological consistency

Regular and irregular English verbs can be described in terms of their phonological consistency. Phonological consistency refers to the extent to which a verb shares with other verbs the phonology of its stem and the type of transformation that associates its stem with its past tense (Pinker & Prince, 1988; McClelland & Patterson, 2002). For example, the regular verb '*walked*' can be said to be phonologically consistent because its stem is similar to its phonological neighbours '*balk*', '*calk*', '*chalk*', '*stalk*', and '*talk*', which all have regular past tenses. Similarly, the irregular verb '*slung*' is relatively consistent since its stem is similar to its phonological neighbours '*cling*', '*fling*', '*sting*', '*string*', '*swing*', and '*wring*', which all follow the same vowel alternation pattern to form their past tense. By contrast, the regular verb '*flowed*' can be said to be phonologically ambiguous. Its stem is similar to its phonological neighbours '*mow*', '*show*', '*slow*', and '*tow*', which have regular past tenses but it is also similar to its phonological neighbours '*blow*', '*grow*', '*know*', and '*throw*', which have irregular past tenses. According to dual-mechanism theories, at least in their original formulations, regulars are derived via rule and thus would not be expected to display phonological consistency effects. Single-mechanisms theories, by contrast, directly predict that both regulars and irregulars should demonstrate phonological consistency effects.

Seidenberg and Bruck (1990; also discussed in Seidenberg, 1992) asked native English speakers to generate the past tense of regular verb stems varying in their

degree of phonological consistency. Participants were slower and less accurate at producing the past tense of phonologically ambiguous regular verb stems compared to phonologically consistent regular verb stems, whilst naming latencies for the verb stems did not vary. Furthermore, for the phonologically ambiguous regular stems, the naming latency effect was related to the number of irregular phonological neighbours. Participants were faster at naming the past tense of regular stems like '*pick - picked*', which only has one phonological neighbour with an irregular past tense, '*stick - stuck*', than regular stems like '*blink - blinked*', which has several phonological neighbours with irregular past tenses such as '*sink - sank*' and '*think - thought*'. Seidenberg and Bruck (1990) argue that phonological consistency effects among regulars are at odds with the predictions made by dual-mechanism theories and support single-mechanism theories.

Ullman (1999a) reported on an experiment in which participants made acceptability ratings for regular and irregular English past tenses varying in their phonological consistency. Acceptability ratings for irregulars were positively correlated with their phonological consistency. In contrast, there was no correlation between acceptability ratings and phonological consistency for regulars. Ullman (1999a) argue that the presence of phonological consistency effects for irregulars but not for regulars is consistent with the predictions of dual-mechanism theories. However, it should be noted that Ullman (1999a) only used regular past tenses whose stems are phonologically dissimilar to the stems of irregulars (e.g. '*walk - walked*') (i.e. phonologically consistent). Seidenberg and Bruck (1990) considered both phonologically consistent and phonologically ambiguous regulars. Ullman (1999a) remark that phonologically ambiguous regulars (e.g. '*glide - glided*'; cf. '*ride - rode*', '*hide - hid*') were tested but not analysed. The conclusions that can be drawn from the results of this study are limited by the exclusion of these phonologically ambiguous regulars. The results may simply reflect the fact that irregulars were more varied in their phonological consistency than regulars.

Daugherty and Seidenberg (1992) trained a network to map the phonological representation of verb stems to their past tense. The training corpus included 60 phonologically consistent regular verb stems (e.g. '*walk - walked*', cf. '*talk - talked*',

'stalk – stalked'), 60 phonologically ambiguous regular verb stems (e.g. 'bake – baked', cf. 'take – took', 'make – made'), and 60 irregular verb stems, all equated in terms of token frequency. The network performed better on the phonologically consistent regular verb stems than on phonologically ambiguous regular verb stems, which in turn, both yielded better performance than irregular verb stems. Daugherty and Seidenberg (1992) argue that these results are consistent with the behavioural results reported by Seidenberg and Bruck (1990) and support a single-mechanism account. A dual-mechanism account has regulars and irregulars being learned and processed by two separate mechanisms. Hence, there is no basis for predicting that they will interfere with each other. A single-mechanism account, however, encodes both regular and irregular past tense in the same set of connection weights. Hence, the learning and processing of a regular form is affected by whether it has an irregular phonological neighbour or not. It should be noted that in order to account for the previously outlined evidence, later formulations of dual-mechanism theories (Ullman et al., 1997, 2005; Pinker, 1999; Pinker & Ullman, 2002b) have suggested that the past tense forms of phonologically ambiguous regulars are stored in the same associative memory network as irregulars. As a result, phonologically ambiguous regulars are argued to be processed more slowly and subject to the same blocking mechanism as irregular forms. With this amendment, dual-mechanism theories are in principle able to account for the previously outlined evidence. However, single-mechanism theories provide a more parsimonious account of the evidence.

1.3. Neural bases of inflectional processing

A single network of distributed representations appears to provide a more parsimonious and psychologically plausible account of children's developmental profile in acquiring past tense inflections as well as the influence of frequency measures and phonological consistency. However, the debate between single- and dual-mechanism theories continued well into the 1990s before reaching somewhat of a stalemate as far as the behaviourally observable properties of the phenomenon are concerned. To distinguish between these opposing views, research shifted its focus towards characterising the neural mechanisms underpinning adult processing of regular and irregular English past tenses in both health and disease. Dual-mechanism

theories postulate that since regulars and irregulars are subserved by distinct mechanisms, it should be possible to identify separate neural subsystems for their processing. By contrast, single-mechanism theories argue that there should be no such dissociation.

Several neuroimaging studies have reported different patterns of brain activation associated with the generation of regular and irregular past tenses in healthy adults using positron emission tomography (PET; Jaeger et al., 1996; Indefrey et al., 1997), functional magnetic resonance imaging (fMRI; Ullman, Bergida, & O'Craven, 1997; Beretta et al., 2003), magnetoencephalography (MEG; Dhond, Marinkovic, Dale, Witzel, & Halgren, 2003), and electroencephalography (EEG; Lavric, Pizzagalli, Forstmeier, & Rippon, 2001). These studies share the assumption that differences in brain activity patterns associated with generating regulars and irregulars provide evidence for separate cognitive and neural subsystems underlying their processing. For example, Jaeger et al. (1996) found participants to show increased activation in the left dorsolateral prefrontal cortex when generating regulars and increased activation in the left middle temporal gyrus (MTG) when generating irregulars. Jaeger et al. (1996) interpreted these results as supporting dual-mechanism theories and suggested that whilst the dorsolateral prefrontal cortex may support rule application, the left MTG may be involved in retrieving stored irregulars.

Neuropsychological studies showing double dissociations in the processing of regulars and irregulars in patient groups have also been taken to support dual-mechanism theories. Across different behavioural tasks, patients with frontal damage as a consequence of stroke (Marslen-Wilson & Tyler, 1997; Tyler, DeMornay-Davies, et al., 2002) or Parkinson's disease (PD) (Ullman et al., 1997) have shown impaired processing of regulars but relatively well-preserved processing of irregulars. Patients with damage to temporal regions, as a result of either stroke (Marslen-Wilson & Tyler, 1997), semantic dementia (Patterson, Lambon Ralph, Hodges, & McClelland, 2001), Herpes Simplex Encephalitis (Tyler et al., 2002), or Alzheimer's disease (AD) (Ullman et al., 1997) have shown the opposite pattern of impairment. For example, Ullman et al. (1997) found that patients with AD had more difficulty generating the past tense of irregular than regular verbs. Patients with PD exhibited the converse pattern of

impairment. Ullman et al. (1997) suggested that this dissociation in behavioural deficits and their association with neurodegeneration in different brain regions provides strong support for distinct cognitive and neural subsystems underlying the processes of regulars and irregulars.

Taking these results as a starting point, Ullman and colleagues proposed the Declarative/Procedural (DP) model of language (Ullman, Corkin, et al., 1997; Ullman, 1998, 2001, 2004, 2006). The DP model postulates that different aspects of language processing can be related to the neurobiological distinction between the declarative and procedural memory systems (Mishkin, Malamut, & Bachevalier, 1984; Schacter & Tulving, 1994; Squire & Knowlton, 2000; Eichenbaum & Cohen, 2001). The ‘mental lexicon’, which contains memorised words such as irregular past tenses, is said to depend on the temporal lobe structures that subserve declarative memory. The ‘mental grammar’, which underpins the rule-based computation of complex forms (e.g. ‘*jump*’ + ‘*ed*’), is argued to rely on the frontal/basal-ganglia circuits that support procedural memory. AD affects temporal lobe structures whilst leaving frontal/basal-ganglia circuits relatively intact. By contrast, PD affects the basal ganglia and the frontal areas⁴ to which it projects whilst sparing temporal lobe structures. Thus, Ullman et al. (1997) suggested that these specific impairments in declarative and procedural memory could account for the double dissociation observed in patients with AD and PD processing regulars and irregulars.

Based on their own neuropsychological and neuroimaging studies, Marslen-Wilson, Tyler, and colleagues proposed the morpho-phonological decomposition model which, like the DP model, allies with dual-mechanism theories (Marslen-Wilson & Tyler, 1997, 1998, 2003, 2007; Tyler, DeMornay-Davies, et al., 2002; Tyler, Randall, & Marslen-Wilson, 2002; Tyler, Marslen-Wilson, & Stamatakis, 2005; Tyler, Stamatakis, Post, Randall, & Marslen-Wilson, 2005). The model proposes a distinction between lexical access processes that require morpho-phonological decomposition of complex forms, such as regular past tenses (e.g. ‘*jump*’ + ‘*ed*’), from

⁴ As well as frontal areas, degeneration of the basal ganglia also inhibits motor areas to which it projects to causing the tremors and muscle rigidity that are characteristic of PD.

those that rely on more direct access to whole forms, such as irregular past tenses (e.g. *'found'*), which have no internal morpho-phonological structure. For example, Marslen-Wilson and Tyler (1997) used auditory-auditory repetition priming to show selective deficits in morpho-phonological decomposition. Two aphasic patients with left-hemisphere lesions affecting parts of the frontal lobe, including the inferior frontal gyrus (IFG), and most of the temporal lobe following stroke showed normal priming for irregulars (e.g. *'found'* priming *'find'*) but reduced priming for regulars (e.g. *'jumped'* priming *'jump'*).

Marslen-Wilson and Tyler (2007) argue that whilst irregulars are equivalent to regulars in terms of their syntactic and semantic properties, they do not require the same decomposition into their constituent stem and affix. Instead, irregulars must be learned, stored, and accessed as whole forms. Preserved priming for irregulars is claimed to demonstrate that the impairment for regulars is specific to their decompositional properties. Furthermore, Marslen-Wilson and Tyler (2007) postulate a decompositional network linking left inferior frontal regions with temporal regions. The DP model (Ullman, 1997, 2001, 2004, 2005, 2006) distinguishes between a memorised mental lexicon and a computational mental grammar, relying on the declarative and procedural memory systems, respectively. The morpho-phonological decomposition model (Marslen-Wilson & Tyler, 2007) instead proposes a fronto-temporal network involved in morpho-phonological decomposition that cuts across these two memory systems.

Both the DP and the morpho-phonological decomposition models have attracted criticism from proponents of single-mechanism models. Arguments have developed on the grounds that differences in the processing of regulars and irregulars reported in neuroimaging and neuropsychological studies do not necessarily reflect distinct neural systems (Joanisse & Seidenberg, 1999; Juola & Plunkett, 2000; McClelland & Patterson, 2002c; Plaut, 1995; Seidenberg & Hoeffner, 1998). One alternative explanation is that differences between regulars and irregulars simply

reflect differences in their processing difficulty⁵ (Seidenberg & Arnoldussen, 2003; Seidenberg & Hoeffner, 1998). For example, Jaeger et al. (1996) used a blocked design such that regulars and irregulars were presented separately, introducing potential experimental confounds. Regulars were more homogeneous and predictable (all formed by affixing ‘*ed*’) than irregulars, which were more varied and less predictable (formed in several different ways). Thus, irregulars could elicit more activity in left temporal regions than regulars because they are less predictable and more difficult to process than regulars not because they rely on different neural systems. This account is also consistent with the latency data reported by Jaeger et al. (1996). Participants were much faster at producing regular past tenses than they were at producing irregular past tenses.

A number of neuroimaging studies have found no differences in the patterns of brain activity elicited by regulars and irregulars, particularly when controlling for factors such as the phonological complexity of the past tenses. These findings have been taken as supporting single-mechanism theories (Desai, Conant, Waldron, & Binder, 2006; Joanisse & Seidenberg, 2005; Sach, Seitz, & Indefrey, 2004). For example, Joanisse and Seidenberg (2005) compared patterns of brain activity elicited by regulars (e.g. ‘*step – stepped*’), true irregulars (e.g. ‘*take – took*’), and pseudo-regular irregulars (e.g. ‘*sleep – slept*’) in an event-related fMRI study. Regulars and irregulars (combining true irregulars and pseudo-regular irregulars) elicited similar patterns of activity in bilateral posterior temporal lobe. By contrast, regulars elicited greater activity in left IFG than irregulars. Previous studies have interpreted similar results as supporting dual-mechanism theories (Pinker & Ullman, 2002). In a follow-up analysis, however, pseudo-regular irregulars were found to elicit similar levels of activation to regulars in the left IFG, which both elicited more activity than true irregulars. Importantly, naming latencies and accuracies did not differ between the pseudo-regular irregulars and the true irregulars so that differences in brain activity pattern could not be attributed to differences in processing difficulty. Thus, pseudo-

⁵ Seidenberg and Arnoldussen (2003) argue that this problem is not restricted to neuroimaging studies of the English past tense but reflects a broader issue in neuroimaging studies which is that differences between experimental conditions are interpretable with reference to the variables of interest (e.g. type of stimuli) only if the stimuli are equated across conditions in terms of processing difficulty, a problem they coined the *Difficulty Matching Error*.

regular irregulars pattern with regulars in terms of eliciting activity in brain regions associated with phonological processing such as the left IFG. As such, there is no strict dichotomy between regulars and irregulars as suggested by single-mechanism theories. Consistent with Joanisse and Seidenberg's (2005) findings, Desai et al. (2006) found that when contrasting regular and irregular past tense generation directly in an fMRI study, no areas showed greater activation for regulars if the regular and irregular past tenses were matched on phonological complexity.

Proponents of dual-mechanism theories have held double dissociations in the processing of regulars and irregulars observed in neuropsychological studies as the crowning piece of evidence for separate neural systems. However, these double dissociations have also been claimed to be compatible with single-mechanism accounts on the basis of network simulations (Plaut, 1995, Juola & Plunkett, 1998, Joanisse & Seidenberg, 1999) and neuropsychological studies (Patterson et al., 2001; Bird, Lambon Ralph, Seidenberg, McClelland, & Patterson, 2003). For example, Joanisse and Seidenberg (1999) noted that the regular past tenses used by Ullman et al. (1997) to show a double dissociation between AD and PD patients appeared to be phonologically more complex than the irregulars, such that any specific impairment with regulars could reflect a more general deficit in phonological processing. PD has been associated with phonological deficits whilst AD has been associated with semantic deficits (Grossman, Carvell, Stern, Gollomp, & Hurtig, 1992; Hodges & Patterson, 1995). Joanisse and Seidenberg (1999) implemented a neural network in which verbs were represented in terms of their phonology and semantics. Once trained, lesioning the semantic units led to specific impairments in the production of irregular past tenses, while lesioning the phonological units led to a specific impairment in the production of regulars. Joanisse and Seidenberg (1999) argue that these results show how knowledge of past tense inflection is represented in a single network that encodes semantic, phonological, as well as other types of information about words. The behavioural dissociations reported by Ullman et al. (1997) and others can reflect phonological and semantic representations being differentially affected by brain damage. Furthermore, in the neuropsychological assessment of 10 aphasic patients, Bird et al. (2003) found that the disadvantage for the production of

regulars disappeared when phonological complexity was controlled consistent with a general impairment in phonological processing.

1.4. Cross-linguistic inflectional processing

The debate between single- and dual-mechanism theories of inflectional morphology, as well as the evidence reviewed thus far, has mostly focused on the English past tense. However, English has an impoverished inflectional system and lacks the complexity of other inflectional systems including, for example, those of Romance and Slavic languages. Furthermore, regular and irregular English verbs are strikingly different in that only regulars undergo affixation (e.g. *walk* – *walked*) whereas irregulars either undergo vowel changes (e.g. *take* – *took*) or suppletion (e.g. *go* – *went*). As such, the consideration of cross-linguistic inflectional processing may provide a stronger test of single- versus dual-mechanism accounts of morphology. An exhaustive review of cross-linguistic inflectional processing is beyond the scope of this thesis. Nevertheless, some illustrative examples are considered below.

Italian, Russian, Norwegian, and Icelandic all have complex verbal conjugational paradigms that have more than one regular verb class. A number of studies have considered the processing of verbal morphology in these languages in adults and children as well as in second language learners (Chernigovskaya & Gor, 2000; Gor & Chernigovskaya, 2003a, 2003b; Orsolini & Marslen-Wilson, 1997; Orsolini, Fanari, & Bowles, 1998; Ragnasdottir, Simonsen, & Plunkett, 1997; Simonsen, 2000). Across these languages, type frequency has been found to be a determining factor of the generalisation of conjugational patterns in all verb classes including regular verb classes. As reviewed earlier in this chapter, only irregulars should be sensitive to type frequency effects according to dual-mechanisms theories. Single-mechanism theories, by contrast, consider type frequency to be a determining factor of the productivity of an affixation pattern. Thus, cross-linguistic behavioural data does not appear to support a categorical distinction between regular and irregular morphological processing as proposed by dual-mechanism theories.

Mirković, Seidenberg, and Joanisse (2011) analysed a large corpus of Serbian nouns and showed that, like English, Serbian inflectional morphology has a number of

partial regularities that vary in type frequency and phonological consistency. Mirković et al. (2011) then trained a network on over 3000 Serbian nouns to produce correctly inflected phonological forms from a specification of a word's lemma, gender, number and case. The model's ability to learn the correct inflected forms was influenced by type frequency and phonological consistency. Mirković et al. (2011) argue that these results suggest that the correct production of inflected forms in Serbian, as in English, reflects the interplay of competing factors within a single network of distributed representations. They further argue that their results suggest that common computational mechanisms may govern the representation of inflectional morphology across languages.

Turning to neuroimaging evidence for cross-linguistic inflectional processing, de Diego Balaguer et al. (2006) compared brain activations associated with the generation of regular and irregular Spanish verbs. In Spanish, both regulars (e.g. *cantar* – *canto* [I sing]) and irregulars (e.g. *sentir* – *siento* [I feel]) undergo suffixation, which arguably offers better control over the potential differences in difficulty that may account for contrasting patterns of neural activity associated with the processing of regular and irregular English verbs discussed in the previous section. De Diego Balaguer et al. (2006) found that both regulars and irregulars activated the left IFG, a region that dual-mechanism theories would expect only regulars to recruit for the retrieval of grammatical features and affixation. Regulars elicited greater activity than irregulars in the anterior STG, the hippocampus, and the insula, which de Diego Balaguer et al. (2006) argue may reflect automatic stem reactivation and maintenance when inflecting regulars. By contrast, irregulars elicited greater activity than regulars in the bilateral IFG and MFG regions, which have been implicated in supporting memory retrieval, manipulation of information, and selection of correct responses. De Diego Balaguer et al. (2006) argue that the overlapping activations between regulars and irregulars are more consistent with single-mechanism theories whereas the distinct activations could be taken to support dual-mechanism theories. However, unlike previous neuroimaging evidence supporting dual-mechanism theories, de Diego Balaguer et al.'s (2006) results suggest differences in the processing of regulars and irregulars are mostly supported by prefrontal areas.

1.5. Summary

The debate surrounding the learning, representation, and processing of the English past tense inscribes itself in a broader debate regarding the nature of mental computations. Dual-mechanism theories argue that regulars and irregulars rely on distinct cognitive and neural subsystems, whilst single-mechanism theories contend that they are subserved by the same architecture. The evidence reviewed here suggests that language users' ability to learn and process inflections is best seen as a complex interplay of competing influences within a single network of distributed representations. Learning and processing are dependent upon the frequency and phonological consistency of inflectional patterns. However, most of the evidence for the influence of these factors comes from child acquisition data, network simulations, and diachronic change data. Behavioural experiments examining these issues in adults are scarce. Furthermore, whilst a large body of neuroimaging studies has investigated the processing of inflections in adults, a paucity of work has considered the neural mechanisms supporting their learning. Investigating the mechanisms that support the learning of inflections requires the consideration of domain-general learning and memory mechanisms. Following initial learning, new memories undergo subsequent reorganisation through consolidation processes. A fairly recent but growing body of work suggests that the development of language representations benefits from periods of consolidation. The next chapter reviews the literature on memory consolidation and its role in language learning.

Chapter 2: Consolidating morphology

2.1. Origins of consolidation theory

Memory consolidation refers to the gradual stabilisation of new memories in long-term memory and the processes that underpin this stabilisation (Dudai, 2002, 2004; Dudai, Karni, & Born, 2015). Reference to the idea of memory consolidation can be traced back to the writings of the Roman rhetorician Quintilian (c. AD 95) who noted the “*curious fact...that the interval of a single night will greatly increase the strength of the memory*” such that “*things which could not be recalled on the spot are easily coordinated the next day*” raising the possibility that memory “*undergoes a process of ripening and maturing during the time that intervenes*”. In the late 19th century, the process of consolidation was described by Ribot (1882) based on his clinical observations of amnesic patients for whom recently acquired premorbid memories are typically more impaired than more distant premorbid memories. This phenomenon is epitomised in Ribot’s Law, which states that the “*progressive destruction advances progressively from the unstable to the stable*” (Ribot, 1882). Hence, Ribot conceptualised a time-dependent process of memory organisation such that recent (unstable) memories are more vulnerable to disruption following trauma or injury than more distant (stable) memories.

The term *consolidation* was coined by Müller and Pilzecker (1900) based on a series of behavioural experiments which suggested that memory requires time to become *fixed* or undergo *Konsolidierung*. Participants learned nonword paired associates and were subsequently tested on the recall of the appropriate nonwords when cued with their pair members. When a new list of paired associates was learned in the first few minutes following learning of the initial list, recall was impaired compared to when the training-retrieval interval was of the same duration but did not include new learning, a phenomenon Müller and Pilzecker (1900) termed *retroactive interference*. No retroactive interference was observed when a gap of a few minutes was introduced between the learning of the two lists. Müller and Pilzecker (1900) suggested that their findings reflected a post-training interval during which associations are consolidated in memory and that introducing interfering material during this interval could disrupt the consolidation process. Jenkins and Dallenbach

(1924) further suggested a potential role for sleep in preventing retroactive interference after observing that participants recalled more nonsense syllables when they slept between presentations and recall than when they remained awake.

The consolidation process was further characterised by Burnham (1903) in a seminal paper integrating findings from experimental psychology and clinical observations of amnesic patients in which he noted that “*considerable time may be necessary*” for a memory to be consolidated and that consolidation is not “*merely a process of making a permanent impression upon the nerve cells, but also a process of association, of organisation of the new impressions with the old ones*”. Burnham’s (1903) characterisation added two important features to the description of the consolidation process. First, his conceptualisation of time referred to two different types of consolidation kinetics; fast, such as shown in the studies of Müller and Pilzecker and slow, as revealed by patterns of impairment in amnesia. Second, he suggested that consolidation is not merely a process through which new memories are strengthened but a process through which they are reorganised and integrated with old memories.

These two features have played an important role in characterising the neurobiological foundations of memory consolidation. A long-standing distinction in learning and memory research invokes separate neurobiological mechanisms for fast and slow consolidation. The former is accomplished within the first minutes to hours after the encoding of novel memories. Donald Hebb’s (1949) seminal work suggested the first neurobiological mechanism for fast consolidation. He proposed post-encoding neuronal reverberation as a means to maintain labile new memories. Through reverberation, neuronal activity persists after the encoding has occurred or practice has ended, such as to allow the strengthening of synapses (i.e. long-term potentiation) among the neuronal ensembles that encoded the new memory. Consequently, fast consolidation is commonly referred to as *synaptic consolidation*

(Frey & Morris, 1997; Sajikumar, Navakkode, Sacktor, & Frey, 2005; Reymann & Frey, 2007; Redondo & Morris, 2011; Clopath, 2012)⁶.

In contrast to fast consolidation, slow consolidation occurs over the course of days, weeks, or even years. In two related papers, Marr (1970, 1971) proposed a neurobiological theory of learning and memory that accounted for slow consolidation. He described how the archicortex, a structure now known as the hippocampus and hippocampal formation, functions as a *simple memory* storing information derived from sensory experience. Simple memory is suggested to perform a direct and temporary memorising function without performing any complex classification or organisation computations. It provides the input to the neocortex, which slowly classifies and organises new information in long-term memory. Marr (1970) further suggested a role for sleep in facilitating the transfer of information between the archicortex and the neocortex. This process is commonly referred to as *systems consolidation* as it involves the reorganisation over time of the brain systems that represent memories. Memory representations spread to neocortical structures and eventually relinquish their dependence from the archicortex, which supported their initial encoding. Thus, Marr's proposal echoes the early ideas of Burnham in that consolidation does not merely entail the strengthening of new memories but also their reorganisation and integration with existing memories.

Since Quintilian noted it, Ribot conceptualised it, and Müller and Pilzecker coined it, memory consolidation has been further characterised by Burnham and its potential underlying neurobiological mechanisms proposed by Hebb and Marr. Marr's ideas, in particular, have made an epochal contribution to contemporary theories of systems consolidation. In the spirit of Marr's original ideas, various modern theories postulate two dissociable neurobiological mechanisms for systems consolidation; a hippocampal system for the initial encoding of new memories and a neocortical system for the long-term storage of memories and their integration with existing memories (Squire, 1992; Alvarez & Squire, 1994; McClelland, McNaughton,

⁶ There is also evidence that fast consolidation is dependent on interactions between synapses and their cell body and nucleus (Martin et al., 1997, Casadio et al., 1999, Dudai & Morris, 2000, Huber, Kayser, & Bear, 2000, Adams & Dudek, 2005) and so the terms "cellular" and "local" consolidation are sometimes used.

& O'Reilly, 1995; Nadel & Moscovitch, 1997; O'Reilly & Norman, 2002; Winocur & Moscovitch, 2011; McClelland, 2013; O'Reilly, Bhattacharyya, Howard, & Ketz, 2014; Kumaran, Hassabis, & McClelland, 2016). One of these theories, the Complementary Learning Systems (CLS) model (McClelland et al., 1995; O'Reilly & Norman, 2002; McClelland, 2013; O'Reilly et al., 2014; Kumaran et al., 2016), offers perhaps the best worked out neuroanatomical and computational description of systems consolidation and is the focus of the next section.

2.2. Consolidation and Complementary Learning Systems

The CLS model shares similarities with other accounts of memory and learning that invoke dissociable neurobiological systems specialised for different types of memory (e.g. Marr, 1970, 1971; Squire, 1992). However, whereas other accounts are mostly descriptive and rather undetermined in computational terms, the CLS model was innovative and influential in applying connectionist modelling to consider the types of computations underlying each system and how systems consolidation emerges from these computations. It also sought to relate these computations to the neurobiological and functional organisation of memory.

Like other theories of systems consolidation, the CLS model is grounded in the striking observation, in both human and animals, that damage to the hippocampus and related medial temporal lobe (MTL) structures causes dramatic amnesia. Many amnesics display temporally graded retrograde amnesia, in which recently acquired premorbid memories are more impaired than older ones, consistent with Ribot's Law (Scoville & Milner, 1957; Marslen-Wilson & Teuber, 1975; Squire & Alvarez, 1995; Rempel-Clower, Zola, Squire, & Amaral, 1996; Squire & Bayley, 2007). Using a multitude of memory tasks, temporally graded retrograde amnesia has been reported in mice, rats, rabbits, and monkeys following selective lesions to the hippocampus and related MTL structures (Zola-Morgan & Squire, 1990; Winocur, 1990; Kim & Fanselow, 1992; Cho, Beracochea, & Jaffard, 1993; Kim, Clark, & Thompson, 1995; Kubie, Sutherland, & Muller, 1999; Clark, Broadbent, Zola, & Squire, 2002). Distant autobiographical memories, however, tend to be well-preserved in amnesics with hippocampal or MTL lesions (Bayley, Hopkins, & Squire, 2003). In contrast, flat extensive retrograde amnesia typically involves neocortical damage in the lateral and

anterior temporal lobes (Squire, Clark, & Knowlton, 2001). Together, these observations suggest that the hippocampus and related MTL structures are required for intact long-term declarative memory but that their role is circumscribed in time such that memories are eventually stored and represented in the neocortex.

Immediate memory span for sequences of digits and unrelated pairs of words often remains within the normal range in amnesic patients (Scoville & Milner, 1957). Together with the retention of old memories, these observations have been suggested to imply that deficits are not in the acquisition or short-term retention of novel information but rather in its consolidation or retrieval. Temporally graded retrograde amnesia suggests that the impairment lies in the consolidation of new memories (Dudai, 2004). In fact, fMRI studies have shown temporally graded activity in MTL structures in the human brain (Haist, Bowden, & Mao, 2001) and in the hippocampal formation of mice (Bontempi, Laurent-Demir, Destrade, & Jaffard, 1999) consistent with the idea that long-term memory relies on the time-dependent reorganisation of brain circuitry.

The CLS model elaborates an account of the functional organisation of memory that is consistent with the neuropsychological evidence described above. The account postulates a hippocampal and a neocortical system with distinct computational functions and representational structures (Figure 1). The first postulate of the CLS model is that the neocortical system houses structured knowledge representations stored in the connections of overlapping, distributed networks of neurons. This postulate originates from the observation that three-layered neural networks learn the shared structure among training items by gradually adjusting their connection weights to minimise the error in the network outputs. For example, the Parallel Distributed Processing (PDP) framework to modeling cognitive processes employs a hidden layer to mediate between the input and output layers. Such networks are trained to map inputs to outputs by gradually adjusting the weights between their connections by using backpropagation algorithms (Rumelhart, Hinton, & Williams, 1986). One major advantage of these neural networks is their ability to generalise. That is, trained neural networks can classify data that they have never encountered before if it pertains to the same category as the learning data. For

example, early neural networks learned to compute the phonological form of words (Sejnowski & Rosenberg, 1987; Plaut, McClelland, Seidenberg, & Patterson, 1996; Guyonneau, VanRullen, & Thorpe, 2004) from exposure to letters and their corresponding phonemes in English words. These networks gradually learned structured knowledge representations in the connection weights among their units, based on the statistics of the training set and used these representations to pronounce a large corpus of monosyllabic words, including exception words, nonwords, and novel examples. Such representations are best described as being distributed and overlapping. According to the CLS model, skill acquisition in a variety of domains including language, perception and motor action as well as the residual learning abilities observed in amnesic patients are dependent on such neural networks (McClelland et al., 1995). This proposal echoes and extends Marr's original ideas by suggesting that the neocortex slowly classifies and organises information in long-term memory based on input statistics.

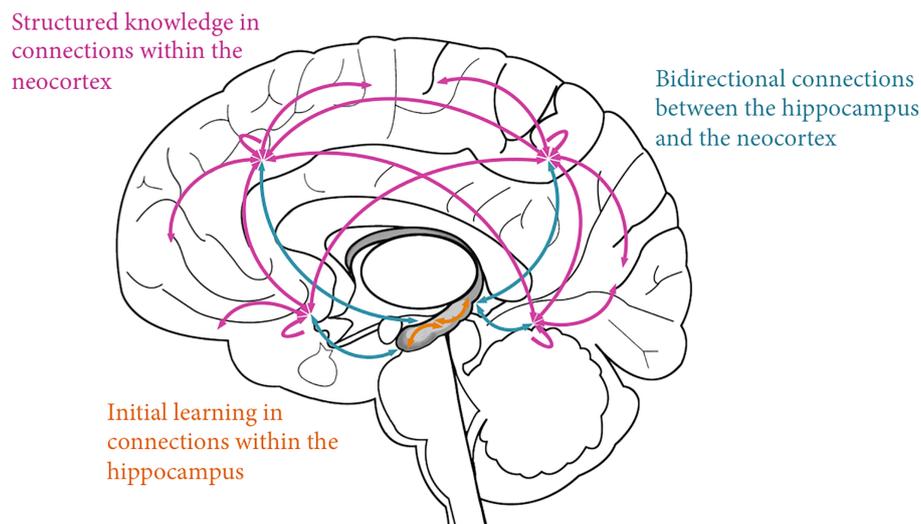


Figure 1. Complementary Learning Systems and their interactions

Schematic depiction of the CLS model on a sagittal view of the left hemisphere of the brain. The hippocampus is shown in dark grey. Orange arrows denote connections within the hippocampus. These connections exhibit rapid synaptic plasticity crucial for the rapid binding of the elements of an item or event into an integrated hippocampal representation. Systems consolidation involves hippocampal activity during replay spreading to neocortical areas via pathways depicted with blue arrows, thereby supporting learning within intra-neocortical connections (pink arrows). Adapted from Kumaran et al. (2016).

Despite the advantages associated with a system using overlapping, distributed representations in terms of robust performance, generalisation capacity and skill acquisition, it also has important limitations. First, the rapid adjustments of

connection weights to accommodate novel input-output mappings can disrupt or replace existing mappings – a phenomenon termed *catastrophic interference* (McCloskey & Cohen, 1989; Ratcliff, 1990; McClelland et al., 1995; French, 1999). McCloskey and Cohen (1989) demonstrated this phenomenon by using simulations of the AB-AC paradigm (Barnes & Underwood, 1959). In this paradigm, participants first learn word paired associates (A-B) and then learn to associate new words with the first words of the original pairs (A-C). Similar to the results obtained by Muller and Pilzecker (1900), learning of the second pairing typically impairs recall performance of the first pairings. While this effect is moderate in human participants, the effect in McCloskey and Cohen's (1989) network was to completely erase memory for the initial mapping. This issue of catastrophic interference is related to the *stability-plasticity* dilemma (Carpenter & Grossberg, 1987), which refers to the idea that learning requires plasticity for the integration of new knowledge but also stability to prevent forgetting of previous knowledge. Another limitation of a system using overlapping, distributed representations is that behaviour sometimes needs to be based on a single experience. For example, one severe allergic reaction to nuts should be sufficient to learn to avoid eating nuts. This type of one-shot learning would be prevented by a system that does not support memory for specific episodes.

To circumvent these limitations, the CLS model suggests a dual-memory system, which, fundamentally, simulates the presence of a short-term and a long-term memory and goes some way to addressing the stability-plasticity dilemma. As described above, the neocortical system employs a learning algorithm that produces overlapping, distributed representations. This system affords stability of representations over long periods of time despite changes in its structure or to the form of the input, which underlies its resistance to damage and ability to generalise, respectively. A second, complementary system, affords plasticity by allowing the rapid acquisition of information about individual items or episodes without interference from previously or subsequently learned ones. Extending Marr's and subsequent proposals (Mcnaughton & Morris, 1987; Treves & Rolls, 1992; O'Reilly & McClelland, 1994), the CLS model suggests that the hippocampus and related MTL structures support the initial storage of information about individual items or episodes. The hippocampal system employs a learning algorithm, which produces non-overlapping

representations allowing individuated storage of new items or episodes without overwriting preexisting knowledge. Evidence for the role of the hippocampus in recognition memory as well as its sensitivity to the specific context in which novel episodes occur is consistent with its role in the initial acquisition of information about individual items or episodes (Cohen & Eichenbaum, 1993; O'Reilly & Rudy, 2001; Norman & O'Reilly, 2003; Squire, Stark, & Clark, 2004; Davachi, 2006; Knierim, Lee, & Hargreaves, 2006; Mayes, Montaldi, & Migo, 2007; O'Reilly et al., 2014; Schiller et al., 2015).

The CLS model further posits a dynamic balance between two processes in the hippocampal system that are critical to avoiding catastrophic interference: *pattern separation* (making similar memories distinct by orthogonalising their neural representations) and *pattern completion* (reestablishing a past pattern of activity in response to partial or degraded input) (Marr, 1971; McNaughton & Morris, 1987; Treves & Rolls, 1992; O'Reilly & McClelland, 1994; McClelland et al. 1995; O'Reilly & Norman, 2002; Hunsaker & Kesner, 2013; Johnston, Shtrahman, Parylak, Gonçalves, & Gage, 2016; Knierim & Neunuebel, 2016). Pattern separation would allow differentiating between two birthday celebrations or two dinner dates, for example. The dentate gyrus (DG) and the CA3 subregion of the hippocampus are thought to be at the heart of pattern separation. With its sparse coding granule cells, the DG is the hypothesised source of the separation signal, which projects to the CA3 subfield of the hippocampus via strong mossy fibres (Treves, Tashiro, Witter, & Moser, 2008). Separation-like signals have been reported in the human DG/CA3 using high-resolution fMRI (Bakker, Kirwan, Miller, & Stark, 2008; Lacy, Yassa, Stark, Muftuler, & Stark, 2011). By contrast, pattern completion would allow remembering a specific birthday celebration or dinner date when returning to the restaurant where it took place. The ability to retrieve previous memories based on environmental cues is thought to require pattern completion, a process that is often ascribed to the CA3 hippocampal subregion (Knierim & Neunuebel, 2016). CA3 has been postulated as an attractor network containing excitatory recurrent connections, which attract the current firing pattern of the network to a stored pattern (Knierim & Zhang, 2012; Renno-Costa, Lisman, & Verschure, 2014). Importantly, this is thought to occur even

when the pattern of firing triggered by an external cue is incomplete (i.e. pattern completion).

According to the CLS, the hippocampal system cannot act alone due to its representational capacity and generalisation limitations. In fact, non-overlapping representations have an associated cost since the use of pattern-separated hippocampal codes for related events or episodes disregards their shared structure thereby limiting generalisation abilities. In contrast, dense similarity-based coding employed by the neocortical system is more adapted to supporting generalisation (Rolls, Treves, & Tovee, 1997; Kiani, Esteky, Mirpour, & Tanaka, 2007; Kriegeskorte et al., 2008; Leibold & Kempter, 2008; McNaughton, 2010; Bengio, Courville, & Vincent, 2013; Clarke & Tyler, 2014, Khaligh-Razavi & Kriegeskorte, 2014; Yamins et al., 2014). For example, neocortical activity patterns in high-level visual regions have been shown to exhibit more densely distributed coding and greater similarity-based overlap compared to hippocampal activity patterns (Barnes, McNaughton, Mizumori, Leonard, & Lin, 1990; Rolls et al., 1997; Kiani et al., 2007; Kriegeskorte et al., 2008; Leibold & Kemper, 2008; McNaughton, 2010; Khaligh-Ravazi & Kriegeskorte, 2014; Yamins et al., 2014).

2.2.1. Sleep and CLS

Within the CLS framework, interactions between the hippocampal and neocortical systems are hypothesised to play an important role in supporting systems consolidation (McClelland et al., 1995). Patterns of hippocampal activity that characterise the encoding of a new episode or event can drive the gradual integration of this new knowledge with existing knowledge in the neocortex. This is thought to occur when the hippocampus replays, or reinstates, the contents of a new episode or event back to the neocortex during offline periods such as sleep (McClelland et al., 1995; O'Neill, Pleydell-Bouverie, Dupret, & Jozsef, 2010; Wikenheiser & Redish, 2015). Hippocampal replay of new episodes is thought to be interleaved with replay of ongoing experiences allowing new memories to gradually become integrated into a neocortical network of existing memories. This is concomitant with a decay of hippocampal memories, either passively or through interference from new memories. Evidence for the role of sleep in supporting hippocampal replay comes primarily from

research on hippocampal place cells in rats (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996; Carr, Jadhav, & Frank, 2011). Hippocampal place cells responding to particular locations show large irregular activity (LIA) patterns during slow-wave sleep (SWS)⁷ that are distinct from the firing patterns observed during wake states (Figure 2). During LIA states, synchronous firing originating from the CA3 hippocampal subregion produces oscillatory patterns called sharp-wave ripples (SWRs) which propagate to neocortical cells. Importantly, SWRs are thought to reflect the reinstatement of recent experiences as they correlate with the firing rates of hippocampal cells to particular locations during wake states. The suggestion that SWRs propagate to the neocortex is supported by findings showing that SWRs are synchronised with neocortical oscillations (Sirota, Csicsvari, Buhl, & Buzsáki, 2003; Battaglia, Sutherland, & McNaughton, 2004). Furthermore, selective elimination of SWRs during SWS results in spatial memory impairments suggesting a causal role for replay in systems consolidation (Girardeau, Benchenane, Wiener, Buzsáki, & Zugaro, 2009; Nakashiba, Buhl, McHugh, & Tonegawa, 2009; Ego-Stengel & Wilson, 2010; Nokia, Mikkonen, Penttonen, & Wikgren, 2012).

Evidence suggesting sleep-dependent memory consolidation in humans has first come from behavioural studies reporting better performance on declarative and procedural memory tasks after overnight sleep compared to equivalent wake intervals (Barrett & Ekstrand, 1972; Plihal & Born, 1997, 1999; Stickgold, James, & Hobson, 2000; Fischer, Hallschmid, Elsner, & Born, 2002; Gais, Mölle, Helms, & Born, 2002; Walker, Brakefield, Hobson, & Stickgold, 2003; Tucker et al., 2006; Rasch, Büchel, Gais, & Born, 2007; Korman et al., 2007; Lahl, Wispel, Willigens, & Pietrowsky, 2008). Several of these studies have focused on comparing the effects of different sleep stages, specifically SWS and rapid eye movement (REM)⁸ sleep, on memory consolidation. The overall pattern of results suggests that SWS benefits the consolidation of

⁷ Slow-wave sleep (n.d.). In Merriam-Webster.com. Retrieved from <http://www.merriam-webster.com/dictionary/slow-wave%2Bsleep>: A state of deep usually dreamless sleep that occurs regularly during a normal period of sleep with intervening periods of REM sleep and that is characterised by delta waves and a low level of autonomic physiological activity.

⁸ REM sleep (n.d.). In Merriam-Webster.com. Retrieved from <http://www.merriam-webster.com/dictionary/rem%2Bsleep>: a state of sleep that recurs cyclically several times during a normal period of sleep and that is characterised especially by increased neuronal activity of the forebrain and midbrain, depressed muscle tone, dreaming, and rapid eye movements.

declarative memory whereas REM benefits the consolidation of procedural memory. For example, Plihal and Born (1997) found better recall of both paired-associated lists (declarative memory) and motor-tracing skills (procedural memory) after a period of sleep when compared to an equivalent period of wakefulness. Importantly, however, better recall of paired-associates lists improved more after intervals of SWS sleep whereas recall of motor-tracing skills improved more after intervals of REM sleep. Furthermore, Marshall, Helgadóttir, Mölle, & Born (2006) showed that the application of a weak electric anodal current with a frequency matching that of oscillations during SWS over prefrontal electrodes, where slow potential oscillations during SWS are thought to originate, improves the retention of a list of paired associates after sleep. These findings suggest a causal role for SWS in the consolidation of declarative memory and that its influence can be artificially enhanced.

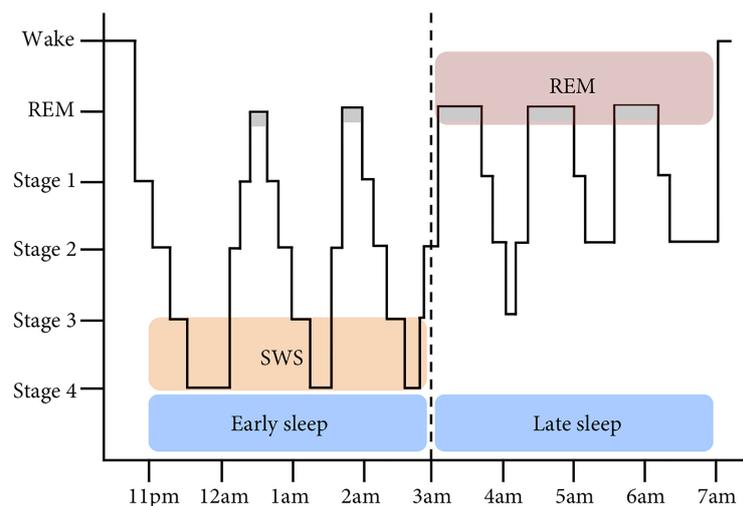


Figure 2. Sleep architecture and sleep stages

Sleep is characterised by the cyclic occurrence of REM sleep and non-REM sleep, which includes SWS (stages 3 and 4) and lighter sleep stages 1 and 2. In humans, early sleep is characterised by high amounts of SWS whilst REM prevails in late sleep. Adapted from Diekelmann and Born (2010).

While the results from several studies are consistent with a *dual-process hypothesis* (Plihal & Born, 1997, 1999), which implicates SWS in the consolidation of declarative memory and REM sleep in the consolidation of procedural memory, others studies have demonstrated that SWS can be associated with the consolidation of procedural memory (Gais, Plihal, Wagner, & Born, 2000, Huber, Ghilardi, Massimini, & Tononi, 2004, Aeschbach, Cutler, & Ronda, 2008) and REM sleep with the consolidation of declarative memory (Rauchs et al., 2004; Fogel, Smith, & Cote,

2007). These divergent results may reflect the fact that different memory tasks involve both declarative and procedural memory to different extents. Such results are also consistent with the *sequential hypothesis* (Giuditta et al., 1995), which suggests that it is the sequential appearance of SWS and REM sleep that confers the optimal benefits of sleep on the consolidation of both declarative and procedural memory. Consistent with this hypothesis, Stickgold, Whidbee, Schirmer, Patel, & Hobson (2000) reported a positive correlation between both the amount of SWS in the first quarter of sleep and the amount of REM sleep in the last quarter of sleep with improvements on a visual texture discrimination task relying on visual-procedural memory. Mednick, Nakayama, and Stickgold (2003) further showed greater improvements in visual texture discrimination following a short nap containing both SWS and REM sleep compared to a nap of the same duration containing only SWS. Similar results have been reported for sleep-related improvements in declarative memory. For example, Ficca & Salzarulo (2004) showed that disrupting the natural SWS-REM sleep cycle while leaving the time spent in each sleep stage unchanged impaired the morning recall of declarative verbal material.

Reports of sleep-dependent hippocampal replay in humans have also provided support for the role of sleep in memory consolidation. Using polysomnography and PET, Peigneux and colleagues (2004) showed that hippocampal regions that are activated while participants learn different routes in a virtual town are similarly activated during subsequent SWS. Improvements in performance in retrieving the different routes on the subsequent day were positively correlated with the amount of hippocampal activity expressed during SWS. Rasch et al. (2007) provided compelling evidence for a causal role for SWS hippocampal replay in memory consolidation. Participants learned spatial locations associated with odours. Re-exposure to the odours during SWS, but not during REM sleep or wakefulness, enhanced subsequent memory for the spatial locations. Importantly, fMRI showed significant hippocampal activation in response to odour re-exposure during SWS suggesting that it is particularly sensitive to inputs that can re-active memories.

As suggested by the CLS model, hippocampal replay is thought to play an important role in systems consolidation and in redistributing hippocampal memories

to the neocortex. Although a direct role for hippocampal replay is not yet established, there is evidence for a shift from hippocampal to neocortical representations with consolidation (Gais et al., 2007; Takashima et al., 2009). For example, Gais et al. (2007) provided evidence for interactions between the hippocampus and the medial prefrontal cortex (mPFC) in the consolidation of new declarative memories. Using fMRI, the authors measured participants' brain activity during the recall of trained word paired-associates, either following a period of post-training sleep or sleep deprivation. Post-training sleep led to greater hippocampal activation at recall and better recall of the word pairs than post-training sleep deprivation. Furthermore, there was greater functional coupling between the right hippocampus and the mPFC following post-training sleep than sleep deprivation. Gais et al. (2007) argue that these results may reflect the early stages of memory consolidation where the hippocampus still plays a central role in memory retrieval but interacts with the mPFC in a way that is consistent with the emergence of hippocampal-neocortical memory transfer. Similarly, Takashima et al. (2009) reported changes in functional connectivity between the hippocampal and neocortical systems following overnight consolidation, which are described in detail in Chapter 7.

In sum, there is ample evidence that new memory traces gradually become integrated into long-term memory through a process of memory consolidation. The CLS model provides a useful framework to characterise this process as involving the complex interplay of two memory systems, particularly during offline periods including sleep. The next section focuses on empirical work that has examined the role of memory consolidation and sleep in the domain of language learning specifically. The studies presented cut across several levels of language learning including lexical and morphological learning.

2.3. Consolidation in language learning

2.3.1. Consolidation of speech sounds and words

Only a few studies have focused on the role of memory consolidation in learning speech sounds (Fenn, Nusbaum, & Margoliash, 2003; Roth, Kishon-Rabin, Hildesheimer, & Karni, 2005; Eisner & McQueen, 2006; Fenn, Margoliash, & Nusbaum, 2013; Earle & Myers, 2015). For example, Fenn et al. (2003) pre-tested and

then trained participants to identify words presented in low-intelligibility synthetic speech. Training required participants to learn and generalise the mapping of synthetic speech to pre-existing phonological categories. One group of participants were pre-tested and trained in the morning and retested on a new set of words after a 12hr delay of waking (wake group). Another group of participants were pre-tested and trained in the evening and retested on a new set of words after a 12hr delay that included sleep (sleep group). A final group of participants were pre-tested, trained and re-tested within a single session without any retention interval. Performance improved immediately after training in the control group but degraded in the wake group. Importantly, performance recovered and was similar to that of the control group in the sleep group. Fenn et al. (2003) suggested that performance recovery in the sleep group implicates a role for sleep in refining and stabilising the new mappings between complex acoustic patterns and pre-existing phonological categories in a manner that supports generalisation to untrained material. As such, this account is consistent with the CLS model to the extent that overnight sleep may play a role in supporting the integration of new speech sounds with existing ones.

Evidence also suggests that the formation of new lexical representations requires memory consolidation (Gaskell & Dumay, 2003; Dumay & Gaskell, 2007; 2012; Tamminen & Gaskell, 2008; Davis, Di Betta, Macdonald, & Gaskell, 2009; Tamminen, Payne, Stickgold, Wamsley, & Gaskell, 2010; Henderson, Weighall, Brown, & Gaskell, 2012). These studies share the common assumption that spoken word recognition requires the activation of candidate words in the mental lexicon followed by the selection of the most likely candidate based on the information contained in the input (see the Cohort Model, Marslen-Wilson, 1987). Words that are simultaneously activated, because they share the same onset phoneme, for example, are said to enter in lexical competition (Gaskell & Marslen-Wilson, 1997). For new words to have a similar status to existing words (i.e. to have been integrated into the mental lexicon), they must, therefore, engage in lexical competition with existing words.

Gaskell & Dumay (2003) reported the first results suggesting a role for memory consolidation in lexical learning. Participants were trained on novel

nonsense words that overlapped with existing words (e.g. ‘*cathedruke*’ and ‘*cathedral*’) through a phoneme-monitoring task. The engagement of the novel words in lexical competition with the existing words was measured using a pause-detection task either immediately after training or one week later. In a pause-detection task, listeners have to identify a silent pause inserted in an existing word (e.g. ‘*cathe_dral*’). Slower detection times are taken as an index of lexical competition reflecting lexical integration of the novel words. The results showed that while recognition memory for the novel words was good immediately after training, lexical competition effects in the pause-detection task emerged only after a delay of one week. Gaskell and Dumay (2003) suggested that the distinct patterns of results for form recognition and lexical competition were consistent with phonological information being learned swiftly but lexicalisation developing more slowly and requiring memory consolidation.

Dumay and Gaskell (2007, 2012) went a step further and attempted to disentangle the effects of the passage of time and sleep after learning on the lexicalisation of novel words. One group of participants was trained on novel words through a phoneme-monitoring task at 8 am (am group) and another group of participants was trained at 8 pm (pm group). Recognition memory and lexical competitions effects, through pause-detection, were measured in both groups immediately after training, as well as 12hr and 24hr later. As in the previous experiment, participants showed good recognition performance immediately after training and at other retests. For the pause-detection task, the results showed that training on the novel words had no effect on lexical competition immediately after training in both groups. After a 12hr delay, the pm group showed a reliable lexical competition effect whereas the am group did not. A reliable lexical competition effect emerged in the am group after a 24hr delay after they had had the opportunity to sleep. The results thus suggest that sleep, rather than the mere passage of time, contributes to the lexicalisation of novel phonological forms since they start engaging in lexical competition only after a period including sleep. In sum, these lexical competition studies (see Bowers, Davis, & Hanley (2005) for similar lexical competition effects emerging when using written words) are consistent with a CLS account of word learning (see Davis & Gaskell, 2009) suggesting that the hippocampus provides the initial means to encode new phonological forms and that

the neocortex integrates these new phonological forms into the lexicon. Using polysomnography, Tamminen et al. (2010) as well as Tamminen, Lambon Ralph, and Lewis (2013) have further suggested a role for sleep in supporting this integration. They showed that spindle activity during sleep is positively correlated with subsequent lexical competition effects. Spindles are temporally aligned with SWRs. As noted earlier in the chapter, SWRs are synchronised with neocortical oscillations and consequently have been implicated in the hippocampal – neocortical crosstalk that supports systems consolidation.

Using fMRI, Davis, Di Betta, Macdonald, and Gaskell (2009) examined the neural bases of lexical learning and consolidation. The studies reporting consolidation effects in lexical learning described above (Gaskell & Dumay, 2003; Bowers et al., 2005; Dumay & Gaskell, 2007, 2012) all used a single training session followed by multiple test sessions potentially confounding consolidation-related changes in lexical learning with practice effects on the test tasks. To circumvent this methodological limitation, Davis et al. (2009) trained participants on two sets of novel words on two consecutive days using a phoneme-monitoring task and tested recognition performance and lexical competition effects only once on the second day. This design allowed testing within a single session for consolidation-related changes in the lexical representation of newly-learned words by comparing recognition performance and lexical competition effects for words trained on the day of testing or on the previous day. In a first behavioural experiment, Davis et al. (2009) found high levels of recognition performance for both sets of words. In contrast, only the sets of words trained on the previous day engaged in lexical competition with phonologically similar existing words confirming that earlier findings were not dependent on practice effects on the test tasks. In a second fMRI experiment, Davis et al. (2009) examined the neural bases of this lexical integration by comparing activations elicited by the set of novel words trained on the day of scanning (unconsolidated) or on the preceding day (consolidated) to similar untrained nonsense words and real words. A region of interest (ROI) analysis revealed that the average response to novel words was significantly reduced for consolidated relative to unconsolidated words in the left STG. Davis et al. (2009) interpreted this reduction in activation in response to consolidated words in areas of the brain involved in phonological processing as

indexing representations becoming more word-like following overnight consolidation. To assess the neural bases of the initial learning of new words, Davis et al. (2009) conducted an ROI analysis in the left hippocampus, which showed higher activation to untrained than unconsolidated words in the first scanning run suggesting hippocampal engagement upon first exposure to new words. Together, these findings are consistent with a CLS account of word learning whereby the hippocampus is involved in the initial learning of new words and the neocortex is involved in the representation of newly-learned words, which are subject to the influence of overnight consolidation.

Further evidence for the role of the hippocampus in word learning comes from an fMRI study by Breitenstein et al. (2005). During scanning, participants learned picture-pseudoword associations over five training blocks and were asked to press a button to report whether the associations were correct or not. This provided a behavioural measure of vocabulary proficiency, which was correlated with neural activity during learning. Increased vocabulary proficiency over the course of scanning was paralleled by a decrease in activity in the left hippocampus and an increase in activity in the left inferior parietal cortex, the latter having been suggested as an amodal semantic hub (Geranmayeh, Wise, Mehta, & Leech, 2014). While this study did not include a re-test after sleep, the results suggest that a consolidation process involving hippocampal-neocortical interactions begins at the onset of word learning. Adaptation to the novel words could not explain the reduction in hippocampal activity over the course of scanning since a control condition where novel words and pictures were randomly paired from trial to trial did not show the effect.

2.3.2. Consolidation of morphology

While there is accumulating evidence for a role for overnight memory consolidation in the domain of word learning, there is still limited behavioural and neural evidence for its role in morphological learning. Merkx, Rastle, and Davis (2011) examined the role of semantic information and memory consolidation in morphological learning. Participants were trained on novel affixes (e.g. ‘-*nept*’), which combined with existing words (e.g. ‘*sleep*’) to form new affixed words (e.g. ‘*sleepnept*’). During training, one group of participants heard the words alone whereas another group of participants

heard the words together with their definition. To assess knowledge of the trained words, participants in both groups completed a recognition task and a lexical decision task. To assess consolidation-related changes in the representation of the newly-learned affixed forms, participants were tested both two days and two months after training. The results showed that episodic affix knowledge, as indexed by performance on the recognition memory task, could be acquired after a single training session in either the form only or semantic learning groups. In contrast, the development of lexicalised affix representations, as indexed by slower reaction times in rejecting words containing trained novel affixes compared to words containing untrained novel affixes, benefited from the provision of semantic information and emerged only at re-test two months after training. Using a similar paradigm, subsequent studies have reported the formation of context-independent lexicalised affix representations, as indexed by participants' ability to generalise semantic affix knowledge to untrained stems, but only two days after training (Tamminen, Davis, Merkx, & Rastle, 2012; Tamminen, Davis, & Rastle, 2015). These studies are considered in more detail when discussing the role of consolidation in generalisation in Chapter 3.

Using MEG, Leminen et al. (2016) recently examined the neural mechanisms underpinning morphological learning and consolidation. Participants were trained on novel affixes (e.g. *-pe*, *-tu*), which combined with existing Finnish words (e.g. *'kukka'*, *'savu'*) to form new affixed words (e.g. *'kukkape'*, *'savutu'*) through a word-picture association task. The affixes referred either to the size or artificiality of familiar objects, two derivational categories that are not present in Finnish. Neural responses to combinations of real word stems (e.g. *'tuma'*, *'savu'*) not encountered in training with either trained (e.g. *-pe*, *-tu*) or untrained (e.g. *-ku*, *-ti*) affixes were measured both immediately after training and the next day. On the day of training, trained affixes combined with real word stems elicited greater activity in the left inferior frontal gyrus (IFG) around 50ms after the onset of the affix compared to untrained affixes. Leminen et al. (2016) suggest that these results are consistent with the involvement of the left IFG in morphological decomposition and the development of affix representations early in training. On the next day, the left STG showed a similar effect around 200 to 300ms after the onset of the affix. The authors argue that this result could reflect the development of a lexico-semantic affix representation

following overnight consolidation. The authors do acknowledge that their data does not permit to disentangle whether the development of these new lexico-semantic affix representations is driven just by the affix, the whole affixed form, or both, but they suggest that overnight consolidation plays a role in this process. Nonetheless, the results are consistent with those reported by Merks et al. (2011) and suggest that the development of lexicalised representations for newly-learned affixed forms benefits from periods of overnight consolidation. The authors augment previous findings by reporting changes in neural responses to newly-learned affixes in the left IFG immediately after training. The authors further suggest that their results extend the morpho-phonological decompositional model of morphology processing (Marslen-Wilson & Tyler, 2007) to morphological learning. The left IFG could immediately support morphological parsing during learning whereas the development of lexico-semantic representations in the left STG may benefit from overnight consolidation.

2.4. Summary

Newly-acquired labile memories become more stable over time through a process of memory consolidation. The CLS model provides a computational and neurobiological framework to characterise the types of computations performed by the hippocampal and neocortical systems and how these systems interact, particularly during sleep, to support learning and generalisation. The ideas articulated in the CLS model have recently been applied to the study of language learning and have been particularly successful in the domain of word learning. Converging behavioural and neuroimaging evidence suggests differential contributions of the hippocampal and neocortical systems to the rapid initial learning and long-term storage of newly-acquired words, respectively. The CLS model has also been applied to the domain of morphology although it remains rather under-explored in this area. Results so far seem to indicate that the learning of new affixed forms also benefits from periods of overnight consolidation. Given that morphology plays a key role in supporting linguistic generalisation and that the CLS model provides a detailed account of the computations supporting learning and generalisation, the latter offers a well-suited framework to study the former. The next chapter details how CLS support generalisation broadly and morphological generalisation more specifically.

Chapter 3: Generalising morphology

3.1. Generalisation

Generalisation refers to the application of knowledge derived from specific experiences to new stimuli or situations (Seger & Peterson, 2013). This general knowledge develops by integrating information across multiple experiences such that their similarities can be detected. These similarities allow experiences to be grouped into meaningful categories or concepts that can guide future behaviour. As such, generalisation forms the basis of adaptive behaviour and skill development. For example, consider a sommelier's ability to correctly identify a wine they have not tried before as a *Rioja*. This ability is predicated on retaining episodic memories for having tried different *Riojas* over the years such that their common appearance, aroma, and taste can support the development of general knowledge about wines from this particular region. Correspondingly, theories of declarative memory have suggested an *episodic – semantic* distinction such that knowledge about specific experiences (episodic memory) and decontextualised, abstract knowledge (semantic memory) are retained in separate memory systems (Tulving, 1972, 1985; Moscovitch, 1995; Moscovitch et al., 2005). A similar distinction has been proposed for recognition memory between *recall* of specific details about previously studied items and *familiarity* with previously studied items in the absence of recollection for the context in which they occurred (Yonelinas, 2002; Rotello, Macmillan, & Reeder, 2004). Both of these proposals share the assumption that different processes and, possibly representations, distinguish specific from general knowledge.

Developing in parallel to these memory theories, different cognitive models of category learning and their computational implementations have debated the nature of the representations supporting generalisation (Smith & Medin, 1981; Estes, 1994; Ashby & Maddox, 1998, 2005, 2010; Joseph, 2001; Humphreys & Forde, 2001; Murphy, 2002; Barsalou, 2003; Cree & McRae, 2003; Martin, 2007; Mahon & Caramazza, 2009; Seger & Miller, 2010). A powerful means to generalisation involves the use of categories, which capture the shared properties of items through organising principles that account for their relatedness. Cognitive models of category learning have mostly disagreed on the nature of these organising principles. Exemplar models

suggest that category representations consist of laying down episodic memories for all exemplars belonging to a category. When a new exemplar is encountered, its similarity to every exemplar within each of the various candidate categories is computed. It is then assigned to the category for which the sum of its similarities is greatest (Homa & Vosburgh, 1976; Brooks, 1978; Medin & Schaffer, 1978; Estes, 1986, 1994; Hintzman, 1986, 1988; Nosofsky, 1986; Nosofsky & Zaki, 1998; Lamberts, 2000; Nosofsky & Johansen, 2000). On this account, a sommelier's knowledge of the wine category '*Rioja*' arises from the knowledge that specific experiences tasting a '*Viña el Pisón*', a '*Gran Reserva 904*', and a '*Baron de Chirel*' are exemplars of '*Rioja*' wines. If a new wine is more similar to these known exemplars than it is to known exemplars of another wine category, it will be classified as belonging to the '*Rioja*' category. Thus, there is no decontextualised, abstract representation of a set of necessary features that define a '*Rioja*', or a derived summary representation of a '*Rioja*'. Generalisation is supported solely by the retrieval of the episodic representations of similar exemplars.

By contrast, abstractionist models suggest that category representations are summary representations abstracted over specific experiences with several category exemplars. A new exemplar is compared to the summary representation of the various candidate categories and assigned to the category of the one it best matches (Posner & Keele, 1968, 1970; Franks & Bransford, 1971; Reed, 1972; Rosch, 1973, 1975, 1977; Homa, Sterling, & Trepel, 1981; Rosch & Mervis, 1975; Smith & Minda, 1998; Minda & Smith, 2001). For example, a summary representation for a '*Rioja*' wine might blend the appearance, aroma, and taste of all experienced exemplars of '*Rioja*' wines. As such, a summary representation of a '*Rioja*' is a decontextualised, abstract representation that can be independent of any specific episodic reference to exemplars, which can be discarded. It is this abstract summary representation and not episodic representations for specific exemplars that supports generalisation.

Both exemplar and abstractionist models are single-system models of category learning and generalisation since they posit a unique form of representation underlying all categories. In fact, they can be considered as extremes on a continuum. At one extreme, exemplar models propose separate representations for each previously encountered category exemplar. At the other extreme, abstractionist

models suggest a single representation summarising all category exemplars. Both models have had success in accounting for learning and generalisation patterns on a range of categorisation tasks such that it is has proven difficult to differentiate them empirically (Ashby & Maddox, 1993; Minda & Smith, 2001, 2002; Smith & Minda, 2000; Zaki, Nosofsky, Stanton, & Cohen, 2003). Furthermore, the observation that participants retain exemplar information in performing classification tasks has posed a challenge to abstractionist models (e.g. Homa & Vosburgh, 1976). Similarly, the fact that summary representations that are not experienced during training can be classified as well or sometimes better than trained exemplars has challenged exemplar models (e.g. Homa, 1984 but see Nosofsky, 1988, 1991 for a dissenting view). These findings led to the suggestion that both episodic and abstract representations may contribute to generalisation.

Several dual-system models positing the operation of two types of representation supporting category learning and generalisation have been put forward (Anderson, 1991; Nosofsky, Palmeri, & McKinley, 1994; Palmeri & Nosofsky, 1995; Vandierendonck, 1995; Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998; Smith & Minda, 1998; Anderson & Betz, 2001; Love, Medin, & Gureckis, 2004). Although differing in detail and implementation, these models share the common assumption that one system computes abstract category representations, whilst a second system stores specific exemplar representations. These models typically perform better and explain a broader range of behavioural findings than single-system models (Nosofsky et al., 1994; Erickson & Kruschke 1998; Ashby et al., 1998). Dual-system models are also consistent with various observations suggesting that the formation of semantic memories largely depends on episodic memory (Rosenbaum, Winocur, & Moscovitch, 2001; Moscovitch et al., 2005 but see Gardiner, Brandt, Baddeley, Vargha-Khadem, & Mishkin, 2008 for a different account) and an intact hippocampal system (Gordon Hayman, Macdonald, & Tulving, 1993; Manns, Hopkins, Reed, Kitchener, & Squire, 2003; Bayley & Squire, 2005). Correspondingly, recent neurobiological models of category learning have posited interactions between the hippocampus and the neocortex in learning specific exemplars and integrating across multiple exemplars, respectively (Ashby & Ell, 2001; Seger, 2008; Kumaran, Summerfield, Hassabis, & Maguire, 2009; Seger & Miller, 2010; Davis, Love, &

Preston, 2012; Seger & Preston, 2013). However, such interplay must withstand the computational constraints associated with having to extract similarities among experiences whilst retaining specific information. The CLS model offers a possible solution by having the hippocampal and neocortical systems interact during systems consolidation as reviewed in Chapter 2. The next section focuses on the computational and neurobiological basis of generalisation derived from the CLS model.

3.2. Generalisation and Complementary Learning Systems

The CLS model suggests that the hippocampal system is chiefly involved in the rapid acquisition of item-specific, context-dependent information based on non-overlapping, pattern-separated representations that minimise interference between similar items. However, as differences between items are emphasised, similarities fall by the wayside, offering little possibility for generalisation (McClelland et al., 1995; Kumaran & McClelland, 2012). This seemingly serious limitation of the hippocampal system can be overcome by the neocortical system's ability to efficiently extract similarities across related items during memory consolidation. Indeed, neural network simulations of the neocortical system can successfully learn the similarity structure underlying sets of items and generalise to novel items. For example, neural networks have shown how semantic facts (e.g. 'can fly') shared by several concepts (e.g. 'robin', 'canary') can generalise to novel concepts (e.g. 'sparrow') (Hinton, 1986; Touretzky, 1987; Touretzky & Geva, 1987; Rumelhart, 1990; McClelland et al., 1995; Rogers & McClelland, 2003; McClelland et al., 2010). Neural networks of reading have shown how similarities in the spelling-to-sound mappings among related words (e.g. 'save', 'cave') can generalise to novel words (e.g. 'mave') (Sejnowski & Rosenberg, 1987; Seidenberg & McClelland, 1989; Plaut, 1995). The neural networks of the English past tense discussed in Chapter 1 have shown how similarities in the transformation from stem to past tense form among related verbs (e.g. 'crept', 'slept') can generalise to novel verbs (e.g. 'flept') (Plunkett & Marchman, 1991, 1993; Daugherty & Seidenberg, 1992). The capacity of these neural networks to generalise derives from the fact that they generate overlapping, distributed representations for

similar items in contrast to the non-overlapping, pattern-separated representations associated with the hippocampal system.

In its earliest formulation, the CLS model proposes that the development of similarity structure that can support generalisation proceeds slowly. The main goal of the neocortical system is to develop overlapping distributed representations that are optimised for an entire domain rather than for individual items. Each encounter with an item gives rise to small adjustments to the synaptic connections among the neurons involved in its representation. A slow learning rate allows the neocortical system to estimate the underlying distribution statistics of an entire domain more accurately by integrating over several samples. Many repetitions of the same item (cf. Chapter 1, token frequency) will lead to changes in the synaptic connections involved in the representation of that specific item (McClelland et al., 1995; Kumaran & McClelland, 2012). While this may improve behaviour relying on knowledge of a specific item, it will not support the discovery of similarity structure. Encountering several items that reflect some sort of systematic relationship (cf. Chapter 1, type frequency) allows the discovery of similarity structure as more examples of the distribution of possible items within a domain can be sampled (McClelland et al., 1995; Kumaran et al., 2016). Consider Plunkett & Marchman's (1993) network simulation of the English past tense discussed in Chapter 1. The ability of the network to discover the similarities shared across regular mappings and generalise to novel verbs required the training corpus to reach a critical mass of regular verbs (i.e. high type frequency). Extensive training on a small corpus was insufficient to support generalisation. Evidence for the importance of type frequency in promoting generalisation also comes from statistical learning studies showing that the learning and generalisation of non-adjacent dependencies between the 'A' and 'B' elements of an 'AXB' structure is facilitated by high variability of 'X' (i.e. high type frequency) (Gómez, 2002; Gómez & Maye, 2005; Onnis, Christiansen, Chater, & Gomez, 2003; Onnis, Monaghan, Christiansen, & Chater, 2004).

The development of similarity structure must also proceed slowly so that new item-specific information that is initially dependent on the hippocampus can become integrated into an existing neocortical network of related items without causing

catastrophic interference. This integration requires slow adjustments to the connection weights of the network structure to incorporate new items. Items that are similar to existing items will be most rapidly integrated because they will require small adjustments to the connection weights of the existing network structure. Items that are less similar to existing items or have idiosyncratic features will be integrated more slowly as they will require larger adjustments to the connection weights of the existing network structure (McClelland et al., 1995). Consequently, the development of similarity structure should be easier for highly regular domains (i.e. domains in which items consistently share the same features) than in quasi-regular domains (i.e. domains in which items do not consistently share the same features). These predictions are consistent with the results from Daugherty and Seidenberg's (1992) network simulation of the English past tense discussed in Chapter 1. The network performed better on phonologically consistent regular verbs (e.g. '*walk – walked*', cf. '*talk – talked*', '*stalk – stalked*') than on phonologically inconsistent regular verbs (e.g. '*bake – baked*', cf. '*take – took*', '*make – made*'). Analogous results have been reported for neural networks of reading where words with consistent spelling-to-sound mappings (e.g. '*dust*', '*must*', '*just*', '*crust*') are learned better than words with inconsistent spelling-to-sound mappings (e.g. '*crown*', '*down*' cf. '*shown*', '*grown*') (Plaut et al., 1996). Furthermore, in generalising to novel nonwords, both human readers and the network showed a greater tendency to produce irregular pronunciations for inconsistent nonwords compared with consistent nonwords (Plaut et al., 1996).

Whilst the early formulation of the CLS model attributes the development of similarity structure that can support generalisation to the slow-learning neocortical system, evidence suggests that both humans and animals are able to rapidly utilise relationships among related items after limited exposure (Bunsey & Eichenbaum, 1996; Dusek & Eichenbaum, 1997; Heckers, Zalesak, Weiss, Ditman, & Titone, 2004; Preston, Shrager, Dudukovic, & Gabrieli, 2004; Shohamy & Wagner, 2008; Zalesak & Heckers, 2009; Kumaran, 2012; Kumaran & McClelland, 2012; Zeithamova, Schlichting, & Preston, 2012; Schlichting, Mumford, & Preston, 2015; Eichenbaum, 2017). For example, in the transitive inference paradigm, participants learn a set of relationships among stimuli through trial and error (e.g. $A > B$, $B > C$, $C > D$, $D > E$).

The ability to generalise to novel inferential test trials (e.g. B ? D) immediately after training requires the ability to appreciate the hierarchical structure of the training pairs. Successful generalisation performance on this task has been shown to depend on the hippocampus in animals (Dusek & Eichenbaum, 1997) and humans (Zalesak & Heckers, 2009). The involvement of the hippocampus in this type of *rapid generalisation* (Kumaran & McClelland, 2012) contrasts sharply with the CLS model's proposal that its non-overlapping representations underlie item-specific memory at the expense of capturing similarity structure. It also runs contrary to empirical evidence for the role of the dentate gyrus and CA3 subregion of the hippocampus in pattern separation (Treves et al., 2008; Bakker et al., 2008; Lacy et al., 2011). Hence, there appears to be a tension between the computational and neurobiological principles underlying hippocampal representations and their proposed role in rapid generalisation.

To address this tension, the CLS model was revised to include a computational model accounting for rapid generalisation while preserving the assumption that the hippocampal system generates non-overlapping representations (Kumaran & McClelland, 2012). On this view, rapid generalisation in the transitive inference paradigm to novel pairs of items (e.g. B ? D) arises from the simultaneous activation of memory traces for trained pairs (e.g. B > C, C > D). As such, this addition to the original CLS model extends classical exemplar models of category learning discussed at the outset of this chapter (e.g. Medin & Schaffer, 1978; Hintzman, 1986, 1988; Nosofsky, 1986). Recurrent connections within the hippocampal system are argued to drive the simultaneous activation of episodic memory traces. This proposal is consistent with neural recordings in rats showing that hippocampal outputs originating in CA1 and subiculum reenter hippocampal subfields via the entorhinal cortex (Kloosterman, van Haeften, & Lopes da Silva, 2004). This recurrence is posited to allow the hippocampal system to compute similarity among related items that can support rapid generalisation whilst maintaining non-overlapping representations. Network simulations implementing recurrence have replicated patterns of rapid generalisation in transitive inference paradigms (Kumaran & McClelland, 2012). Furthermore, whilst in its original formulation the CLS emphasises hippocampal replay of individual episodes during memory consolidation, the proposal of

recurrence yields the prediction that it may replay multiple related traces (i.e. *generalised replay*, Kumaran & McClelland, 2012). This proposal is supported by neural recordings in rats showing that path trajectories replayed during SWRs can occasionally represent the construction of novel path trajectories, including ‘shortcuts’, never experienced during wakefulness (Gupta, van der Meer, Touretzky, & Redish, 2010; Wu & Foster, 2014). It remains unclear, however, whether emergent generalisations created through recurrence within the hippocampal system are themselves stored in memory or indeed whether they enhance neocortical learning in comparison to the replay of individual episodes (Walker & Stickgold, 2010; McClelland & Kumaran, 2012; Kumaran, 2012).

In an early account of memory consolidation supporting the development of similarity structure, McClelland et al. (1995) referred to a protracted process lasting days to months to even years. However, more recent proposals suggest that memory consolidation could be accomplished much more quickly, within days or even hours. Recent evidence suggests that new information that is highly consistent with existing structured knowledge can be rapidly integrated into neocortical networks without catastrophic interference (McClelland, 2013; Tse et al., 2007, 2011). For example, Tse et al. (2007) trained rats to associated six unique flavour cues with six specific locations where they could find food rewards in a complex spatial environment. Two weeks of daily training sessions were required for rats to learn all six flavour – location associations. They were then trained on four of the learned flavour – location associations whilst the remaining two associations were replaced with two new ones whose locations were close to the old replaced locations. A test the next day revealed that the rats had learned the two new associations as they searched for food rewards at the new cued locations far more than at the old cued locations. The following day, they received either hippocampal or sham control lesions. Rats that had received hippocampal lesions retained both the original and the new flavour – location associations. Importantly, Tse et al. (2007) showed that this learning was dependent on the presence of prior structured knowledge, or *schema*. No such learning was observed when rats that had been trained in one spatial environment were trained on new flavour – location associations in a novel spatial environment. The results suggest that new memories that are consistent with an existing schema become hippocampally

independent much faster than new memories that lack prior schema. Recent fMRI studies in humans also suggest that new information becomes independent of the hippocampus more rapidly when it conforms to a schema. During a rest period shortly after encoding new visual information, functional connectivity between the hippocampus and the medial prefrontal cortex (mPFC) has been found to decrease for schema consistent compared to schema inconsistent information (van Kesteren, Fernández, Norris, & Hermans, 2010) whilst successful retrieval of schema consistent information has been found to be associated with increased functional connectivity between mPFC and a network of somatosensory regions coding the new information (van Kesteren, Rijpkema, Ruiters, & Fernández, 2010). These results support the view that schema consistent novel information is integrated more rapidly into neocortical networks than schema inconsistent novel information and that this process may be mediated by the mPFC.

A separate line of research has also shown that a single period of sleep could enhance generalisation performance (Ellenbogen, Hu, Payne, Titone, & Walker, 2007; Durrant, Taylor, Cairney, & Lewis, 2011; Lau, Alger, & Fishbein, 2011; Coutanche, Gianessi, Chanals, Willison, & Thompson-Schill, 2013). For example, Ellenbogen et al. (2007) used a transitive inference paradigm to test two groups of participants on their memory for the set of relationships learned during training (e.g. $A > B$, $B > C$, $C > D$, $D > E$) as well as on their ability to generalise to novel inferential test trials (e.g. $B ? D$) after an interval of sleep or wakefulness. Both groups performed similarly on the trained pairs, but only the sleep group showed evidence of generalisation on inferential tests. These results may be interpreted as providing evidence for the role of sleep in supporting the rapid development of overlapping neocortical representations that can support generalisation. However, as discussed above, the CLS model would predict the development of overlapping experiences to proceed over a considerably longer period of time, particularly in the absence of prior schema. Amendments to the original CLS model have suggested that recurrence within the hippocampal system can compute similarity among related experiences to support rapid generalisation. On this view, replay during sleep may increase the strength, robustness, and rate of activation of new hippocampally-dependent memories and this strengthening may account for the capacity for inference emerging only after sleep (Kumaran &

McClelland, 2012; Kumaran et al., 2016). By implementing their recurrent network architecture, Kumaran and McClelland (2012) replicated Ellenbogen et al.'s (2007) results by appealing to these simple hippocampal mechanisms rather than to the development of overlapping neocortical representations.

An alternative interpretation put forward by Lewis and Durrant (2011) proposes that replay of new hippocampal memories can actively support the development of overlapping neocortical representations. On this view, neurons that code for overlapping elements across memories will be more strongly activated than neurons that code for the idiosyncratic elements of individual memories. Through long-term potentiation, the synaptic connections between neocortical neurons coding for overlapping elements become stronger than the synaptic connections between neocortical neurons coding for idiosyncratic elements of individual memories. According to the *synaptic homeostasis hypothesis* (Tononi & Cirelli, 2003, 2006), SWS promotes a generalised downscaling of recently potentiated synaptic connections returning total synaptic weight to an energetically sustainable baseline level. All synaptic connections converging into the same neuron are downscaled proportionally such that total synaptic weight can be reduced whilst preserving the relative differences in the strength of synaptic connections that are important for preserving memories. Lewis and Durrant (2011) suggest that the more strongly potentiated synaptic connections for overlapping elements have a higher chance of withstanding synaptic downscaling (Figure 3). The repeated replay of newly-acquired memories in different combinations, acting synergistically with synaptic downscaling, allows the progressive development of overlapping representations. Thus, the sleep-related gains in generalisation performance reported by Ellenbogen et al. (2007) may be accounted for by overlapping replay of trained pairs allowing the construction of a representation of the full hierarchy, which can support generalisation to novel inferential trials.

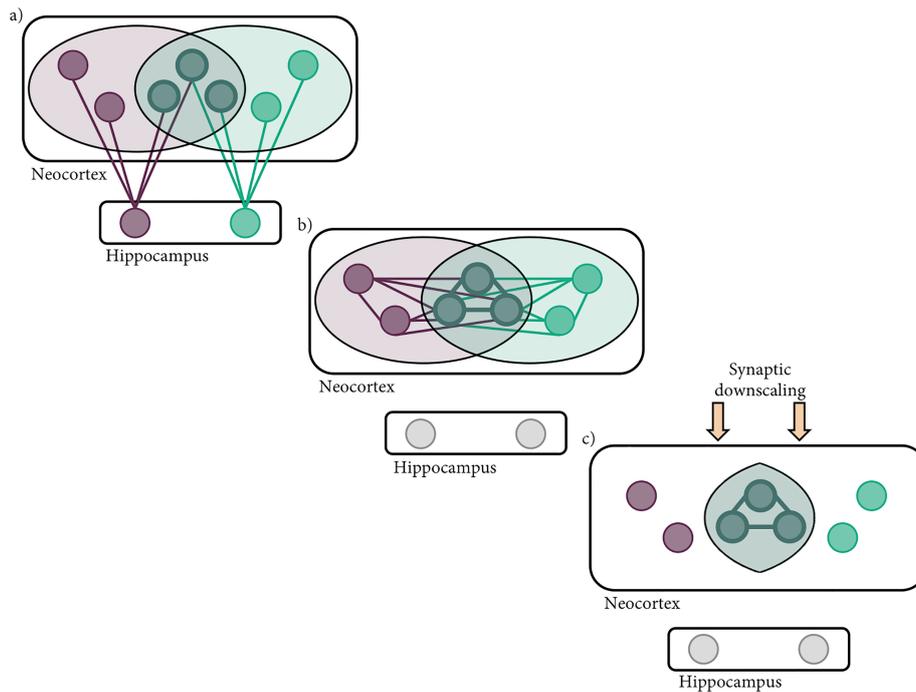


Figure 3. Hippocampal replay and neocortical consolidation

Schematic representation of the development of overlapping neocortical representations through hippocampal replay. a) two hippocampal memories are reactivated simultaneously during sleep. This reactivation includes some neocortical neurons unique to each memory, and some that are shared by both memories. b) shared neurons are potentiated more strongly and, as a result, develop stronger connections to each other than to neurons not involved in this overlap. c) following synaptic downscaling, only these strong connections between the shared neurons remain intact. This set of neurons now represents abstraction from the two memories, encoding their shared information only. Thicker lines indicate stronger synaptic connections. Adapted from Lewis and Durrant (2011).

3.3. Morphological generalisation

The ability to generalise from a limited number of experiences lies at the heart of morphological generalisation. From exposure to a finite number of inflected words (e.g. ‘walked’, ‘jumped’, ‘laughed’), language users can readily generalise to novel words that enter the language (e.g. ‘Skyped’) or nonwords (e.g. ‘wugged’). The debate introduced in Chapter 1 between single- and dual-mechanism theories of the acquisition, representation, and processing of regular and irregular English past tenses has naturally also considered the issue of generalisation. Both regular and irregular inflections have been shown to generalise to nonwords (e.g. *wug* – *wugged*; ‘*splung* – *splung*’) (Bybee & Slobin 1982; Bybee & Moder, 1983; Prasada & Pinker 1993). Dual-mechanism theories attribute regular generalisations to the application of symbolic rule and irregular generalisations to analogical processing in an associative memory network (Pinker & Prince, 1988, 1994; Pinker, 1991, 1999; Prasada & Pinker, 1993; Clahsen, 1999). By contrast, single-mechanism theories attributed all generalisations

to regularities extracted across mappings between stems and inflected forms within a single network of connections (Rumelhart & McClelland, 1986; Bybee, 1985, 1988; 1995, 2001). The influence of frequency measures and phonological consistency discussed in Chapter 1 also bears on the issue of generalisation.

3.3.1. Frequency

The regular English past tense applies productively to novel forms much more than irregular inflections (Bybee & Moder, 1983; Prasada & Pinker, 1993; Albright & Hayes, 2003). Consequently, it has been characterised by proponent of dual-mechanism theories as a *default* inflection reflecting the application of a rule (Prasada & Pinker, 1993). Single-mechanism theories have argued that default generalisation can be captured within a single network of connections without invoking a rule-based process. Network simulations of the English past tense are capable of treating the regular past tense as the default, generalising the ‘*ed*’ affix to the majority of untrained verbs and producing overregularisation errors similar to those produced by children when type frequency is sufficiently high (Rumelhart & McClelland, 1986; Plunkett & Marchman, 1991, 1993). Thus, single-mechanism theories argue that high type frequency is a determining factor of default generalisation.

Prasada and Pinker (1993) have criticised these network simulations on the basis that their ability to treat the regular past tense as a default is an artefact of the idiosyncratic frequency measures of the English past tense. Regular verbs have high type frequency, which allows networks to construct a well-populated default category that covers large portions of the language’s phonological space. Novel forms are thus more likely to fall close to a known regular verb and be regularised. By contrast, irregular verbs have low type frequency but relatively high token frequency allowing their memorisation in terms of a number of narrow phonological subcategories (Bybee & Slobin, 1982; Bybee, 1985; Hooper, 1976; Hare & Elman, 1995). However, other languages are argued to have low type frequency inflections that act as defaults. These are commonly referred to as *minority defaults*. The German –s noun plural, the German –t past participle and the Modern Standard Arabic sound plural have all been argued to represent examples of minority defaults (McCarthy & Prince, 1990; Clahsen & Rothweiler, 1992; Clahsen, Rothweiler, Woest, & Marcus, 1992; Marcus, Brinkman,

Clahsen, Wiese, & Pinker, 1995; Clahsen, 1999). Dual-mechanism theories argue that minority defaults provide evidence that default generalisation can be independent of high type frequency and further argue that network simulations would be unable to account for this phenomenon (Pinker, 1991, 1999; Prasada & Pinker, 1993; Marcus et al., 1995).

Daugherty and Hare (1993) as well as Hare, Elman, and Daugherty (1995) trained networks to map the phonological representation of verb stems to their past tense form given a training corpus representative of Old English. The training corpus comprised five classes of highly phonologically structured irregular verbs and a sixth class of phonologically varied regular verbs. Each class contained the same number of verbs presented the same number of times in training. Both networks sought to show that the sixth class could act as default inflection by virtue of its phonological diversity despite having low type frequency. Generalisations to the first five classes were shown to be dependent on the phonological similarity of the novel forms to trained forms. By contrast, phonological similarity was irrelevant for generalisations to the sixth class. Hare and Daugherty (1993) and Hare et al. (1995) argue that the sixth class acts as an attractor for novel forms even when they do not resemble trained regulars. Thus, connectionist models are able to learn and generalise a default inflection even in the absence of higher type frequency.

Careful examination of the German *-t* participle and the Modern Standard Arabic sound plural call into question their description as minority defaults based on methodological flaws in calculating type frequencies (Bybee, 1995; Plunkett & Nakisa, 1997; Boudelaa & Gaskell, 2002). In fact, both systems are more accurately described as majority defaults in which high type frequency plays a determining role in default generalisation. Whilst the German *-s* noun plural has been more accurately described as a minority default, it does not apply as uniformly as would be predicted by dual-mechanism theories. Marcus et al. (1995) reported that German speakers were more likely to use the German *-s* noun plural to inflect nonwords that were intended to represent German nouns, borrowings, and proper names. They argued that the heterogeneity of these contexts suggests that the German *-s* noun plural serves as a default despite having low type frequency. However, reanalyses of Marcus et al.'s

(1995) results show that for nonwords intended to be like German common nouns and borrowings, the –s plural was preferred but only if the nonwords did not resemble existing German words. Surnames were the only case where the –s plural was used uniformly (Bybee, 1995; Hahn & Nakisa, 2000). Such specificity in application undermines the claim that the German –s noun plural acts as a default inflection and instead rallies more strongly with single-mechanism theories arguing that it is sensitive to phonological and contextual factors.

Bybee's usage-based model (1985, 1988, 1995, 2001) provides a psycholinguistic account of the ways in which type and token frequency influence the generalisation of morphological patterns. When several distinct words share an affix (i.e. high type frequency), these words will be related to one another and the representation of the affix will emerge. The higher the type frequency, the stronger the affix representation will be and the easier it will be to access when generalising to new words. Words with high individual token frequency are more autonomous and less likely to be acquired by forming relations with other words. The higher the token frequency, the stronger the lexical representation of the word and the less likely it will be to contribute to the generalisation of a morphological pattern. For example, Bybee (1995) notes that the '*strung*' and '*swept*' classes of irregular verbs in English have approximately the same type frequency but the '*strung*' class is much more productive than the '*swept*' class. The total token frequency of the '*swept*' class is almost three times that of the '*strung*' class. Thus, for equal type frequency, higher token frequency may impede generalisation by preventing the emergence of relations among class members.

3.3.2. Phonological consistency

Whilst English speakers generalise the regular past tense much more, they sometimes generalise irregular patterns. For example, English speakers typically rate the form '*splung*' as an acceptable past tense of the novel verb '*spling*' (Bybee & Moder, 1983; Prasada & Pinker, 1993; Albright & Hayes, 2003). This willingness to generalise an irregular pattern to the novel verb '*spling*' is likely attributable to the fact that it is phonologically similar to its neighbours '*cling*', '*fling*', '*sting*', '*string*', '*swing*', and '*wring*', which all follow the same vowel alternation pattern to form their past tense.

According to proponents of dual-mechanism theories, these occasional irregular generalisations occur because the associative memory network can block the application of the regular rule if novel forms are similar to known irregulars. On this view, phonological similarity to existing regular verbs should not influence the application of the rule.

Prasada and Pinker (1993) tested English speakers' acceptability judgements for novel past tense forms that varied in their phonological similarity to existing regular and irregular verbs. Participants' willingness to generalise irregular patterns was related to the phonological similarity between the novel forms and existing irregulars. That is, novel forms that were phonologically similar to existing irregulars (e.g. '*spling - splung*' as in '*cling - clung*') were more highly rated than novel forms that were phonologically dissimilar to existing irregulars (e.g. '*nist - nust*'). By contrast, participants' willingness to generalise the regular past tense was unrelated to the phonological similarity between the novel forms and existing regulars. Novel forms that were phonologically similar to existing regulars (e.g. '*plip - plipped*' as in '*slip - slipped*') were rated similarly to novel forms that were phonologically dissimilar to existing forms (e.g. '*ploamph - ploamphed*'). Prasada and Pinker (1993) argue that these results provide evidence that regular generalisations are insensitive to phonological consistency. McClelland and Patterson (2002b) note that Prasada and Pinker's (1993) results may be attributed to a confound in their stimuli. Indeed, a novel form can only be phonologically dissimilar to existing regulars by being phonologically dissimilar to all English words. Consequently, the novel forms that were phonologically dissimilar to existing regulars were also phonologically deviant as English words (e.g. '*ploamph*'). In fact, these verbs received low ratings suggesting that regulars are in fact sensitive to phonological consistency. In an attempt to correct for this, Prasada and Pinker (1993) subtracted ratings of the stem (e.g. '*ploamph*') from ratings of the past tense forms (e.g. '*ploamphed*'). This may, however, have corrected a real effect.

Albright and Hayes (2003) avoided this confound by using novel forms that were high in phonological acceptability. Participants were asked to provide acceptability judgements for novel past tense forms that either fell in an island of

reliability or not. Albright and Hayes (2003) refer to a phonological context of relatively high phonological consistency for a particular inflectional pattern as an island of reliability. For example, every English verb stem that ends in a voiceless fricative has a regular past tense (e.g. ‘*miss – missed*’, ‘*wish – wished*’, ‘*laugh – laughed*’). Islands of reliability exist for regulars and irregulars. Albright and Hayes (2003) showed that participants were sensitive to such islands of reliability not only for novel irregular verbs (e.g. ‘*fleep – flept*’, ‘*gleed – gled*’) but also for novel regular verbs (e.g. ‘*bredge – bredged*’, ‘*nace – naced*’). These results go against the predictions of dual-mechanism accounts in which all regulars are derived by rule, and thus would not be expected to show phonological consistency effects.

3.3.3. Generalisation of morphology and CLS

The evidence reviewed in the previous section supports the view that frequency measures and phonological consistency influence how regular and irregular morphological patterns generalise. As discussed at the outset of this chapter, the CLS model also predicts that both of these factors influence the development of overlapping representations that capture the similarity structure among items in a given domain. The CLS model thus offers a well-suited framework to characterise the development of representations that can support morphological generalisation. As reviewed in Chapter 2, the CLS model has been successful in accounting for the integration of new words in the mental lexicon as well as morphological learning, albeit to a more limited extent. However, the application of the CLS model to the problem of morphological generalisation has received very little attention.

Tamminen et al. (2012) examined the development of generalisable affix knowledge in adults. Using a similar training procedure as in Merkx et al. (2011), participants learned the meaning of novel derivational affixes (e.g. ‘*-afe*’), which combined with existing words (e.g. *sleep*) to form new affixed words (e.g. ‘*sleepafe*’). To assess participants’ knowledge of the trained affixes, Tamminen et al. (2012) used a speeded shadowing task in which participants had to repeat aloud a spoken affixed word as quickly as possible. Importantly, the task involved the novel affixes presented either in their training context (e.g. ‘*sleepafe*’) or in a new stem context (e.g. ‘*sailnafe*’). Shadowing of these stimuli was compared to matched controls with untrained affixes

(e.g. ‘-*oke*’ as in ‘*floathoke*’, ‘*friphoke*’). To assess the role of memory consolidation in the development of context-independent affix representations, a group of participants was tested immediately after training and another group of participants 2 days after training. Participants tested immediately after training showed an advantage for repeating affixes presented in their training context (e.g. ‘*sleepafe*’) compared to untrained affixes (e.g. ‘*floathoke*’). Only participants tested two days after training showed an advantage in shadowing affixes presented in novel contexts (e.g. ‘*sailafe*’) compared to untrained affixes (e.g. ‘*friphoke*’). Tamminen et al. (2012) argue that these results are consistent with context-independent affix representations developing with overnight memory consolidation. In a second, non-speeded generalisation task, participants had to select between two possible definitions for untrained stems combining with trained affixes (e.g. ‘*sailafe*’). Correct performance required participants to select the definition that was consistent with the affix meaning as encountered in training. Foil definitions combined the meaning of the stem with the meaning of a different trained affix. Participants performed well on this task regardless of the day of training. Tamminen et al. (2012) suggest that immediate generalisation on the definition task may reflect the use of context-dependent representations that may not be sufficient for online linguistic processing on the speeded task. This interpretation is reminiscent of the distinction in updated versions of the CLS model outlined above whereby generalisation based on retrieval of multiple episodic traces may depend on slower recurrent processes within the hippocampus compared to the more efficient generalisation possible with overlapping neocortical representations.

In another series of experiments, Tamminen et al. (2015) used a similar paradigm to examine some of the factors influencing the generalisation of newly-learned affixes. Of particular relevance here, one experiment investigated whether type frequency influences morphological generalisation. Participants were trained on eight novel affixes, half of which were trained with high type frequency (combining with 8 stems each) and half of which were trained with low type frequency (combining with 2 stems each). To keep the number of affix presentations the same during training, the low type frequency affixes were presented more often (high token frequency) than the high type frequency affixes (low token frequency). Tamminen et

al. (2015) argue that this manipulation should enhance episodic memory for the novel words and affixes in the low type frequency condition, allowing to potentially dissociate episodic memory strength from emerging generalisation effects. The generalisation test involved a sentence congruency task in which participants read a sentence context and had to repeat the final word aloud. The final words were formed by combining untrained stems with trained affixes (e.g. '*sailafe*'). The sentence context could either be congruent or incongruent with the meaning of the affix. For affixes trained with high type frequency, participants were slower at reading novel affixed words presented in an incongruent sentence context whereas there was no such congruency effect for affixes trained with low type frequency. Performance on a recognition memory task, however, showed that participants were significantly more accurate at recognising affixes trained with low type frequency (but high token frequency) than affixes trained with high type frequency. Tamminen et al. (2015) argue that the accumulation of learning episodes (i.e. high token frequency) for the low type frequency affixes appears to benefit episodic memory, but not generalisation. By contrast, successful generalisation of semantic knowledge about affixes depends on high type frequency. These results are consistent with the CLS model in that high type frequency facilitates the development of overlapping representations that can support generalisation.

3.4. Summary

In sum, the ability to group items and experiences into meaningful categories or concepts is fundamental to generalisation. The CLS model provides a computational and neurobiological framework that balances the advantages and disadvantages of having to develop structured knowledge representations whilst maintaining item-specific information. However, the way in which these different types of knowledge may support generalisation or the developmental course of overlapping neocortical representations remains unclear. The development of general knowledge that can be applied to novel instances is central to morphological generalisation. The factors that are known to influence morphological generalisation, such as type and token frequency and phonological consistency, are also suggested by the CLS model to influence the development of overlapping neocortical representations. Thus, the CLS

model offers an ideal framework to investigate the types of representations that support morphological generalisation. Thus far, the types of representations that support morphological generalisation have been mostly inferred on the basis of network simulations and behavioural experiments. However, no work has looked at characterising the neural representations that might support generalisation. The next chapter details how the work presented in this thesis aims to address this empirical lacuna by drawing upon the literature reviewed thus far.

Chapter 4: Learning, consolidating, and generalising morphology

4.1. Building bridges

The previous introductory chapters have provided a background to three ostensibly distinct literatures: morphological learning, memory consolidation, and generalisation. However, these chapters have also highlighted ways in which these three strands of research share important commonalities and how they can cross-fertilise. The empirical work presented in this thesis makes an original contribution to the existing body of knowledge by building bridges between these three strands of research. In doing so, it aims to contribute to a better understanding of the cognitive and neural mechanisms supporting the learning, consolidation, and generalisation of morphology.

The evidence reviewed in Chapter 1 supports the view that morphological learning relies on the complex interplay of multiple competing factors, such as type and token frequency and phonological consistency, within a single network of connections. However, a paucity of behavioural studies has considered the influence of these factors on morphological learning with most evidence coming from child acquisition data, network simulations, and diachronic change data. As yet, there is no empirical demonstration of the influence of these factors on the neural representations and mechanisms underpinning morphological learning or how consolidation processes may modulate this influence.

A characterisation of the neural representations and mechanisms supporting morphological learning requires considering the functional organisation of memory. The CLS model introduced in Chapter 2 provides a computational and neurobiological framework to characterise hippocampal and neocortical representations and how these interact during consolidation to support learning. The CLS model has been successfully applied to word learning to account for the process of integrating new words into the mental lexicon (see Davis & Gaskell, 2009). This account is now well-supported by behavioural and neuroimaging evidence suggesting that the hippocampal system supports the rapid acquisition of new words whilst the

neocortical system supports the integration of new words into existing neocortical networks through memory consolidation processes. It is less clear, however, whether CLS principles can also provide an explanatory framework for morphological learning. To date, there is limited behavioural evidence concerning consolidation effects in morphological learning (e.g. Merx et al., 2011; Tamminen et al., 2012, 2015) and only one study has considered their neural basis (Leminen et al., 2016). It is thus clear that further functional imaging can provide crucial evidence for the role of consolidation in the formation of neocortical representations of new morphological knowledge and how it may modulate the influence of factors such as type and token frequency and phonological consistency.

At the lexical level, morphology offers a powerful example of linguistic generalisation. Thus, considering morphological learning and its underlying neural representations and mechanisms provides a window into the cognitive and neural bases of generalisation. Chapter 3 considered how CLS principles account for the development of representations that can support generalisation. It also highlighted some of the tensions relating to rapid generalisation based on hippocampal representations and the developmental course of overlapping neocortical representations. Thus, characterising the neural representations and mechanisms supporting morphological learning and how these may relate to generalisation could offer insight into resolving these tensions. The evidence reviewed in Chapter 3 also considered how morphological generalisation is subject to the influence of type and token frequency, phonological consistency, and overnight consolidation (e.g. Hare et al., 1995; Albright & Hayes, 2003; Tamminen et al., 2012, 2015). Again, how these factors influence the neural representations and mechanisms supporting new morphological knowledge and how they may influence generalisation, remains unexplored.

Across the three previous chapters, a particular emphasis has been put on network simulations of morphological learning and generalisation (e.g. Rumelhart & McClelland, 1986; Plunkett & Marchman, 1991, 1993; Hare & Elman, 1995; Hare et al., 1995). These networks capture important features of morphological learning and generalisation including type and token frequency effects, phonological consistency

effects, default generalisations, overregularisations, and irregular generalisations. These types of behaviour arise as emergent properties of neural networks discovering regularities shared among similar stem – past tense mappings using overlapping distributed representations. Usage-based models (e.g. Bybee, 1985, 1988, 1995, 2001) make similar claims. The formulation of the CLS model is grounded in the observation that neural networks can learn and generalise structured knowledge representation by gradually adjusting connections in overlapping, distributed networks of neurons. These observations are borne out by recent neuroimaging work, particularly in the domain of visual object recognition, showing more densely distributed coding and greater similarity-based overlap in the neocortex compared to the hippocampus. Thus, the CLS model and existing computational accounts can be fruitfully combined to derive predictions concerning the neural representations and mechanisms supporting morphological learning and generalisation.

4.2. Thesis outline

On the basis of these empirical gaps, the experiments described in this thesis focus on characterising the cognitive and neural representations supporting morphological learning and generalisation within a CLS framework. The first behavioural experiment explores the extent to which the contribution of overnight consolidation to morphological generalisation varies with the type and token frequency and the phonological consistency of new inflected words. It also introduces the behavioural paradigm used in subsequent experiments. This is the focus of the next chapter. The second experiment combines this behavioural paradigm with fMRI to examine the role of overnight consolidation in influencing the neural representations underlying new inflected words. Advanced multivariate fMRI analysis methods, namely Representational Similarity Analysis and task-based functional connectivity, are used to test CLS predictions. Overviews of these analysis methods are provided in Chapters 6 and 7, respectively. Neuroimaging results are presented in Chapter 8. Finally, the third behavioural experiment attempts to shed some light on inconsistencies in the behavioural results obtained in the first two experiments. This is addressed in Chapter 9. Chapter 10 discusses the conclusions that can be drawn from the experimental data and suggests directions for future research.

Chapter 5: Experiment 1

5.1. Introduction

As discussed in the introductory chapters, child acquisition data, neural network simulations, and a more limited number of behavioural studies have shown that type and token frequency, as well as phonological consistency, play a role in the learning and generalisation of morphology. However, the role of memory consolidation in modulating the influence of these factors has largely been ignored. Only one study has considered the influence of type and token frequency on the generalisation of newly-learned derivational affixes (Tamminen et al., 2015). The few experiments that have considered the influence of phonological consistency on generalisation have focused on proficient speakers generalising to novel forms sharing phonological similarities with existing forms (Prasada & Pinker, 1993; Albright & Hayes, 2003). Thus, it remains unclear how the underlying representations of phonologically consistent and ambiguous forms come to be represented shortly after learning and how overnight consolidation may influence these representations. Relatedly, the only neuroimaging experiment that has considered the role of consolidation in morphological learning and generalisation has not considered the role of type and token frequency and phonological consistency (Leminen et al., 2016).

One previous behavioural study has examined the role of overnight consolidation in modulating the influence of type and token frequency as well as phonological consistency on the generalisation of novel morphology. Mirković and Gaskell (in prep.) trained participants on an artificial language containing 3 novel affixes (e.g. *-aff*, *-eem*, *-esh*), which combined with 18 novel singular nouns (e.g. *nork*) to form the plural forms (e.g. *norkaff*) of familiar objects (e.g. *lemons*). This novel morphological system was designed to mimic the type and token frequency and phonological consistency properties of the English past tense. The majority of nouns (12) took a high type frequency regular affix (e.g. *-aff*) whilst a minority of nouns (6) took one of two low type frequency irregular affixes (e.g. *-esh*, *-eem*). A subset of nouns (9) was phonologically diverse (e.g. *nork*, *plass*, *thilt*) and associated with the regular affix (e.g. *norkaff*, *plassaff*, *thiltaff*) similar to phonologically varied regular English verbs (e.g. *called*, *turned*, *played*). Another subset of nouns (3) contained a

phonological cue (e.g. *-isp*; *jisp*, *slisp*, *tisp*), which was consistently associated with one of the irregular affixes (e.g. *-isp*; *jispeem*, *slispeem*, *tispeem*) analogous to phonologically consistent irregular English verbs (e.g. *clung*, *flung*, *stung*). A final subset of stems (6) contained an ambiguous phonological cue (e.g. *-arb*; *farb*, *clarb*, *yarb*, *harb*, *blarb*, *varb*), which was associated with the regular affix for three nouns (e.g. *-arb*; *farbaff*, *clarbaff*, *yarbaff*) and with the second irregular affix for the other three nouns (*-arb*; *harbesh*, *blarbesh*, *varbesh*) mimicking phonologically ambiguous English verbs (e.g. *show-showed*; cf. *blow-blew*). Mimicking the type – token frequency distinction between English regular and irregular verbs, irregular plurals were presented with higher token frequency than regulars during training. Specifically, each irregular plural was presented 24 times during training whilst each regular plural was presented 6 times. The number of affix presentations was kept the same across regular and irregular plurals such that only their relative type and token frequencies were manipulated (Figure 4).

a)

Phonological consistency	Type/Token frequency	Singular nouns	Plural nouns
Diverse	High type/ Low token (regular)	nork	norkaff
		plass	plassaff
		thilt	thiltaff
		heef	heefaff
		groll	grollaff
		shiln	shilnaff
		dowth	dowthaff
Ambiguous	High type/ Low token (regular)	farb	farbaff
		clarb	clarbaff
		varb	varbaff
	High token/ Low type (irregular)	harb	harbesh
		blarb	blarbesh
		varb	varbesh
Consistent	High token/ Low type (irregular)	jisp	jispeem
		slisp	slispeem
		tisp	tispeem

b)

Type frequency	Token frequency	Affix frequency
12	6	72
3	24	72
3	24	72

c)

Figure 4. Mirković and Gaskell (in prep.): training stimuli set

a) example training stimuli set used by Mirković et al. (in prep.): phonologically diverse and phonologically ambiguous novel plurals taking a regular affix are in blue, phonologically ambiguous novel plurals taking an irregular affix are in green, and phonologically consistent novel plurals taking an irregular affix are in red. b) the type, token, and affix (type x token) training frequencies associated with each plural affix. c) example novel singular and plural nouns associated with a familiar object.

After training participants on this new morphological system, Mirković and Gaskell (in prep.) examined the degree to which type and token frequency and phonological consistency influenced generalisation and whether these influences

changed with overnight consolidation. A set of 18 novel singulars was created. Six were phonologically diverse (e.g. *jeech*), 6 were phonologically consistent (e.g. *zisp*), and 6 were phonologically ambiguous (e.g. *narb*). Participants were presented with these novel singulars and asked to supply what they considered to be appropriate plural forms. To assess consolidation-related changes in generalisation, one group of participants were tested immediately after plural training whilst another group of participants were tested after a 24hr delay (Figure 5). When presented with phonologically diverse novel singulars, both groups of participants produced mostly regular generalisations (e.g. *jeechaff*). When presented with phonologically consistent novel singulars, both groups of participants produced mostly irregular generalisations (consistent) (e.g. *zispem*). When presented with phonologically ambiguous novel singulars, participants tested after a 24hr delay produced more irregular generalisations (ambiguous) than regular generalisations (e.g. *narbesh* > *narbaff*). By contrast, participants tested immediately after training produced a similar number of irregular generalisations (ambiguous) and regular generalisations (e.g. *narbesh* \approx *narbaff*) (Figure 6).

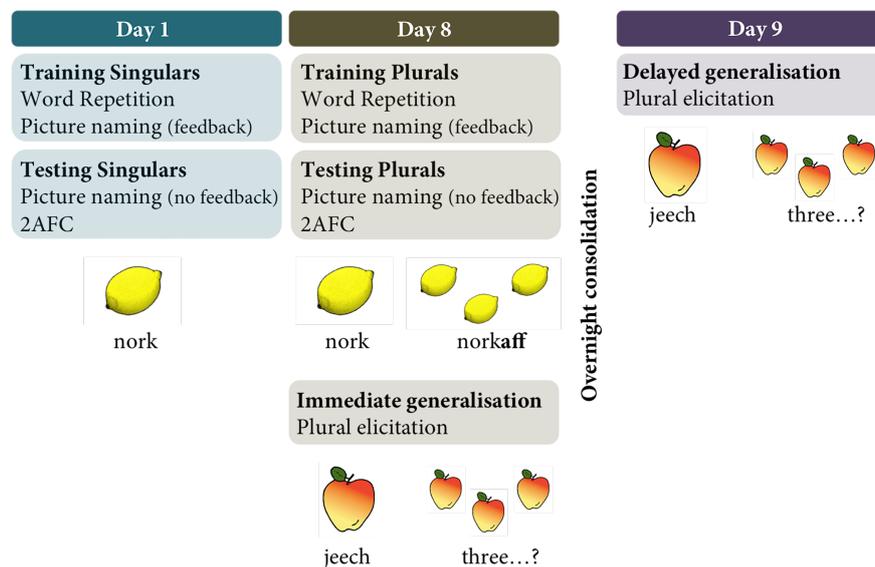


Figure 5. Mirković and Gaskell (in prep.): experimental design

Participants were trained on a novel morphological system over 9 days in a between-subject design. Both groups were trained on the singulars on Day 1 and on the plurals of Day 8. One group performed the generalisation task immediately after training (immediate generalisation) and the other group performed the generalisation task on Day 9 after a 24hr delay containing sleep (delayed generalisation).

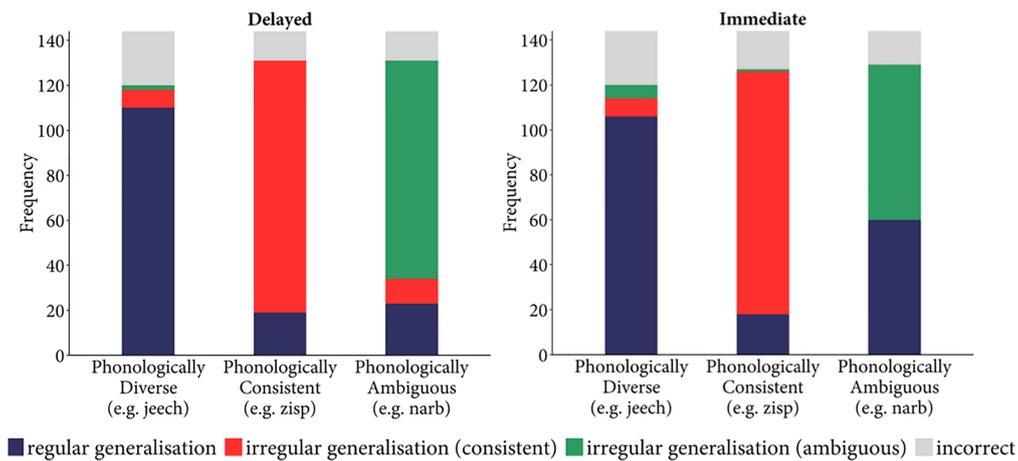


Figure 6. Mirković and Gaskell (in prep.): generalisation results

Frequency of regular and irregular (consistent and ambiguous) generalisations in response to novel phonologically diverse, consistent, or ambiguous singulars produced by participants tested immediately after training (right panel) or after a delay of 24hr including sleep (left panel).

These results suggest that both groups of participants exploited the type and token frequency and phonological consistency properties of the training set to derive knowledge that they could generalise appropriately to novel singulars. The regular affix was applied preferentially to novel singulars that were phonologically diverse (e.g. *jeechaff*) consistent with network simulations suggesting that high type frequency and phonological diversity play an important role in supporting the development of default generalisation (Plunkett & Marchman, 1991, 1993; Hare & Daugherty, 1993; Hare et al., 1995). Participants' willingness to apply the appropriate irregular affix to novel phonologically consistent nouns (e.g. *zisp_{eeem}*) is in line with previous evidence of sensitivity to islands of reliability and high token frequency acting as a protective factor against regularisation (Plunkett & Marchman, 1991; Bybee, 1995; Albright & Hayes, 2003). There was no evidence for overnight changes in these two generalisation behaviours. For novel phonologically ambiguous nouns, participants tested immediately after training produced a mixture of regular (e.g. *narbaff*) and irregular (e.g. *narbesh*) generalisations whilst participants tested after a 24hr delay including sleep volunteered mostly irregular generalisations (e.g. *narbesh*). Albright and Hayes (2003) reported that novel regular verbs that were phonologically similar to existing irregulars were rated lower than novel regular verbs that were not phonologically similar to existing irregulars. In other words, novel regulars were rated lower if existing regulars and irregulars competed to influence participants' ratings. These trade-offs effects are consistent with similar forms competing within a single network of connections. In Mirković and Gaskell's (in prep.) study, participants tested

immediately after plural training produced a similar number of regular generalisations and irregular generalisations (ambiguous) in line with these competition effects. However, participants tested after a 24hr delay produced more irregular generalisations (ambiguous). These results suggest a role for overnight consolidation in modulating competition effects between regular and irregular mappings such that high token frequency irregular mappings are strengthened and preferentially generalised.

Why might overnight consolidation strengthen high token frequency irregular mappings for phonologically ambiguous plurals? Mirković and Gaskell (in prep.) suggest that their results can be explained within a CLS framework. They suggest that systematic mappings between stems and affixes, as is the case for phonologically diverse plurals and phonologically consistent plurals, may be rapidly encoded neocortically such that overnight consolidation may not provide any further benefit. In the case of arbitrary mappings between stems and affixes, as is the case for phonologically ambiguous plurals, learning may initially be more dependent on hippocampal pattern separation and subsequently more prone to overnight consolidation. Mirkovic & Gaskell (2016) report findings that go some way to supporting this proposal. Participants were trained on an artificial grammatical system containing arbitrary stem – referent mappings (e.g. *scoiff* = ballerina, *heef* = priest) and systematic determiner/affix – natural gender mappings (e.g. *tib*_[fem] *scoiffesh*_[fem] = ballerina, *ked*_[masc] *heefaff*_[masc] = priest). After training on this new grammatical system, participants were tested on their knowledge of the arbitrary and systematic mappings. One group of participants were tested after a 2hr nap whilst another group of participants were tested after a 2hr period of wakefulness. Participants in the nap group outperformed participants in the wakefulness group on tests assessing memory for the arbitrary mappings whilst both groups performed equally well on tests assessing knowledge of the systematic mappings. Mirković and Gaskell (2016) suggest that these results are consistent with the preferential involvement of sleep in prioritising the most hippocampally-reliant components over the multi-item generalisations that can also be supported by the neocortical system (McClelland et al., 1995; Stickgold & Walker, 2013).

The aim of Experiment 1 was to build upon and extend Mirković and Gaskell's (in prep.) design by developing a within-subject design better suited for the functional imaging study reported in Experiment 2. To this end, the experimental design used by Davis et al. (2009) in which participants were trained on two sets of novel words on two consecutive days and tested only once on the second day was adapted to examine the effect of overnight consolidation on the generalisation of newly-learned inflected words. Two sets of plural affixes, distinguished by grammatical gender, were trained on two consecutive days following the same type and token frequency and phonological consistency manipulations as in Mirković and Gaskell (in prep.). A plural elicitation task assessing generalisation to novel forms was administered following training of the second set of plural affixes. Consistent with the CLS model, differences in generalisation behaviour are predicted for generalisations requiring the use of plural affixes tested following a 24hr delay containing sleep and those requiring the use of plural affixes tested immediately after training. In line with the results obtained by Mirković and Gaskell (in prep.), when generalising to novel phonologically ambiguous singulars, participants are predicted to produce more irregular generalisations (i.e. based on token frequency) compared to regular generalisations (i.e. based on type frequency) for plural affixes having undergone overnight consolidation compared to those trained immediately prior to generalisation.

5.2. Materials and methods

5.2.1. Participants

Eighteen participants (6 males) aged between 18 and 29 (mean age = 23, SD = 3) were recruited from the MRC Cognition and Brain Sciences Unit Participant Panel and provided their informed consent to take part in the study. Participants were tested under the approval of the Cambridge Psychology Research Ethics Committee. All were monolingual speakers of British English with little or no knowledge of a second language, no known hearing or language impairments, and no neurological or psychiatric disorders. Participants were paid to take part in the study.

5.2.2. Experimental stimuli

5.2.2.1. Training stimuli

Building upon and extending the stimuli set developed by Mirković and Gaskell (in prep.), a novel morphological system was implemented in an artificial language in which novel inflected words referred to the occupation of single or multiple male and female characters. Two sets of 18 novel stems (e.g. set 1: *gleet*; set 2: *shiln*) were generated from the ARC nonword database (Rastle, Harrington, & Coltheart, 2002). All were pronounceable monosyllabic monomorphemic English nonwords and 4 to 5 phoneme long. For set 1, each novel stem was combined with a novel affix (e.g. *-i*) to create singular nouns referring to the occupation of single female characters (e.g. *gleeti*_[fem:sing] = *doctor*_[fem:sing]). For set 2, each novel stem was combined with another novel affix (e.g. *-u*) to create singular nouns referring to the occupation of single male characters (*shilnu*_[masc:sing] = *painter*_[masc:sing]). For set 1, each novel stem was combined with one of three novel affixes (*-aff*, *-imm*, *-esh*) to create plural nouns referring to three female characters depicting the corresponding occupation (e.g. *gleetaff*_[fem:plur] = *doctor*_[fem:plur]). For set 2, each novel stem was combined with one of three novel affixes (*-opp*, *-oot*, *-ull*) to create plural nouns referring to three male characters depicting the corresponding occupation (e.g. *shilnopp*_[masc:plur] = *painters*_[masc:plur]) (Figure 7). Thus, like many Indo-European languages, the new morphological system contained portmanteau affixes, which simultaneously specified number and gender information (Hockett, 1954). For instance, Italian expresses both gender and number simultaneously by single affixes (e.g. *mela*_[fem:sing] = *apple*, *mele*_[fem:plur] = *apples*). Each training set contained the same phonological consistency and type and token frequency manipulations as in Mirković and Gaskell (in prep.) with some minor alterations to the phonological cues. Twelve lists were created to partially counterbalance the assignment of the plural affixes to regular and irregular conditions and the assignment of phonological cues (i.e. *arb*, *isk*) to phonologically ambiguous and consistent conditions. The assignment of each set of plural affixes (*-aff*, *-imm*, *-esh*; *-opp*, *-oot*, *-ull*) to each gender was counterbalanced across participants.

a)

Phonological consistency	Type/Token frequency	Training set 1		Training set 2	
		Singular nouns	Plural nouns	Singular nouns	Plural nouns
Diverse	High type/ Low token (regular)	gleeti	gleetaff	shilnu	shilnopp
		torthi	torthaff	plassu	plassopp
		dulti	dultaff	blornu	blornopp
		vonti	vontaff	hulfu	hulfopp
		zolli	zollaff	rutchu	rutchoff
		mowli	mowlaff	grollu	grollopp
		sleni	slenaff	crephu	crephopp
		chiffi	chiffaff	thiltu	thiltopp
Ambiguous	High type/ Low token (regular)	harbi	harbaff	jarbu	jarbopp
		yarbi	yarbaff	larbu	larbopp
		narbi	narbaff	glarbu	glarbopp
	High token/ Low type (irregular)	tarbi	tarbimm	blarbu	blarboot
		clarbi	clarbimm	marbu	marboot
		slarbi	slarbimm	farbu	farboot
Consistent	High token/ Low type (irregular)	hiski	hiskesh	tisku	tiskull
		liski	liskesh	visku	viskull
		fiski	fiskesh	visku	viskull

b)

Type frequency	Token frequency	Affix frequency
12	6	72
3	24	72
3	24	72

c)

Figure 7. Experiment 1: training stimuli set

a) example training stimuli set: phonologically diverse and phonologically ambiguous novel plurals taking a regular affix are in blue, phonologically ambiguous novel plurals taking an irregular affix are in green, and phonologically consistent novel plurals taking an irregular affix are in red. b) the type, token, and affix (type x token) training frequencies associated with each plural affix. c) example novel singular and plural nouns associated with familiar occupations for female and male characters.

The singular affix *-i* was always associated with the *-aff*, *-imm*, and *-esh* plural affixes. Similarly, the singular affix *-u* was always associated with the *-opp*, *-oot*, and *-ull* plural affixes. These groupings were designed to keep the onset vowels of each set of plural affixes as far apart as possible but similarly distributed in the vowel space to facilitate learning and reduce possible gender confusions (Figure 8). The assignment of the two sets of plural affixes to female or male characters was counterbalanced across participants. A male native speaker of southern British English recorded spoken forms of the novel singulars and plurals in a soundproof booth at a sampling rate of 44.1 kHz. The recording was divided into single audio files, one for each novel word, which were trimmed to length and normalised to equate their intensities. Each singular noun was paired with the picture of a single male or female character depicting a familiar occupation and each plural noun was paired with the picture of three male or female characters depicting the corresponding occupation⁹. The pictures were downloaded from a stock image website (www.dreamstime.com) after purchasing a Royalty-Free license, edited and resized where necessary. There was no phonological overlap between the novel nouns and the

⁹ These were pre-tested in an independent sample of 10 participants to ensure there was complete agreement on the gender and occupation of each character included in the stimuli set.

*shilnu*_[masc:sing] = *painter*_[masc:sing]). On Day 8, one week after learning the singulars, participants were trained on the first set of plural affixes by learning the plural forms of half of the singulars trained on Day 1 for one gender only. This session lasted approximately 1hr15. On Day 9, approximately 24 hours after learning the first set of plural affixes, participants were trained on the second set of plural affixes by learning the plural forms of the other half of singulars trained on Day 1 for the opposite gender to Day 8. Participants then took a short 10-minute break, which included a walk in the vicinity of the testing room, before completing the generalisation task in which they were asked to supply what they considered to be appropriate plural forms for new feminine (e.g. *woathi*_[fem:sing]) and masculine (e.g. *woathu*_[masc:sing]) singular nouns (Figure 9). This last session lasted approximately 2hr. All participants were tested individually in the same testing cubicle for all 3 sessions. They sat comfortably at approximately 60cm of a 17-inch computer monitor. Spoken words were delivered via high-quality headphones and participants' verbal responses were audio recorded. All tasks were implemented in E-prime 2.0 (Psychology Software Tools Inc., Pittsburgh, PA).

Four important features of this design should be noted. First, it builds upon and extends the morphological system developed by Mirković and Gaskell (in prep.) by instantiating the morphosyntactic feature of gender in order to distinguish between two sets of affixes. This allowed training two sets of *different* affixes, distinguished by gender, on two successive days. Effects of initial learning (Day 9 affixes) and overnight consolidation (Day 8 affixes) could thus be compared in a within-subject (i.e. between-affix) design better suited for the fMRI study reported in Chapter 8. Second, effects of initial learning (Day 9 affixes) and overnight consolidation (Day 8 affixes) could be assessed within a single generalisation plural elicitation task administered on Day 9 comparing generalisations that required participants to use Day 8 or Day 9 affixes. This design is similar to the one used by Davis et al. (2009) in the context of word learning. As a consequence of testing generalisation only once, it controls for potential practice effects that would arise from administering several generalisation tests to the same participants. It is also an efficient design for the functional imaging study reported in Experiment 2. This is discussed in more detail in Chapter 8. Third, each set of plural affixes was trained using a different set of stems, both trained on Day

1, to avoid introducing interference from onset overlapping competitors when training the second set of plural affixes on Day 9 (e.g. *gleetaff*_[fem:plur] = *doctors*_[fem:plur] and *gleetopp*_[masc:plur] = *doctors*_[masc:plur] were not trained in the same participants). Fourth, singulars were trained on Day 1 such that both sets of stems were learned and consolidated prior to training the plurals. Based on the evidence concerning the consolidation and lexicalisation of novel words reviewed in Chapter 2, training singulars on Day 1 and allowing a week to pass before training the plurals should be sufficient for their consolidation (Davis & Gaskell, 2009). Consolidation-related changes in generalisation behaviour could thus be attributed more confidently to changes in the representation of the affixes rather than the stems.

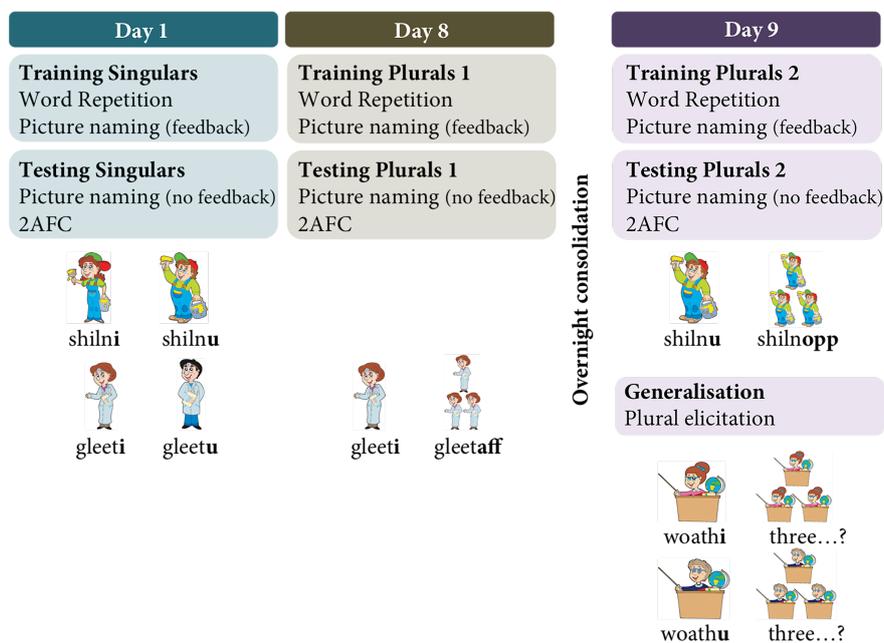


Figure 9. Experiment 1: experimental design

Participants were trained on a novel morphological system over 9 days in a within-subject design. Participants were trained on the singulars on Day 1. They returned to the lab one week later to be trained on one set of plurals for one gender on Day 8. They were trained on a second set of plurals for the other gender on Day 9 following a 24hr delay containing sleep. They then performed a generalisation task, which required them to use plural affixes trained immediately before or on the previous day.

5.2.3.1. Day 1

Participants completed 4 blocks of training on the singulars. Each training block comprised 1) a word repetition task followed by 2) a picture naming task with feedback. Each singular was presented once per block in each task for a total of 8 presentations over the course of all training blocks. Memory for the singulars was assessed by 1) a picture naming without feedback task followed by 2) a two-alternative

forced choice (2AFC) recognition memory task in a testing block presented immediately after the last training block (Figure 10).

Word repetition task. Each trial began with a black fixation cross presented for 500ms in the centre of a white background, which was followed by the auditory presentation of a novel singular. The picture of the character corresponding to the spoken word together with the written form of the word presented below the picture were displayed in the centre of the screen 300ms after the onset of the spoken word and remained on the screen for 4000ms. Participants were instructed to repeat the word out loud. There was an inter-trial interval of 500ms. Each singular was presented once per block. The order of presentation of the trials was randomised for each participant.

Picture naming with (without) feedback. Each trial started with a black fixation cross presented for 500ms in the centre of a white background, which was replaced by the target picture which remained on the screen for 4500ms. Participants were instructed to name the picture using the novel words they had learned in the previous training block. After 4000ms, the correct spoken word was played over the headphones. There was an inter-trial interval of 500ms. Each target picture and its associated spoken word were presented once per block. The order of presentation was randomised for each participant. The picture naming task without feedback used in the testing block was exactly the same but participants did not hear the correct response at the end of each trial. For both training and test picture naming tasks, participants' verbal responses were audio recorded.

2AFC recognition memory. Each trial started with a black fixation cross presented for 500ms in the centre of a white background, which was followed by the auditory presentation of one of the novel singulars. Two pictures were presented on each side of the screen 300ms after the onset of the spoken word and remained on the screen until participants provided a response. There was an inter-trial interval of 500ms. Participants were instructed to press “z” on the keyboard to indicate that the spoken word corresponded to the picture on the left and “m” to indicate that it corresponded to the picture on the right. The two pictures were always of the same gender. The character corresponding to each word was presented once as a target and once as a

distractor. The position of the target picture on the screen was counterbalanced across trials. The order of presentation of the trials was randomised for each participant.

5.2.3.1. Day 8

Participants learned the first set of plural affixes for one gender (Figure 9). First, participants completed one top-up training block on the singulars alone comprising 1) a word repetition task and 2) a picture naming task with feedback as described above. Each singular was presented 6 times during the word repetition task and once in the picture naming task. Participants then performed 3 training blocks including both the singulars trained in the first block and their corresponding plurals. Each plural training block comprised 1) a word repetition task and 2) a picture naming task with feedback as described for Day 1. During the word repetition task, each singular was presented twice per training block and the plurals were presented following the token frequencies presented in Figure 7. The regulars were presented a total of 6 times over the 3 training blocks (x2 per block) while each irregular was presented a total of 24 times over the 3 training blocks (x8 per block). Memory for the singulars and plurals was assessed by 1) a picture naming without feedback task followed by a 2) 2AFC recognition memory task in a testing block presented immediately after the last plural training block. The 2AFC recognition memory task was the same as the one described for Day 1 but this time included both singulars and plurals. The two pictures were always of the same gender and number. The character corresponding to each word was presented once as a target and once as a distractor. The position of the target picture on the screen was counterbalanced across trials. The order of presentation of the trials was randomised for each participant.

5.2.3.1. Day 9

Participants learned the second set of plural affixes for the opposite gender to the one trained on Day 8 (Figure 9). The training tasks were the same as for Day 8. The order in which the two genders were trained across the two plural training days (Day 8, Day 9) was counterbalanced across participants.

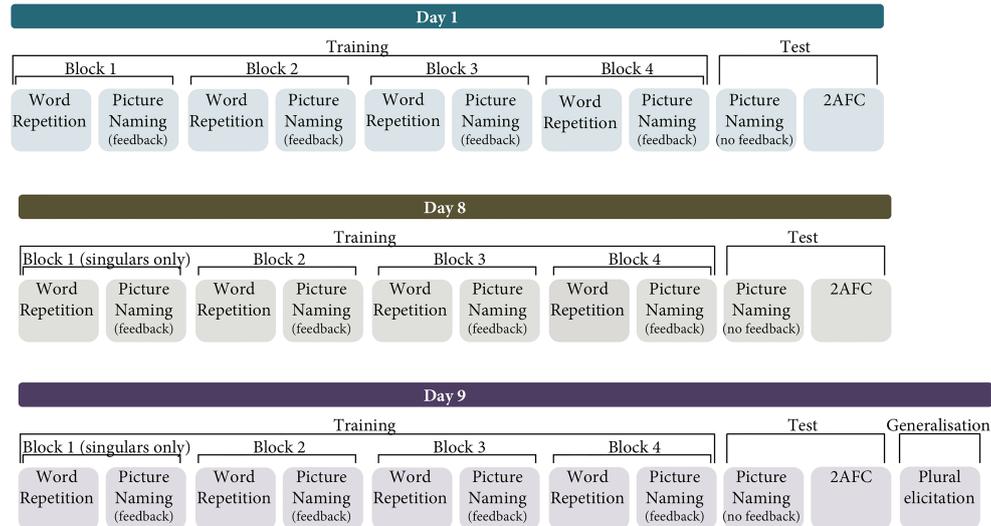
Generalisation. Each trial began with a black fixation cross presented for 500ms in the centre of a white background, which was followed by the auditory presentation of a novel singular word not previously presented in training. The picture of the

corresponding single character with the phrase ‘one [novel wordsing]’ (e.g. ‘one woath_i’) displayed underneath it was presented in the centre of the screen 300ms after the onset of the spoken word and remained on the screen for 2000ms. A blank screen was then presented for 500ms, which was followed by the presentation of the same character depicted as a triplet with the phrase ‘three...?’ underneath it. This remained on the screen for 4500ms. Participants were instructed to say out loud what they thought was the most appropriate plural form of the new singular. There was an inter-trial interval of 500ms. The experimenter checked with the participant to make sure they understood the instructions and went through a practice trial offline together. Each generalisation item was presented once and the order of presentation of the trials was randomised for each participant.

5.2.4. Data analysis

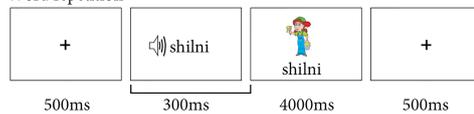
Data were analysed by fitting generalised linear models using the *glm2* package (Marschner, 2011) in R version 3.3.1 (R Development Core Team, 2016). Accuracy on the picture naming and 2AFC recognition memory tasks were analysed using logistic regression with a binomial distribution since participants could either produce a correct or an incorrect response. Logistic regression is more appropriate and less prone to type I or type II errors than applying an analysis of variance (ANOVA) on proportional data even if the proportional data is arcsine transformed (Jaeger, 2008). For the picture naming tasks, only responses where participants produced the appropriate stem and affix (e.g. *gleet* + *i* for *gleet*_{i(fem:sing)}) were deemed correct. For the plural elicitation generalisation task, data were analysed using loglinear regression with a Poisson distribution applied to the frequency counts of each response type. Loglinear regression is also less likely to result in type I or type II errors than an ANOVA applied to proportional data, even if the proportional data is arcsine transformed (Agresti, 2002; Jaeger, 2008). Loglinear regression also offers advantages over the most frequently used Chi-square tests used for the analysis of language acquisition data including its applicability to complex designs with multiple variables of the kind used in this thesis (Li, 2002). For the primary analysis of the plural elicitation generalisation task, responses in which participants correctly reproduced the stem of a novel singular (e.g. *woath* for *woath*_{i(fem:sing)}) and inflected it with a gender-appropriate trained plural affix were deemed correct.

a) Experimental procedure

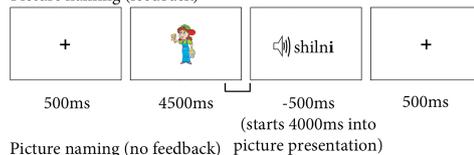


b) Experimental tasks

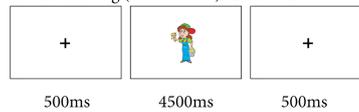
Word repetition



Picture naming (feedback)



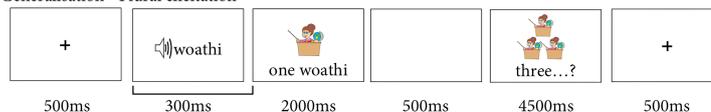
Picture naming (no feedback)



2AFC



Generalisation - Plural elicitation

**Figure 10. Experiment 1: experimental procedure and tasks**

a) summary of the experimental procedure. On Day 1, participants completed 4 blocks of training following by a test block on all singulars. On Day 8, participants were trained on the plurals for one gender. They began by completing one top-up block on the singulars only followed by three training blocks on the singulars and plurals and a test block on the singulars and plurals. The same procedure was followed on Day 9 for the plurals of the other gender. Additionally, participants completed a generalisation task after the testing block. b) summary of all experimental tasks: word repetition, picture naming with and without feedback, 2AFC, and generalisation plural elicitation task.

These correct responses were classified as regular generalisations (e.g. *woathaff*_[fem:plur]), irregular generalisations (consistent) (e.g. *woathimm*_[fem:plur]), and irregular generalisations (ambiguous) (e.g. *woathesh*_[fem:plur]) on the basis of the plural affix used. All other responses were classified as incorrect. This approach was preferred to

analysing *accurate* responses (e.g. producing a regular generalisation, *woathaff*_[fem:plur], in response to a phonologically diverse singular, *woathi*_[fem:singl]) to examine differences in the use of different types of generalisations as a function of the training day of the novel affixes and the phonological consistency of the novel singulars. In all analyses, the terms phonologically diverse, phonologically consistent, and phonologically ambiguous are used to refer to both singulars and plurals. It should be noted that this classification refers to the type of phonological cues contained in the stems.

A model simplification approach was followed in building the generalised linear models. Full models with all the predictor variables and their interactions were considered first. The significance of each predictor was evaluated by carrying out likelihood ratio tests (LRT) comparing a model including a predictor of interest to an identical model excluding the predictor of interest (Barr, Levy, Scheepers, & Tily, 2013). Elimination of non-significant predictors began with the highest order interactions moving backwards to lower order interactions and main effects. For LRT, Chi Square and p values are reported. The statistics for non-significant effects are not reported in the text, except if there are marginally significant ($p \leq .06$). For significant predictors ($p < .05$), pairwise comparisons were performed using the *multcomp* package (Hothorn, Bretz, & Westfall, 2008). For these, the estimated coefficient (b), the z statistic associated with the coefficient, and the p-value based on the z statistic are reported. The p-values reported in the text are uncorrected for multiple comparisons. However, Bonferroni corrections were also applied based on the number of contrasts examined for each significant predictor. Uncorrected p-values that did not survive Bonferroni correction are marked with the symbol †. This approach provides information about the significance level of each contrast as well as its robustness to multiple comparisons. The same data analysis approach is used for similar tasks in experiments presented in subsequent chapters.

5.3. Results

5.3.1. Training

Accuracy on the picture naming tasks presented at the end of each training block in each training session was analysed to examine how participants' mastery of the new words improved over the course of training (Figure 11). Singulars and plurals were

analysed separately as there was no plural training on Day 1, there were only 3 training blocks for plurals on Day 8 and Day 9 (compared to 4 blocks for singulars), and because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix.

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8, Day 9), training block (Block 1, Block 2, Block 3, Block 4), and their interaction was fitted. The three-way interaction between condition, training day, and training block was not significant. The two-way interaction between condition and training block was significant, $\chi^2_{(6)} = 16.72$, $p = .010$. Accuracy increased significantly over the course of training blocks for phonologically diverse singulars (Block 2 vs. Block 1: $b = 0.71$, $z = 7.58$, $p < .001$, Block 3 vs. Block 2: $b = 0.68$, $z = 6.70$, $p < .001$, Block 4 vs. Block 3: $b = 0.62$, $z = 5.25$, $p < .001$) and for phonologically ambiguous singulars (Block 2 vs. Block 1: $b = 0.96$, $z = 8.23$, $p < .001$, Block 3 vs. Block 2: $b = 0.59$, $z = 4.69$, $p < .001$, Block 4 vs. Block 3: $b = 0.53$, $z = 3.61$, $p = .003$), but did not increase further after the third training block for phonologically consistent singulars (Block 2 vs. Block 1: $b = 0.97$, $z = 5.88$, $p < .001$, Block 3 vs. Block 2: $b = 1.14$, $z = 5.77$, $p < .001$). The two-way interaction between training block and training session was significant, $\chi^2_{(6)} = 50.10$, $p < .001$. Participants improved significantly across training blocks on Day 1 (Block 2 vs. Block 1: $b = 1.48$, $z = 15.50$, $p < .001$, Block 3 vs. Block 2: $b = 0.86$, $z = 10.30$, $p < .001$, Block 4 vs. Block 3: $b = 0.60$, $z = 6.42$, $p < .001$) but not on Day 8 or Day 9. The two-way interaction between condition and training session was marginally significant, $\chi^2_{(4)} = 9.33$, $p = .053$. For phonologically ambiguous singulars, participants were marginally more accurate on Day 9 compared to Day 8 (Day 9 vs. Day 8: $b = 0.60$, $z = 1.92$, $p = .060$) but were equally accurate across Day 8 and Day 9 for phonologically diverse and phonologically consistent singulars. The main effect of condition was significant, $\chi^2_{(2)} = 8.25$, $p = .020$. Participants were more accurate on phonologically consistent than phonologically diverse ($b = 0.24$, $z = 2.86$, $p = .004$) and phonologically ambiguous ($b = 0.19$, $z = 2.06$, $p = .040^+$) singulars but performed similarly on phonologically diverse and phonologically ambiguous singulars. The main effect of training day was significant, $\chi^2_{(2)} = 1141.23$, $p < .001$. Participants were more accurate on Day 8 compared to Day 1 ($b = 3.04$, $z = 24.06$, $p < .001$) and on Day

9 compared to Day 1 ($b = 3.02, z = 24.13, p < .001$) but performed similarly on Day 8 and Day 9. The main effect of training block was significant, $\chi^2_{(3)} = 1173.84, p < .001$. Participants' accuracy increased significantly over the course of training blocks (Block 2 vs. Block 1: $b = 1.23, z = 14.80, p < .001$, Block 3 vs. Block 2: $b = 0.86, z = 10.62, p < .001$, Block 4 vs. Block 3: $b = 0.61, z = 6.74, p < .001$).

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), training block (Block 1, Block 2, Block 3, Block 4), and their interaction was fitted. Note that phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular affix. The three-way interaction between condition, training day, and training block was not significant. None of the two-way interactions were significant. The main effect of condition was significant, $\chi^2_{(3)} = 27.29, p < .001$. Participants were significantly more accurate on phonologically diverse than on phonologically ambiguous (regular) ($b = 0.86, z = 5.02, p < .001$) and phonologically ambiguous (irregular) plurals ($b = 0.51, z = 2.80, p = .005$). Participants were also more accurate on phonologically consistent than on phonologically ambiguous (regular) plurals ($b = 0.73, z = 3.36, p < .001$). The main effect of training block was also significant, $\chi^2_{(2)} = 83.08, p < .001$. Participants' accuracy improved significantly across training blocks (Block 2 vs. Block 1: $b = 0.83, z = 5.45, p < .001$, Block 3 vs. Block 2: $b = 0.66, z = 3.37, p < .001$). The main effect of training day was not significant.

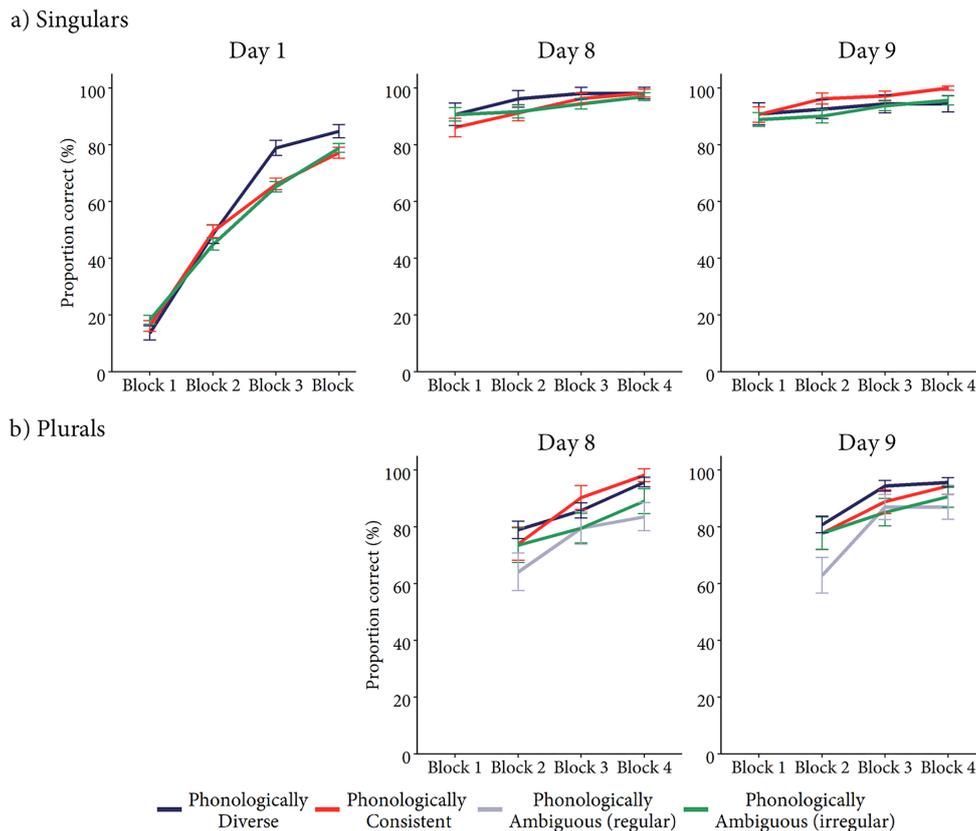


Figure 11. Experiment 1: training accuracy

Mean accuracy on the picture naming task presented at the end of each training block (Block 1, Block 2, Block 3, Block 4) on each training day (Day 1, Day 8, and Day 9) for a) singulars and b) plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (light blue line) and those associated with an irregular (green line) affix. All phonologically ambiguous singulars are depicted by a green line.

5.3.2. Testing

Accuracy on the picture naming (Figure 12) and 2AFC recognition memory (Figure 13) tasks presented in a testing block at the end of each training session was analysed to ensure that words with different phonological consistencies and training frequencies (plurals) were learned equally well. Singulars and plurals were analysed separately as there was no plural training on Day 1, there were only 3 training blocks for plurals on Day 8 and Day 9 (compared to 4 blocks for singulars), and because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix.

5.3.2.1. Picture naming

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8,

Day 9), and their interaction was fitted. There was no significant interaction between condition and training day. The main effect of condition was significant, $\chi^2_{(2)} = 7.79$, $p = .020$. Participants were significantly more accurate on phonologically consistent compared to phonologically diverse ($b = 0.51$, $z = 2.19$, $p = .028^+$) and phonologically ambiguous singulars ($b = 0.64$, $z = 2.66$, $p = .008$) but performed similarly on phonologically diverse and phonologically ambiguous singulars. The main effect of training day was also significant, $\chi^2_{(2)} = 116.58$, $p < .001$. Participants were significantly more accurate on Day 8 compared to Day 1 ($b = 2.78$, $z = 5.46$, $p < .001$) and on Day 9 compared to Day 1 ($b = 1.95$, $z = 5.63$, $p < .001$) but performed similarly on Day 8 and Day 9.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction was fitted. Note that phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular affix. The interaction between condition and training day was not significant. The main effect of condition was significant, $\chi^2_{(3)} = 8.17$, $p = .043$. Participants were significantly more accurate on phonologically diverse than on phonological ambiguous regular ($b = 9.29$, $z = 2.48$, $p = .013^+$) and on phonologically consistent than on phonological ambiguous regular plurals ($b = 1.35$, $z = 2.32$, $p = .021^+$). The main effect of training day was not significant.

5.3.2.2. 2AFC recognition memory

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8, Day 9), and their interaction was fitted. Condition and training day did not interact significantly. There was no significant main effect of condition or training day.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction was fitted. There was no significant interaction between condition and training day. Neither the main effect of condition nor the main effect of training day was significant.

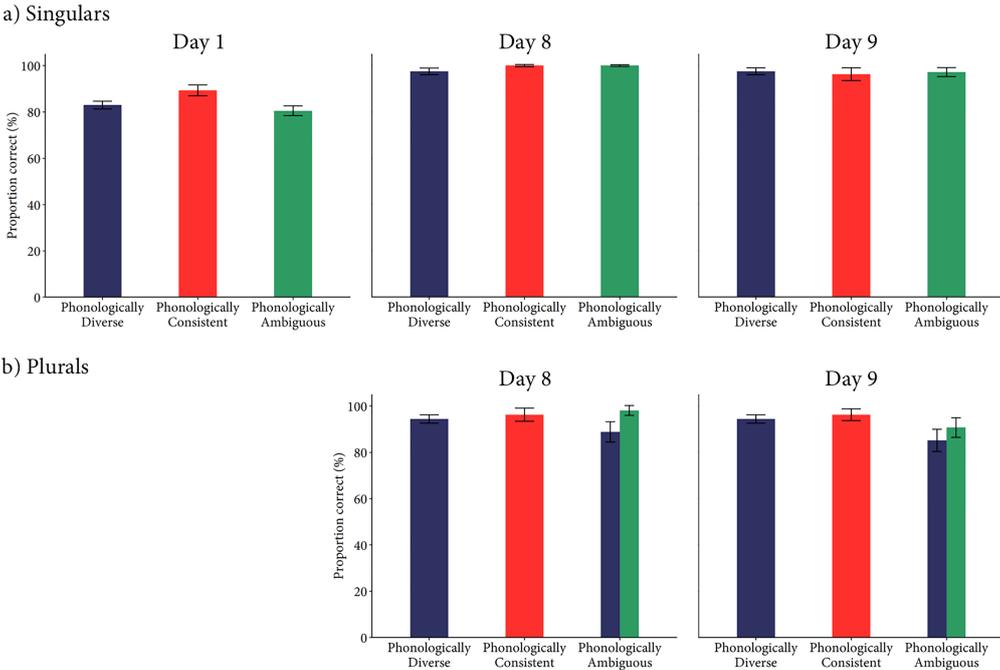


Figure 12. Experiment 1: test accuracy (picture naming)
Mean accuracy on the picture naming task presented in the testing block the end of each training day (Day 1, Day 8, Day 9) for a) singulars and b) plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (green bar) affix.

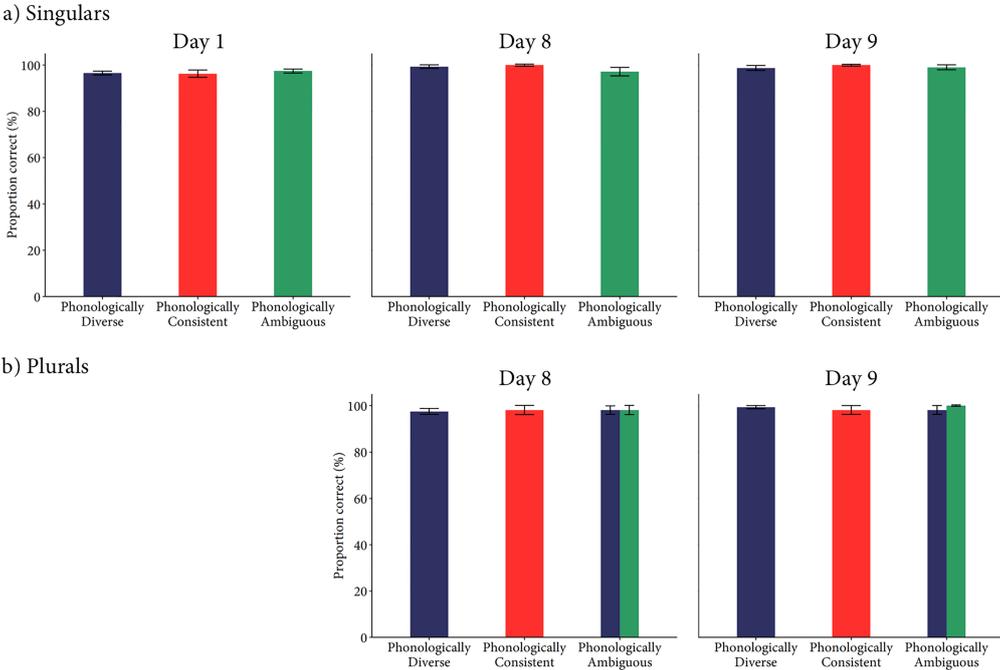


Figure 13. Experiment 1: test accuracy (2AFC)
Mean accuracy on the 2AFC recognition memory task in the test block presented at the end of each training day (Day 1, Day 8, Day 9) for a) singulars and b) plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (green bar) affix.

5.3.3. Generalisation

5.3.3.1. Response type analysis

The frequency of each response type (regular generalisations, irregular generalisations (consistent), irregular generalisations (ambiguous), incorrect) for novel phonologically diverse, consistent, and ambiguous singulars requiring participants to use Day 8 and Day 9 affixes are shown in Figure 14. A loglinear model with phonological consistency (diverse, consistent, ambiguous), affix training day (Day 8, Day 9), response type (regular generalisations, irregular generalisations (consistent), irregular generalisations (ambiguous), incorrect), and their interaction was fitted. The three-way interaction between phonological consistency, affix training day, and response type was significant, $\chi^2_{(6)} = 17.60$, $p = .007$. To break down this interaction, interactions between training day and response type were examined for each level of phonological consistency separately. For phonologically diverse novel singulars, the interaction between affix training day and response type was significant, $\chi^2_{(3)} = 47.69$, $p < .001$. Participants produced fewer regular generalisations for Day 8 compared to Day 9 affixes ($b = -3.45$, $z = -2.43$, $p = .015^+$), fewer irregular generalisations (consistent) for Day 8 compared to Day 9 affixes ($b = -1.20$, $z = -2.59$, $p = .010$), and more incorrect responses for Day 8 compared to Day 9 affixes ($b = 1.50$, $z = 5.09$, $p < .001$). For phonologically consistent novel singulars, the interaction between training day and response type was also significant, $\chi^2_{(3)} = 65.68$, $p < .001$. Participants produced fewer regular generalisations for Day 8 compared to Day 9 affixes ($b = -1.20$, $z = -4.09$, $p < p.001$) and more incorrect responses for Day 8 compared to Day 9 affixes ($b = 1.79$, $z = 5.50$, $p = < .001$). For phonologically ambiguous novel words, the interaction between training day and response type was also significant, $\chi^2_{(3)} = 48.68$, $p < .001$. Participants produced fewer regular generalisations ($b = -8.60$, $z = -4.78$, $p < .001$), more irregular generalisations (ambiguous) ($b = 5.39$, $z = 2.53$, $p = .011$), and more incorrect responses ($b = 1.01$, $z = 3.87$, $p < .001$) for Day 8 affixes compared to Day 9 affixes.

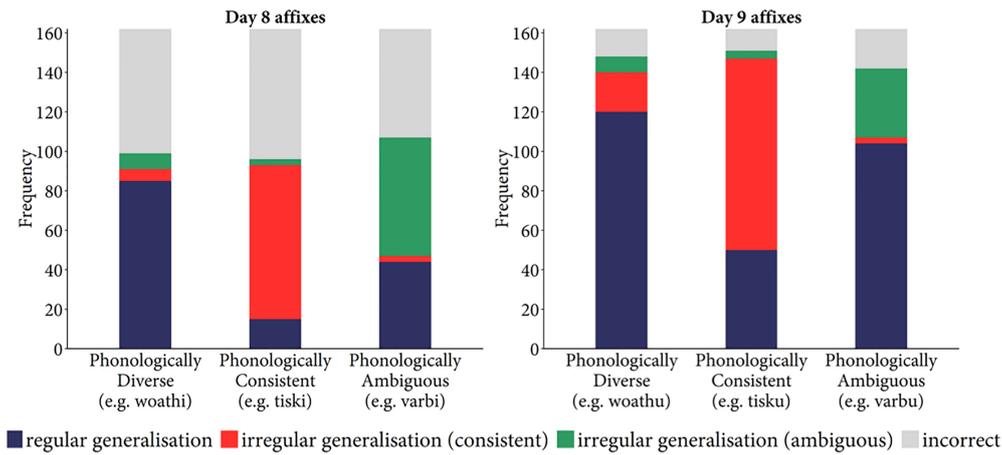


Figure 14. Experiment 1: generalisation results

Frequency of each response type supplied within each phonological consistency category of novel words requiring participants to use affixes of Day 8 (left) and on Day 9 (right).

To mitigate potential confounds due to a difference in the number of correct responses for Day 8 compared to Day 9 affixes, the analysis was repeated with a random sup-sample of correct responses for Day 9 affixes to artificially match the number of correct responses produced for Day 8 affixes (see Liu et al., 2016 for a similar approach in the context of emotional memory consolidation) (Figure 15). The three-way interaction between phonological consistency, affix training day, and response type was replicated, $\chi^2_{(4)} = 15.00$, $p = .005$. For phonologically diverse novel singulars, the two-way interaction between affix training day and response type was not significant. For phonologically consistent novel singulars, there was a significant two-way interaction between affix training day and response type, $\chi^2_{(2)} = 21.09$, $p < .001$. Participants produced fewer regular generalisations for Day 8 compared to Day 9 affixes ($b = -1.17$, $z = -3.81$, $p < .001$) and more irregular generalisations (consistent) for Day 8 compared to Day 9 affixes ($b = 3.51$, $z = 2.03$, $p = .042^\dagger$). For phonologically ambiguous singulars, the two-way interaction between affix training day and response type was significant, $\chi^2_{(2)} = 21.26$, $p < .001$. Participants produced fewer regular generalisations ($b = -0.54$, $z = -2.83$, $p = .005$) and more irregular generalisations (ambiguous) (0.82 , $z = 3.44$, $p = .001$) for Day 8 compared to Day 9 affixes (Figure 15).

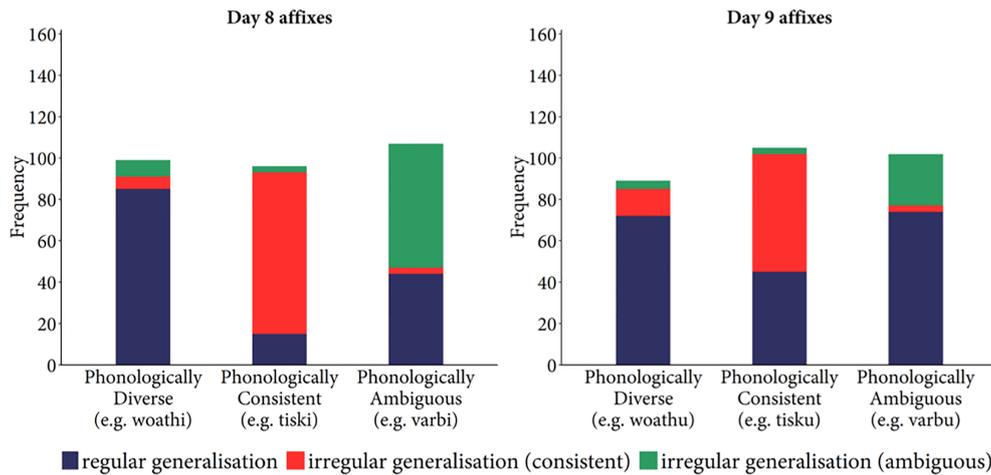


Figure 15. Experiment 1: generalisation results for equalised correct responses

Frequency of each response type supplied within each phonological consistency category of novel words requiring participants to use affixes of Day 8 (left) and on Day 9 (right).

5.3.3.2. Error analysis

A secondary analysis focused on the errors produced by participants. Incorrect responses were subcategorised as a failure to provide any response (no response), the production of a legal but gender-inappropriate affix (gender error), and the production of an illegal affix (illegal affix) (Figure 16). A loglinear model with phonological consistency (diverse, consistent, ambiguous), affix training day (Day 8, Day 9), error type (no response, gender error, illegal affix), and their interaction was fitted. There was no significant three-way interaction between phonological consistency, affix training day, and error type. The two-way interaction between affix training day and error type was significant, $\chi^2_{(2)} = 8.29$, $p = .016$. To break down this interaction, differences between the numbers of each error type were examined for Day 8 and Day 9 affixes separately. For Day 8 affixes, the main effect of error type was significant, $\chi^2_{(2)} = 106.25$, $p < .001$. Participants produced more illegal affix errors than no response errors ($b = 1.43$, $z = 6.28$, $p < .001$) and more illegal affix errors than gender errors ($b = 1.61$, $z = 9.86$, $p < .001$) but a similar number of gender errors and no response errors. For Day 9, the main effect of error type was also significant, $\chi^2_{(2)} = 10.95$, $p = .004$. Participants produced more no response errors than gender errors ($b = 0.97$, $z = 2.64$, $p = .008$) and more illegal affix errors than no response errors ($b = 1.03$, $z = 2.88$, $p = .004$) but a similar number of no response errors and illegal affix errors. The two-way interaction between affix training day and phonological consistency was not significant nor was the two-way interaction between phonological

consistency and error type. The main effect of phonological consistency was not significant.

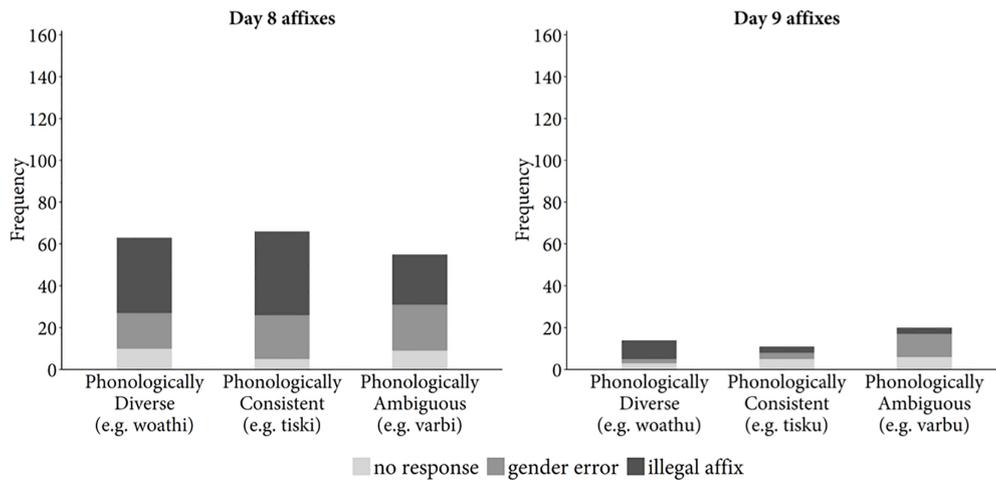


Figure 16. Experiment 1: generalisation errors

Frequency of each error type (no response, gender error, illegal affix) supplied within each phonological consistency category of novel singulars requiring participants to use affixes trained on Day 8 (left) or on Day 9 (right).

5.4. Discussion

Experiment 1 investigated the role of overnight consolidation in influencing the contribution of type and token frequency and phonological consistency to the generalisation of newly-learned plural affixes. It adapted and extended the between-subject design used by Mirković and Gaskell (in prep.) to develop a within-subject design better suited for the fMRI study reported in Experiment 2. Three experimental sessions were run over the course of 9 days. On Day 1, participants learned novel singular nouns corresponding to the occupation of single female and male characters. One week later, on Day 8, participants learned one set of plural affixes for half of the singulars trained on Day 1 for one gender. On Day 9, participants learned a second set of plural affixes for the other half of the singulars trained on Day 1 for the opposite gender to the one trained on Day 8. After plural training on Day 9, participants completed a generalisation plural elicitation task during which they were presented with novel feminine and masculine singulars and asked to supply what they considered to be appropriate plural forms. Effects of initial learning and overnight consolidation could then be assessed by comparing generalisations requiring the use of plural affixes trained on Day 8 to those requiring the use of plural affixes trained on Day 9. This design adapts the method used by Davis et al. (2009) in the context of

lexical learning and consolidation and extends it to consider morphological learning and generalisation.

Behavioural results from the training sessions confirm that the training schedule was successful in allowing participants to become highly proficient on the novel morphological system. Participants became increasingly more accurate at naming the novel singulars over the course of four training blocks on Day 1. There was evidence that learning phonologically consistent singulars (e.g. *hiski*, *hisku*) was easiest as performance plateaued at the penultimate training block. These singulars were indeed the smallest and most consistent set to learn. After completing one top-up training block on singulars on Day 8 and Day 9, participants performed with high levels of accuracy, which did not increase further across subsequent training blocks or differ between Day 8 and Day 9. These results suggest that participants had retained some knowledge of the singulars trained on Day 1 such that only one top-up training block was sufficient to yield accuracy levels comparable to those achieved at the end of training on Day 1. Overnight changes in generalisation behaviour could thus be attributed more confidently to changes in the representations of the plural affixes rather than to changes in the representations of the stems (which are shared with singulars). For plurals, participants' naming accuracy also improved across training blocks. Participants showed sensitivity to the phonological consistency of the plurals. Phonologically diverse plurals were named more accurately than phonologically ambiguous plurals (regular and irregular) in line with the role of high type frequency in facilitating the learning of inflectional affixes. Phonologically consistent plurals were named more accurately than phonologically ambiguous regular plurals. This is likely attributable to the fact that phonologically consistent plurals were a small group of highly consistent mappings supported by high token frequency. Phonologically ambiguous regulars, by contrast, were a small group of ambiguous mappings trained with low token frequency.

These patterns of results were further borne out in the picture naming tasks presented in the testing blocks administered at the end of each training session. Performance on the 2AFC recognition memory tasks presented during these testing blocks was at ceiling indicating that participants had mastered the meaning of the

novel words by the end of each training session. Importantly, across picture naming tasks and 2AFC recognition memory tasks, participants performed similarly for both sets of plurals as evidenced by a lack of significant main effects or interactions involving affix training day across all analyses. This suggests that training two sets of plural affixes with different sets of stems was a successful strategy to avoid interference that could have arisen from having onset overlapping competitors when training the second set of plural affixes (e.g. *gleetaff*, *gleetopp*). It also suggests that learning the first set of plural affixes did not facilitate the learning of the second set of plural affixes. Thus, generalisations requiring participants to use plural affixes trained immediately before the generalisation plural elicitation task (Day 9 affixes) or trained before a period of overnight consolidation (Day 8 affixes) could be compared for two sets of plural affixes that were learned equally well. Differences in generalisation behaviour could not be attributed to differences in the initial learning of the plural affixes.

In the generalisation plural elicitation task, participants were presented with novel singulars in both their feminine (e.g. *woathi*_[fem:sing]) and masculine (e.g. *woathu*_[masc:sing]) forms and asked to supply what they considered to be appropriate plural forms. Thus, the task required participants to use Day 8 and Day 9 affixes as required by the gender of the novel character. For novel phonologically ambiguous singulars (e.g. *varbi*, *varbu*) containing a phonological cue that had been associated with both high type frequency regular affixes (e.g. *harbaff*, *jarbopp*) and high token frequency irregular affixes (e.g. *tarbimm*, *blarboot*) during training, participants produced more irregular generalisations (ambiguous) (e.g. *varbimm*) than regular generalisations (e.g. *varbaff*) for Day 8 affixes. For Day 9 affixes, participants produced more regular generalisations (e.g. *varbopp*) than irregular generalisations (ambiguous) (e.g. *varboot*). The increase in irregular generalisations (ambiguous) following overnight consolidation replicates the result that Mirković and Gaskell (in prep.) obtained using a between-subject design. The results found here further suggest that the increase in irregular generalisations (ambiguous) is coupled with a decrease in regular generalisations. For novel phonologically diverse singulars, participants produced fewer irregular generalisations (consistent) for Day 8 affixes (e.g. *woathesh*) compared to Day 9 affixes (e.g. *woathull*). For novel phonologically consistent

singulars, participants produced fewer regular generalisations for Day 8 affixes (e.g. *tiskaff*) compared to Day 9 affixes (e.g. *tiskopp*). These latter two effects could reflect the strengthening of the regular and irregular mappings for Day 8 affixes following overnight consolidation which manifests in a reduction of inappropriate generalisations perhaps reflecting more ‘optimal’ processing (Stickgold & Walker, 2013).

As discussed in Chapter 2 and Chapter 3, memory consolidation has been implicated in the strengthening of memories for individual items, as evidenced by improvements on declarative memory tasks following overnight sleep, (e.g. Plihal & Born, 1997) and the extraction of similarities across related memories, as evidenced by improvements in generalisation performance (e.g. Ellenbogen et al., 2007; Lewis & Durrant, 2011). Stickgold and Walker (2013) argue that in cases where neither strengthening individual memories nor extracting similarities across several memories is clearly preferable as a ‘*path to integration*’ into long-term memory, memory consolidation may be particularly important in selecting one path. This account would be consistent with the shift from regular generalisations to irregular generalisations (ambiguous) observed for phonologically ambiguous novel singulars seen for Day 8 compared to Day 9 affixes. In such cases, an increase in irregular generalisations (ambiguous) could be governed by salience tags attached to memories for high-token frequency irregulars during encoding. A growing body of evidence suggests that memory consolidation is ecologically guided by the properties of the items to be remembered. For example, emotionality (Wagner, Gais, & Born, 2001; Hu, Stylos-Allan, & Walker, 2006; Atienza & Cantero, 2008; Nishida, Pearsall, Buckner, & Walker, 2009), reward motivation (Fischer & Born, 2009; Wilhelm et al., 2011; van Dongen, Thielen, Takashima, Barth, & Fernández, 2012), and explicit instructions to remember items (Fischer, Drosopoulos, Tsen, & Born, 2006; Saletin, Goldstein, & Walker, 2011), have all been shown to lead to preferential memory consolidation. It is conceivable that high token frequency irregulars were ‘*tagged*’ as being of particular importance or relevance to future behaviour and thus preferentially consolidated (but see Drosopoulos, Schulze, Fischer, & Born, 2007). As suggested by Stickgold and Walker (2013), this may be particularly important for phonologically ambiguous plurals where two types of mappings are competing for consolidation.

It should be noted that affix training day had an effect on overall performance on the generalisation plural elicitation task. Indeed, significantly more incorrect responses were produced for Day 8 compared to Day 9 affixes. A more detailed analysis of the responses produced in error showed that participants produced more gender errors and illegal affixes for Day 8 affixes compared to Day 9 affixes. This could be due to forgetting of Day 8 affixes or interference from Day 9 affixes. The present data do not allow distinguishing between these possibilities (Mensink & Raaijmakers, 1988; Anderson, 2003). However, the fact that for phonologically ambiguous novel singulars the two-way interaction between regular generalisations and irregular generalisations (ambiguous) remained after equalising the number of correct trials across the two training days suggests that the effect is not confounded by forgetting of Day 8 affixes or interference from Day 9 affixes. Furthermore, Mirković and Gaskell (in prep.) report similar results in a between-subject design in which interference between newly-learned and consolidated affixes can be ruled out. It should also be noted that the design employed in the current experiment and the subsequent experiments reported in this thesis is not intended to directly test whether overnight consolidation effects are driven by sleep or the passage of time since learning. As discussed at the outset of this chapter, however, previous research in which participants remained awake or slept for equivalent periods of time suggests a role for sleep in the learning of arbitrary mappings (Mirković & Gaskell, 2016). It thus seems likely that the increase in irregular generalisations (ambiguous) and decrease in regular generalisations for Day 8 affixes is related to overnight, sleep-related consolidation of the newly-learned affixes. It will be for further research to assess the role of sleep more directly. It is also worth mentioning that the circadian time of testing was not controlled in this experiment or the subsequent experiments reported in this thesis. Young adults have been shown to be at the peak of their cognitive abilities in the afternoon or evening (Hasher, Goldstein, & May, 2005). The possibility that time of testing may have affected the results thus cannot be ruled out. However, it is likely that all circadian times, encompassing morning and afternoon, were represented in the group of participants tested in this and subsequent experiments minimising the effects of circadian factors.

In sum, in line with the results obtained by Mirković and Gaskell (in prep.) and the prediction outlined at the outset of this chapter, when generalising to novel phonologically ambiguous singulars, participants produced more irregular generalisations (high token) compared to regular generalisations (high type) for novel plural affixes having undergone overnight consolidation compared to those trained immediately prior to generalisation. This lends support to the hypothesis that overnight consolidation may provide greater benefit in ambiguous cases where no single path to integration is clearly preferable. The results obtained here suggest that high token frequency mappings may be preferentially consolidated in such cases. A characterisation of the neural representations and mechanisms underpinning the newly-learned plural affixes is necessary to flesh out these proposals. This is the focus of the fMRI study presented in Chapter 8. The next two chapters introduce the methods, namely Representational Similarity Analysis (RSA) (Chapter 6) and task-based functional connectivity (Chapter 7), which were used to analyse the fMRI data.

Chapter 6: The representation of morphology

6.1. Decoding mental representations

Information coming from the environment or generated by inner thought is represented in patterns of neural activity. One of the key challenges for cognitive neuroscience lies in decoding these mental representations from their associated patterns of neural activity. The recent introduction of powerful pattern-classification algorithms to the analysis of fMRI data, grouped under the umbrella term ‘*multivariate pattern analysis*’ or ‘*MVPA*’, offers useful tools to decode mental representations (Norman, Polyn, Detre, & Haxby, 2006; Haynes & Rees, 2006; Kriegeskorte, Mur, Ruff, et al., 2008; Kriegeskorte, Mur, & Bandettini, 2008; Pereira, Mitchell, & Botvinick, 2009; Naselaris, Kay, Nishimoto, & Gallant, 2011; Tong & Pratte, 2012; Kriegeskorte & Kievit, 2013; Haxby, Connolly, & Guntupalli, 2014).

Most past and current fMRI studies have employed conventional univariate analyses to identify how experimental tasks or stimuli affect the overall activity of individual voxels or the mean engagement of voxels across the whole brain or an ROI (Friston, Jezzard, & Turner, 1994; Poldrack, 2011). That is, activity from many thousands of individual voxels is measured and analysed separately to identify those voxels that show a statistically significant activation in response to the experimental tasks or stimuli. Thus, the goal of this type of analysis is very much to relate brain topography to function. Brain regions that are activated by specific experimental tasks or stimuli are inferred to be functionally relevant to the perceptual or mental states that are being targeted. Conventional univariate analyses have several limitations. For example, to increase sensitivity, a smoothing kernel is typically applied to the data to spatially average across voxels. While this reduces noise in the data, it also reduces signal. Voxels with weaker activations (i.e. non-significant) may still be carrying information about the presence or absence of a particular experimental task or stimuli. Spatial smoothing also smears out fine-grained patterns of activation that may distinguish between experimental tasks or stimuli.

The MVPA approach also seeks to increase sensitivity in detecting differences between experimental tasks or stimuli but looks at the contribution of *multiple* rather

than *individual* voxels. Spatial smoothing is not typically applied to the data, which addresses some of the limitations relative to noise reduction strategies in univariate analyses. Pattern-classification algorithms are used to extract the signal that is present in the pattern of neural activity across multiple voxels, even if voxels are not individually significantly responsive to experimental tasks or stimuli (Norman et al., 2006). The core concept of high-dimensional representational vector space underlies all MVPA methods. Patterns of neural activity across voxels are converted to vectors in neural representational space. Each dimension of the neural representational space corresponds to a local feature in the distributed pattern of neural activity. For example, in fMRI, local features are usually voxels. If fMRI responses are measured for 100 voxels, then the response vectors are analysed in a 100-dimensional neural representational space (Haxby et al., 2014). Experimental tasks or stimuli also need to be converted to vectors in high-dimensional representational space. For example, spoken syllables can be converted to vectors of phonological and acoustic features in stimuli representational space. Different algorithms can then be applied to relate vectors in neural representational space to vectors in stimuli representational space.

A seminal study by Haxby et al. (2001) illustrates how the MVPA approach to fMRI analysis can be used to decode mental representations of visual stimuli. Patterns of neural activity were measured whilst participants viewed pictures of faces, cats, and other categories of objects. The pattern of neural activity elicited by each category in the ventral temporal (VT) cortex was examined to determine whether it could be distinguished from the pattern of neural activity elicited by all other categories. To this end, the data from each participant was split into two sets (i.e. odd and even scanning runs). By correlating patterns of activity elicited in the odd runs with those elicited in the even runs, Haxby et al. (2001) demonstrated that each category was associated with a distinct pattern of neural activity. For example, the pattern of neural activity elicited by viewing houses was more similar across odd and even runs than it was to the patterns of neural activity elicited by viewing other categories of objects. Importantly, the distinction between stimulus categories was found even in VT regions that responded maximally to only one category in a univariate analysis. This latter result highlights the increased sensitivity of MVPA methods over conventional univariate analyses. Several studies have since demonstrated that categories of visual

objects can be distinguished using MVPA methods (Spiridon & Kanwisher, 2002; Tsao, Freiwald, Knutsen, Mandeville, & Tootell, 2003; Carlson, Schrater, & He, 2003; Cox & Savoy, 2003; Hanson, Matsuka, & Haxby, 2004; O'Toole et al., 2005; Diana, Yonelinas, & Ranganath, 2007).

Subsequent work extended these findings to show that distinct patterns of neural activity can distinguish low-level visual features such as edge orientation (Haynes & Rees, 2005; Kamitani & Tong, 2005) and motion direction (Kamitani & Tong, 2006), the categorical content of sounds (Formisano, De Martino, Bonte, & Goebel, 2008; Staeren, Renvall, De Martino, Goebel, & Formisano, 2009), whether participants are looking at pictures or sentences, reading ambiguous or non-ambiguous sentences, and the semantic category of the words they are viewing (Mitchell, Hutchinson, Niculescu, Pereira, & Wang, 2004). MVPA approaches have also been used to decode more abstract mental states that are not directly related to the features of the experimental stimuli including lying about the identify of a playing card (Davatzikos et al., 2005), which of two tasks participants are covertly intending to perform (Haynes et al., 2007; Soon, Brass, Heinze, & Haynes, 2008), the contents of visual working memory (S. A. Harrison & Tong, 2009), and which categories of stimuli participants are thinking about during memory retrieval (Polyn, Natu, Cohen, & Norman, 2005).

6.2. Representational Similarity Analysis

RSA is a more recently introduced type of MVPA, which applies the notion of representational geometry to the analysis of neural activity patterns (Kriegeskorte et al., 2008; Nili et al., 2014). In order to characterise the geometry of a mental representation, RSA examines the dissimilarity between the patterns of neural activity elicited by different experimental stimuli or conditions. The dissimilarity between two patterns of neural activity represents the distance between their associated vectors in neural representational space. The dissimilarity between each pair of experimental stimuli or conditions is computed and yields a representational dissimilarity matrix (RDM). In its simplest form, a RDM is a square symmetric matrix with experimental stimuli or conditions indexed horizontally and vertically. The diagonal represents comparisons between pairs of identical stimuli or conditions and thus equals 0. All

other values represent the dissimilarity (distance) between the patterns of neural activity elicited by pairs of experimental stimuli or conditions. Different measures can be used to compute the dissimilarity (distance) between two vectors including correlation ($1 - \text{Pearson correlation}$, computed across voxels of two vectors), the Euclidean distance (the square root of the sum of squared differences between two vectors), and the Mahalanobis distance (the Euclidean distance measured after linearly recoding the space to whiten noise) (Kriegeskorte et al., 2008; Kriegeskorte & Kievit, 2013; Nili et al., 2014; Diedrichsen, Provost, & Zareamoghaddam, 2016). To draw inferences about the processing and representational properties of these neural RDMs, they can be compared against theoretical models, expressed as model RDMs (Figure 17).

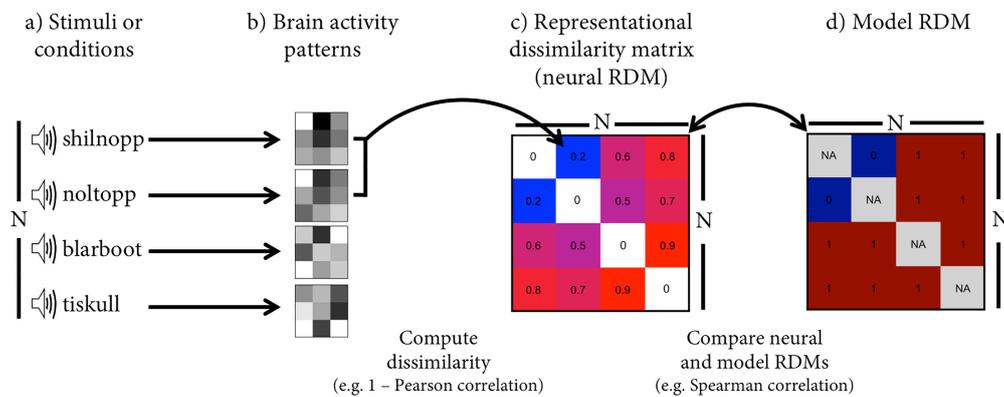


Figure 17. Representational Similarity Analysis (RSA)

Schematic of RSA steps for a simple design with four auditory stimuli. a) experimental stimuli (or, more generally, the experimental conditions to which the stimuli belong). b) the multivariate pattern of brain activity (here depicted over 9 voxels) elicited by each of the stimuli (white to black colour scale for small-to-large activity). c) the neural representational dissimilarity matrix (RDM) contains the dissimilarity values for all pairwise stimuli comparisons (blue-to-red colour scale for small-to-large dissimilarities). Here the RDM is symmetric about a diagonal of zeros (white elements). The first two stimuli 'shilnopp' and 'noltopp' (curved black arrow) have similar multivariate patterns of brain activity (i.e. low dissimilarity value) whilst all other stimuli have dissimilar multivariate patterns of brain activity (i.e. high dissimilarity values). d) The neural RDM is correlated with a model RDM (curved double arrow) to assess the extent to which the model RDM (blue elements = similar; red elements = dissimilar) reflects the multivariate patterns of brain activity elicited by the stimuli. Here, the model RDM is a good fit for the neural RDM. Several model RDMs can be tested against the same neural RDM. Model RDMs can be derived from theory, computational models, behaviour, other brain regions, other participants or even other species (Kriegeskorte et al., 2008a, 2008b).

The introduction of similarity analysis to fMRI data dates back to a study by Edelman, Grill-Spector, Kushnir, and Malach (1998) investigating object category representation in the visual cortex. The authors presented participants with pictures of object from five categories and converted patterns of neural activity for each category to multi-dimensional vectors in neural representational space. To visualise the representational geometry of these vectors, Edelman et al. (1998) used

multidimensional scaling¹⁰ (MDS) (Torgerson, 1958; Kruskal & Wish, 1978; Shepard, 1980; Borg & Groenen, 2005) to embed them in two-dimensional space. The configuration obtained from the MDS distinguished some clustering among object categories. Subsequent work applied explicit similarity analyses to patterns of neural activity in human fMRI and monkey electrophysiology. Reanalyses of the first MVP classification by Haxby et al. (2001) described in the previous section found similarity structure in the representation of visual objects based on intermediate-layer weights in a neural network (Hanson et al., 2004) and misclassifications (O'Toole et al. 2005, 2007). In both cases, greater dissimilarity between animate and inanimate objects was found. The similarity structure of visual object categories was corroborated and further characterised by Kiani and colleagues (2007) in the monkey inferior temporal (IT) cortex. Single-cell recordings in monkey IT in response to different categories of visual objects were obtained. The authors calculated correlations among response vectors as measures of the similarity of the population response vector. The results were consistent with those of Edelman et al. (1998) and showed greater dissimilarity between animate and inanimate stimuli. Using a subset of the stimuli from Kiani et al. (2007) in an fMRI study in humans, Kriegeskorte et al. (2008a, 2008b) reported a similarity structure of object representations in human VT cortex that was consistent to that of object representations in monkey IT cortex and formally introduced RSA as a data-analytical framework.

The distinction between first-order and second-order isomorphism (Shepard & Chipman, 1970) lies at the heart of RSA and underpins many of its advantages over more traditional MVPA methods. A *first-order isomorphism* refers to the relationship between an individual object and its corresponding internal representation. For example, when participants view images in an fMRI study, the patterns of neural activity that are elicited by the images in the visual cortex bear a topological resemblance to the images themselves. That is to say, a first-order isomorphism requires that a neural representation (e.g. multivariate BOLD response pattern) be

¹⁰ MDS refers to a set of methods that allow visualising the level of (dis)similarity between data points by embedding them in a n-dimensional space such that the distances between the points in the n-dimensional space match, as well as possible, the original (dis)similarities among the data points (Borg & Groenen, 2005).

related to the structural properties of the objects represented (e.g. line orientation). By exploring the relationships between experimental stimuli or conditions and patterns of neural activity, univariate analyses and early MVP classifications focused on first-order isomorphisms. A *second-order isomorphism* does away with the requirement of structural similarity between internal representations of objects and objects themselves. Instead, second-order isomorphism posits a functional equivalence between the relationships among objects and the corresponding relationships among their representations. RSA examines these second-order isomorphisms by comparing neural activity patterns across every pair of experimental stimuli or conditions in RDMs.

By converting patterns of neural activity to a set of relative distances among vectors (i.e. second-order isomorphism) in representational space rather than vectors in a feature-based representational space (i.e. first-order isomorphism), RSA allows to examine the geometry of the representational space in a way that is independent from the feature coordinate axes (Haxby et al., 2014). Exploiting second-order isomorphism in this way makes RSA a flexible method, which allows quantitatively relating representations in different brain regions in the same individual, corresponding brain regions in different individuals, corresponding brain regions in different species (e.g. humans and monkeys), a brain region to a computational model representation, a brain region to behavioural data, and representations derived from different measurement modalities (e.g. single-cell recordings and fMRI) (Kriegeskorte et al., 2008a, 2008b) by computing correlations between RDMs.

6.2. RSA and morphology

While RSA has been first and mostly used to characterise the similarity structure of visual representations, the method has been used in other fields including auditory perception (Giordano, McAdams, Zatorre, Kriegeskorte, & Belin, 2013; Evans & Davis, 2015; Blank & Davis, 2016), memory (Schurger, Pereira, Treisman, & Cohen, 2010; Ritchey, Wing, Labar, & Cabeza, 2012; Staresina, Henson, Kriegeskorte, & Alink, 2012; LaRocque et al., 2013; Liu et al., 2016), action and motor control (Wiestler, Mcgonigle, & Diedrichsen, 2011; Diedrichsen, Wiestler, & Krakauer, 2013) as well as language (Bozic, Tyler, Su, Wingfield, & Marslen-Wilson, 2013; Carota,

Bozic, & Marslen-Wilson, 2013; Devereux, Clarke, Marouchos, & Tyler, 2013; Tyler, Cheung, Devereux, & Clarke, 2013; Bozic, Fonteneau, Su, & Marslen-Wilson, 2015; Klimovich-Gray, Bozic, & Marslen-Wilson, 2017).

In the domain of morphology, a series of recent fMRI studies have used RSA to further examine the morpho-phonological decompositional model introduced in Chapter 1 (Bozic et al., 2013, 2015; Carota et al., 2016; Klimovich-Smith et al., 2017). For example, Bozic et al. (2015) used RSA to compare the neural representations supporting the processing of two types of grammatically complex sequences, namely inflected words (e.g. *sings*) and simple phrases (e.g. *I sing*). Model RDMs coding for the representational similarity among inflected words or simple phrases were compared to the neural RDM of all pairwise comparisons across complex words and simple phrases. The model RDM for inflected words provided a significant fit to the neural RDM in the left IFG and posterior superior temporal gyrus (STG), regions that have previously been associated with the processing of inflectionally complex words (Tyler, Stamatakis, et al., 2005). The model RDM for simple phrases generated a significant fit to the neural RDM in bilateral posterior and anterior MTG, consistent with previous report of these regions being involved in the processing of simple canonical sequences (Tyler et al., 2010). Bozic et al. (2015) suggest that the engagement of the left IFG in the representation of inflected words reflects its role in parsing phonologically separable affixes. Simple phrases, by contrast, do not require this type of morpho-phonological decomposition and accordingly do not engage left frontal regions. While these results provide further evidence for the morpho-phonological decompositional model of morphology, they also provide evidence for RSA's sensitivity to detect differences in the representations underlying the processing of morphologically complex words. Studies so far have applied RSA to characterise the representations of morphologically complex words in fluent speakers. However, no work has used RSA to characterise the neural representations of newly-learned morphologically complex words, or how these may change with consolidation.

6.4. RSA and consolidation

Recall from Chapter 2 that the CLS model posits that the main goal of the neocortical system is to develop overlapping distributed representations that capture the

similarity structure across sets of items. As mentioned previously, RSA of the brain activity patterns elicited during passive viewing of visual objects has provided evidence for this similarity-based coding in both primate IT cortex and human VT cortex (e.g. Kiani et al., 2007; Kriegeskorte et al., 2008a, 2008b). As discussed in Chapter 3, this similarity structure may be particularly important in allowing the development of generalisable knowledge. The developmental course of such similarity structure remains largely unexplored but it likely depends, at least in part, on overnight sleep (Lewis & Durrant, 2011).

To date, only one study has applied RSA to probe consolidation-related changes in the neural representations of experimental stimuli. Liu et al. (2016) examined the neural representations underpinning the suppression of aversive memories before and after overnight memory consolidation. Participants were trained to memorise unfamiliar face-aversive picture associations. In a design similar to the one used in Experiment 1 reported in Chapter 5, half of the face-aversive picture associations were trained 24hr and the other half 30 min before an fMRI scanning session. During scanning, participants were presented with faces trained either 24hr or 30min previously and cued to either recall the previously learned associated aversive picture or to suppress it. Liu et al. (2016) found that aversive memories become more resistant to suppression after overnight consolidation as evidenced by less-suppression induced forgetting and enduring skin conductance levels for face-aversive picture associations trained 24hr before compared to 30 min before scanning. This was concurrent with an increase in representational similarity in bilateral hippocampi for face-aversive picture associations trained 24hr before scanning and having potentially undergone consolidation.

These results may seem to be at odds with the predictions of the CLS model in that similarity-based coding would be expected in neocortical regions but not in the hippocampus, which encodes pattern-separated memories. Indeed, LaRocque et al., (2013) found representational similarity within categories of visual objects in human VT cortex but not in any hippocampal subregion. Instead, the authors found similarity-based coding in the perirhinal and parahippocampal cortex which both project to the entorhinal cortex. As mentioned in Chapter 3, the CLS model considers

the entorhinal cortex as being an intermediate between the hippocampal and neocortical systems (Kumaran et al., 2016). However, as mentioned in Chapter 2, the hippocampus also performs pattern completion, which has been argued to lead to integration in some circumstances (McClelland et al., 1995; O'Reilly & Rudy, 2001; Guzowski, Knierim, & Moser, 2004; Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004). Using RSA, Schlichting et al., (2015) showed that the anterior hippocampus maintains distinct memories for newly-learned object pairs as indexed by low representational similarity. By contrast, the posterior hippocampus was found to integrate across memories for newly-learned object pairs as indexed by high representational similarity. Liu et al.'s (2016) hippocampal results have a peak in posterior hippocampi consistent with the integration results of Schlichting et al. (2015). They further highlight a role for overnight-consolidation in supporting the development of similarity structure. Thus, the role of the hippocampus in integrating over memories, or in generalisation more broadly as discussed in Chapter 3, remains unclear. Liu et al. (2016) nonetheless present results suggesting that RSA is sensitive to overnight-changes in the neural representations of new memories.

6.5. Summary

Before the emergence of MVPA methods, conventional fMRI analyses were mostly focused on describing which experimental task or stimuli activates a brain region globally. The introduction of MVPA methods shifted the focus away from describing *where* the information was represented in the brain to asking *what* type of information is represented, in terms of perceptual or mental states associated with distinct patterns of activity, and how this information is encoded and organised. RSA is a flexible MVPA method that allows characterising the similarity structure among sets of items. As such, it is ideally suited to test predictions derived from the CLS model regarding the development of similarity structure in the neocortical system. RSA has been successfully applied to the study of morphological processing but no work so far has looked at morphological learning. One study so far has provided evidence for consolidation-related changes in the neural representation of new emotional memories but no work has considered memories for new morphological knowledge.

Chapter 7: The connectivity of morphology

7.1. Functional connectivity

As discussed at the outset of Chapter 6, conventional univariate fMRI analyses have mostly focused on characterising how functional specialisation is anatomically segregated to specific brain regions. RSA allows characterising what type of information is represented in specific brain regions in terms of multivariate patterns of brain activity associated with stimuli or experimental conditions. However, both conventional univariate fMRI analyses and RSA disregard how several brain regions may work together to process information (Friston, 2011). One way to characterise how brain regions may work together, or be functionally *integrated*, is in terms of functional connectivity. Functional connectivity refers to the correlations between the activity fluctuations of spatially remote brain regions (Biswal, Van Kylen, & Hyde, 1997; Fox & Raichle, 2007; Guye, Bartolomei, & Ranjeva, 2008). Resting-state functional connectivity measures spontaneous brain activity fluctuations at rest, in the absence of any task (Fox & Raichle, 2007; Smith et al., 2013). Task-dependent functional connectivity, by contrast, focuses on statistical dependencies between the activity fluctuations of different brain regions whilst participants perform particular tasks or in response to different experimental stimuli or conditions (Friston et al., 1997). The next section is concerned with psychophysiological interactions (PPI) analysis, a particular type of task-dependent functional connectivity analysis (Friston et al., 1997; Friston, 2011; O'Reilly, Woolrich, Behrens, Smith, & Johansen-Berg, 2012).

7.2. Psychophysiological interaction analysis

The goal of PPI analysis is to identify a set of voxels in the brain whose activity fluctuations are correlated with the activity fluctuations of a seed ROI in a given context, such as during a particular behavioural task or experimental condition. That is, PPI analysis aims to identify brain regions whose activity depends on an interaction between psychological factors (i.e. behavioural task or experimental conditions) and physiological factors (i.e. the activity fluctuations of a seed ROI). A PPI effect is reflected by a task-dependent (or experimental condition-dependent) increase in the

functional coupling between brain regions (Friston et al., 1997; O'Reilly et al., 2012). It should be noted that PPI does not allow causal inference. In other words, a PPI effect does not imply that the seed region is the driver rather than driven brain region. Incidentally, PPI does not allow to address whether the functional coupling between a seed region and another brain region is direct or mediated by other brain regions. Nonetheless, PPI affords the possibility to characterise how the functional interplay between brain regions is changed as a function of tasks or experimental conditions (Figure 18).

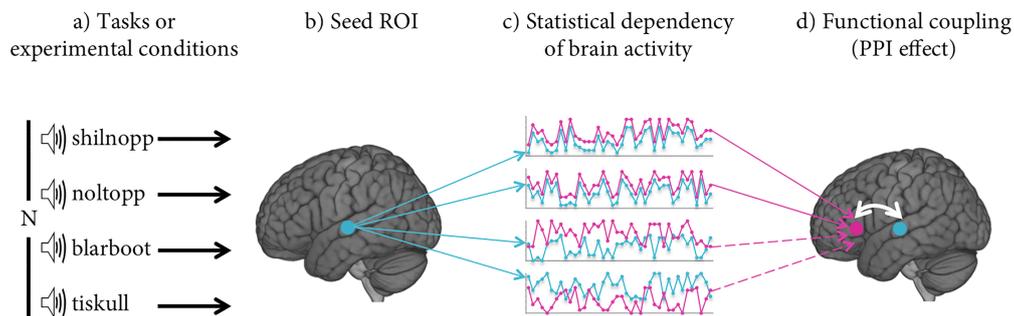


Figure 18. PPI functional connectivity analysis

Schematic of PPI functional connectivity analysis steps. a) experimental auditory stimuli belonging to four (N) experimental conditions b) a left STG seed region is selected for the PPI analysis c) brain activity fluctuations (i.e. timecourse over scanning events) for each condition are extracted from the seed region (blue arrows) c) brain activity fluctuations from the seed region (blue lines) are correlated with every other brain voxel (ROI-to-voxel PPI) or with voxel within another ROI (ROI-to-ROI PPI). Here, the activity fluctuations of the left STG seed region are correlated with the activity fluctuations of a set of voxels in the left IFG for the first two experimental conditions (full pink lines) but not for the bottom two experimental conditions (dashed pink lines) d) the functional coupling between the left STG seed region and the left IFG (white double arrow) is increased in the context of the first two conditions (here reflecting a main effect of the regularity of the plural affix 'opp' shared by 'shilnopp' and 'noltopp').

7.3. PPI and morphology

Only two studies have applied PPI analysis to the study of morphology. Stamatakis, Marslen-Wilson, Tyler, and Fletcher (2005) examined differences in the functional connectivity patterns elicited by the processing of regular and irregular English verbs. Participants performed an auditory same – different task in which they were presented with two successive spoken words and asked to judge whether the second word was the same as the first. Spoken words included regular (e.g. *played* – *play*) and irregular past tenses (e.g. *taught* – *teach*). Stamatakis et al. (2005) used the left IFG and the anterior cingulate cortex (ACC) as seeds as they were previously reported to show differential activation in response to regular and irregular verbs in a univariate

analysis reported by Tyler et al. (2005). The left IFG seed was found to be functionally connected to a cluster of voxels in the left posterior MTG for regular verbs. The ACC seed itself was not significantly connected to any voxels. However, the functional connectivity between the left IFG and ACC seeds covaried with activity in the left posterior MTG more for regular than irregular verbs. Stamatakis et al. (2005) argue that these findings are consistent with regular verbs being more strongly dependent on the fronto-temporal network than irregular verbs. Regular verbs are argued to require modulation by the left IFG morpho-phonological parsing functions. Stamatakis et al. (2005) further argue, on the basis of previous work showing strong reciprocal connections between the ACC and prefrontal areas (Vogt, Rosene, & Pandya, 1979; Goldman-Rakic, 1988; Petrides & Pandya, 1988), that the ACC may play a monitoring role on the fronto-temporal network.

In an fMRI study of Russian morphology, Kireev, Slioussar, Korotkov, Chernigovskaya, & Medvedev (2015) replicated Stamatakis et al.'s (2005) results. The authors examined differences in the functional connectivity patterns elicited during the production of regular and irregular Russian verbs. A left IFG seed region showed greater functional coupling with clusters of voxels in bilateral STG for regular verbs compared to irregular verbs. Kireev et al. (2015) argue that their findings provide some cross-linguistic validation of the results obtained by Stamatakis et al. (2005). Kireev et al. (2015) attribute their findings of bilateral temporal connectivity, compared to the left lateralised connectivity reported by Stamatakis et al. (2005), to self-monitoring required by the production task, which has previously been associated with bilateral STG activation (Indefrey, 2011). Together, these results provide further evidence for the morpho-phonological decomposition model of morphology. They also provide evidence that PPI analyses can reveal differences in the functional networks underlying the processing of regularly and irregularly inflected words. Similarly to RSA, no work has applied PPI analyses to characterise the functional networks underlying newly-learned inflected words, or how these may change with overnight consolidation.

7.4. PPI and consolidation

As described in Chapter 2, the CLS model suggests that new memories undergo notable changes in functional organisation between the hippocampal and neocortical systems (McClelland et al., 1995). However, there has been very little work looking at consolidation-related changes in functional connectivity within and between these systems. Some studies have examined the role of sleep in influencing functional connectivity patterns among putative memory brain areas across different groups of subjects and different experimental conditions but without directly addressing the role of memory consolidation (Addis, Wong, & Schacter, 2007; Gais et al., 2007; Sterpenich et al., 2009). As mentioned in Chapter 2, Paz et al. (2007) provided evidence that the mPFC can facilitate hippocampal-neocortical interactions during reward-based learning. Using single-cell recordings in cats, Paz and colleagues (2007) found that after several days of training, mPFC activity was found to enhance entorhinal to perirhinal communication leading them to suggest that the mPFC is involved in the slow iterative processing supporting the integration of hippocampal memories into neocortical networks.

In an elegant study, Takashima et al. (2009) provided the first human evidence for consolidation-related changes in functional connectivity between the hippocampal and neocortical systems. Participants were trained, on two consecutive days, to associate the pictures of unfamiliar male and female faces with one of six spatial locations. One subset of face – location associations was trained on Day 1 and another subset was trained on Day 2. After being trained on the second set of face – location associations, participants performed a cued recall task during an fMRI scanning session in which they had to retrieve the correct spatial location when cued with a trained face. Using PPI, Takashima et al. (2009) compared functional connectivity patterns for correctly recalled face – locations associations trained on Day 1 and Day 2 using a bilateral hippocampal seed region showing reduced activation in a univariate analysis for correctly retrieved spatial locations for Day 1 compared Day 2 associations. The functional connectivity between the bilateral hippocampal and early visual areas extending into the fusiform face area (FFA) and between the bilateral hippocampal and the posterior parietal cortex (PPC) decreased with overnight

consolidation. Conversely, the authors noted increased functional connectivity with overnight consolidation between the FFA and the PPC as well as between the FFA and early visual areas. Takashima et al. (2009) argue that these patterns of functional connectivity are consistent with the CLS model and provide evidence that memory retrieval networks shift from being initially dependent on hippocampal – neocortical interactions towards relying on neocortical – neocortical interactions.

Only two studies have assessed consolidation-related changes in functional connectivity in the context of language learning (Takashima, Bakker, van Hell, Janzen, & McQueen, 2014, 2016). For example, Takashima et al. (2014) examined consolidation-related changes in the functional networks involved in learning new words in their phonological form only or with their visual referent. On Day 1, participants learned novel words via a phoneme-monitoring task, in which half of the words were phonological forms only (form-only) and half were associated with a picture referent (picture-present). Results showed lexical competition effects emerging 24hr later for form-only words but not for picture-present words. Takashima et al. (2014) argued that this delay in lexicalisation reflects reduced phonological processing during learning as a result of the need to learn semantic associations (cf. Leach & Samuel, 2007; Hawkins & Rastle, 2016). Participants also performed recognition memory tests during fMRI scanning sessions on Day 1 and again on Day 2. Since the recognition memory tests required participants to match the novel spoken words with their memory representations, Takashima et al. (2014) focused on the functional connectivity arising from a seed region in bilateral STG identified from contrasting all novel words with baseline in a univariate analysis. On Day 1, there was significantly greater functional connectivity between bilateral STG and a cluster of voxels within the right STG for correctly recognised form-only compared to correctly recognised picture-present words. On Day 2, the bilateral STG was not significantly connected to any voxels. When restricting the analysis to phonological representational areas in left posterior MTG and STG, Takashima et al. (2014) found that the magnitude of the lexical competition effect for form-only words was positively correlated with the magnitude of the functional connectivity between bilateral STG and a cluster of voxels in left posterior MTG. The positive correlation between lexical competition effects and bilateral STG – left posterior MTG connectivity on Day 2 for form-only words is

argued by Takashima et al. (2014) to reflect stronger associations between newly-learned words and existing words leading to lexical competition. However, unlike the design used by Davis et al. (2009), the design used by Takashima et al. (2014) does not allow to differentiate consolidation-related changes in the functional networks underlying word learning from changes due to re-exposure to the words on Day 2.

7.5. Summary

PPI analyses allow characterising the functional networks involved in performing different tasks or experimental conditions. As such, it provides a complementary method to traditional univariate analysis and MVPA methods in that it seeks to understand how several brain regions are functionally integrated in processing information. PPI is well-suited to examine consolidation-related changes in functional organisation between the hippocampal and neocortical system that are predicted by the CLS model. PPI has been successfully applied to characterising the functional networks involved in morphological processing but there is a lack of evidence concerning the functional networks involved in the processing of newly-learned morphology. Previous studies have provided evidence for consolidation-related changes in the functional networks underlying newly-learned visuospatial associations and words but no work has considered novel morphology. The next chapter reports on Experiment 2, which applied RSA and PPI analyses to fMRI data in the context morphological learning and consolidation.

Chapter 8: Experiment 2

8.1. Introduction

The overnight changes in generalisation behaviour observed in Experiment 1 suggest overnight changes in the representations underlying knowledge of the plural affixes. Indeed, one cognitive marker of morphological knowledge is the ability to generalise to novel forms (Bybee & Slobin, 1982, Bybee & Moder, 1983, Prasada & Pinker, 1993, Tamminen et al., 2015). The results of Experiment 1 as well as those obtained by Mirković and Gaskell (in prep.) further suggest a role for overnight consolidation in modulating the influence of type and token frequency and phonological consistency on the development of this morphological knowledge. Recall that, in both experiments, participants produced more irregular generalisations (ambiguous) (e.g. *varbimm*) than regular generalisations (e.g. *varbaff*) for novel phonologically ambiguous singulars (e.g. *varbi*) after a period of overnight consolidation. As discussed in Chapter 2, these results are consistent with the suggestion that effects of overnight consolidation may be particularly apparent in cases where different types of mappings are competing for consolidation (Stickgold & Walker, 2013). In such cases, high token frequency mappings may be preferentially consolidated. To flesh out this proposal and since new morphological knowledge must be represented neurally, the main goal of Experiment 2 was to build upon the paradigm developed in Experiment 1 to characterise the neural representations underlying the newly-learned plurals.

RSA and PPI analyses described in the previous two chapters were applied to fMRI data collected in a new independent group of participants. The same experimental design as in Experiment 1 was used with some minor changes, which are described in section 8.2. Participants were trained on two sets of plural affixes, with varying phonological consistency and type and token frequency, on Day 8 and Day 9. Effects of initial learning and overnight consolidation could be assessed by comparing the neural representations of plurals trained on Day 8 to those trained on Day 9. This design adapts the method developed by Davis et al. (2009) in the context of lexical learning and consolidation. Training the same participants on two sets of plurals and testing for effects of overnight consolidation in a single fMRI session is an efficient design to address potential variability of functional imaging results arising when

scanning the same participants across different sessions (Noll et al., 1997; Rombouts et al., 1997; Mcgonigle et al., 2000; Waldvogel, van Gelderen, Immisch, Pfeiffer, & Hallett, 2000; Maitra, Roys, & Gullapalli, 2002). It also avoids multiple exposures to the same training materials such that overnight changes in neural representations can be more confidently attributed to consolidation processes rather than to additional learning taking place (cf. Takashima et al., 2014, Chapter 7).

Previous work, reviewed in Chapter 6, has shown that RSA can pick out similarity structure among inflected words in putative language representational areas in the context of natural language processing (e.g. Bozic et al., 2015). Here, the method is applied for the first time to characterise the similarity structure underpinning newly-learned inflected words trained with varying phonological consistency and type and token frequency in an artificial morphological system. Furthermore, the role of overnight consolidation in modulating the influence of these experimental factors on the development of similarity structure is considered. On the basis of the overnight changes in generalisation behaviour observed in Experiment 1 and the similar results reported by Mirković and Gaskell (in prep.), overnight changes in similarity structure were predicted to be mostly apparent for phonologically ambiguous plurals. In such cases, representational differences between phonologically ambiguous (regular) (e.g. *har**ba**ff*) and phonologically ambiguous (irregular) (e.g. *tar**bi**mm*) plurals were predicted to grow larger after a period of overnight consolidation. On the basis of increased irregular generalisations (ambiguous) observed for Day 8 compared to Day 9 affixes in Experiment 1, the similarity structure among phonologically ambiguous (irregular) plurals (e.g. *tar**bi**mm*, *clar**bi**mm*, *slar**bi**mm*) was predicted to increase for Day 8 compared to Day 9 plurals. As reviewed in Chapter 2, the CLS model predicts that high type frequency should support the development of overlapping neocortical representations (i.e. similarity structure). It also predicts that the development of such similarity structure should proceed slowly and thus benefit from overnight consolidation. Consistent with this prediction, Tamminen and colleagues (2012, 2015) have provided behavioural evidence in the context of derivational morphology that both consolidation and high type frequency support generalisation to novel forms. Hence, greater similarity structure among Day 8 compared to Day 9 high type frequency regular plurals was also expected.

As reviewed in Chapter 7, the CLS model also predicts consolidation-related changes in the functional networks supporting learning and memory. With consolidation, new memories should gradually become less reliant on the hippocampus and integrated into neocortical networks such that hippocampal-neocortical functional connectivity should become weaker and neocortical-neocortical functional connectivity should become stronger. As discussed in Chapter 7, only two studies have reported consolidation-related changes in functional connectivity in the context of lexical learning (Takashima et al., 2014, 2016). Here, changes in the functional networks supporting newly-learned inflected words are considered for the first time. In line with the CLS model, plurals trained on Day 9 were predicted to involve stronger functional connectivity between the hippocampus and putative language representational areas. By contrast, plurals trained on Day 8 were predicted to involve stronger functional connectivity between putative language representational areas, while hippocampal connectivity to these areas was predicted to be weaker. On the basis of the overnight changes in generalisation behaviour observed in Experiment 1 and reported by Mirković and Gaskell (in prep.), overnight changes in functional connectivity were predicted to be mostly apparent for phonologically ambiguous plurals. In such cases, functional connectivity differences between phonologically ambiguous (regular) (e.g. *har**ba**ff*) and phonologically ambiguous (irregular) (e.g. *tar**bi**mm*) plurals were predicted to change with overnight consolidation. An increase in irregular generalisations (ambiguous) may reflect the development of similarity structure among high token frequency phonologically ambiguous (irregular) plurals. To the extent that this similarity structure may reflect the development of a new schema (i.e. an inflectional pattern), stronger neocortical-neocortical functional connectivity between language representational areas was predicted for Day 8 compared to Day 9 phonologically ambiguous (irregular) plurals. Recall from Chapter 3 that greater neocortical-neocortical functional connectivity has been implicated in the retrieval of schema consistent information (Van Kesteren et al., 2010a, 2010b) and that sleep has been implicated in supporting schema formation (Lewis & Durrant, 2011). As such, stronger neocortical-neocortical functional connectivity between language representational areas was also predicted for phonologically consistent and high type frequency regular plurals.

8.2. Materials and methods

8.2.1. Participants

Twenty-two participants aged between 18 and 34 (mean age = 23, SD = 4, 8 males) were recruited from the MRC Cognition and Brain Sciences Unit Participant Panel and provided their informed consent to take part in the study. Participants were tested under the approval of the Cambridge Psychology Research Ethics Committee. All were native monolingual speakers of British English with little or no knowledge of a second language, no known hearing or language impairments, and no neurological or psychiatric disorders. Participants were paid to take part in the study. One participant had to be excluded from behavioural and fMRI analyses due to floor performance on the training tasks and excessive movement in the scanner. The final dataset included 21 participants.

8.2.2. Experimental stimuli

8.2.2.1. Training stimuli

The training stimuli described in Experiment 1 were used with some minor changes. Within each set of stems, 3 phonologically varied stems were replaced with stems containing a new phonological cue (e.g. *-olt*; *zolt*, *grolt*, *tolt*), which was consistently associated with a regular affix (e.g. *zoltaff*, *groltaff*, *toltaff*) analogous to phonologically consistent regular English verbs (e.g. *walked*, *talked*, *stalked*). This modification was made to embed a factorial design manipulating the phonological consistency (consistent, ambiguous) and the affix regularity (regular, irregular) of the novel plurals (Figure 19). That is, phonological consistency and affix regularity were fully crossed such that observed changes in the neural representations of the novel plurals could be attributed to these experimental factors rather than reflecting word-specific idiosyncrasies. As these changes were made to both sets of plurals, the experimental factors were further crossed with affix training day (Day 8, Day 9) such that the role of overnight consolidation in modulating their influence could be assessed. The same type and token frequency manipulations as in Mirković and

Gaskell (in prep.) and Experiment 1 were used. Twenty-four training lists¹¹ were created to counterbalance the assignment of plural affixes to regular and irregular conditions and the assignment of phonological cues to consistent and ambiguous conditions. The assignment of each set of plural affixes (*-aff*, *-imm*, *-esh*; *-opp*, *-oot*, *-ull*) to each gender was counterbalanced across participants, as was the order of gender training across Day 8 and Day 9. As in Experiment 1, the pairing of nouns with specific occupations was varied across different training lists to further control for any unintended associations between the novel nouns and pictures.

a)

Phonological consistency	Type/Token frequency	Training set 1		Training set 2	
		Singular nouns	Plural nouns	Singular nouns	Plural nouns
Diverse	High type/ Low token (regular)	gleeti	gleetaff	shilnu	shilnopp
		torthi	torthaff	plassu	plassopp
		vontti	vontaff	blornu	blornopp
		mowli	mowlaff	dowthu	dowthopp
		sleni	slenaff	rutchu	rutchopp
		pelfi	pelfaff	crephu	crephopp
Consistent	High type/ Low token (regular)	skolti	skoltaff	zoltu	zoltopp
		grolti	groltaff	noltu	noltopp
		wolti	woltaff	quoltu	quoltopp
Ambiguous	High type/ Low token (regular)	harbi	harbaff	jarbu	jarbopp
		yarbi	yarbaff	larbu	larbopp
	narbi	narbaff	glarbu	glarbopp	
	High token/ Low type (irregular)	tarbi	tarbimm	blarbu	blarboot
clarbi		clarbimm	marbu	marboot	
		slarbi	slarbimm	farbu	farboot
Consistent	High token/ Low type (irregular)	hiski	hiskesh	tisku	tiskull
		liski	liskesh	visku	viskull
		fiski	fiskesh	yisku	yiskull

b)

Type frequency	Token frequency	Affix frequency
12	6	72
3	24	72
3	24	72

c)

Figure 19. Experiment 2: training stimuli set

a) example training stimuli set. Phonologically diverse, phonologically consistent, and phonologically ambiguous novel plurals taking a regular affix (blue), phonologically ambiguous novel plurals taking an irregular affix (green), and phonologically consistent novel plurals taking an irregular affix (red). The factorial design embedded in the stimuli set is outlined in orange. b) type, token, and affix (type x token) training frequencies associated with each plural affix. c) example novel singular and plural nouns associated with familiar occupations for female and male characters.

8.2.2.2. Generalisation stimuli

The generalisation stimuli were the same as described for Experiment 1 except that they now included a further 9 stems containing the new phonological cue (e.g. *-olt*; *crolt*) for a total of 36 novel stems. Each novel stem (e.g. *woath*) was combined with both singular gender affixes (e.g. *woathi*, *woathu*) yielding a total of 72 generalisation items. Each new singular noun was paired with the picture of a single female and male

¹¹ Twenty-four participants (one per training list) were due to be scanned. Twenty-two participants were scanned before an unexpected MRI scanner quench. As the MRI scanner required a new calibration and because of time constraints, it was decided not to scan an additional 2 participants.

character depicting a familiar occupation not used in the training stimuli sets (e.g.

$woathi_{[fem:sing]} = teacher_{[fem:sing]}$, $woathu_{[masc:sing]} = teacher_{[masc:sing]}$).

8.2.3. Experimental design and procedure

The training schedule was the same as described for Experiment 1 (see Chapter 5, section 5.2.3). On Day 9, participants additionally completed a 3AFC generalisation task and an fMRI scanning session (Figure 20). The procedure for these additional tasks is described in section 8.2.3.1. All other tasks followed the same procedure as described in Experiment 1.

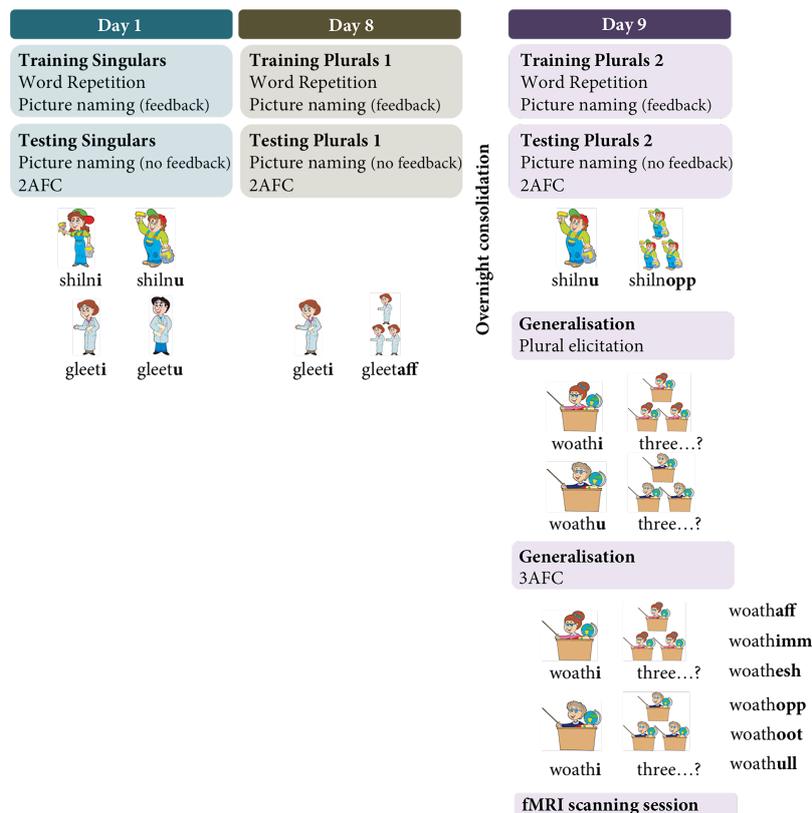


Figure 20. Experiment 2: experimental design

Training took place over 9 days. Participants were trained on the singulars on Day 1. They returned to the lab one week later to be trained on one set of plurals for one gender on Day 8. They were trained on a second set of plurals for the other gender on Day 9 following a 24hr delay containing sleep (i.e. overnight consolidation). They then performed a plural elicitation generalisation task and a 3AFC generalisation task, which assessed knowledge of the plural affixes trained immediately before or on the previous day. Participants then completed an fMRI scanning session.

8.2.3.1. Day 9

8.2.3.1.1. Generalisation - 3AFC

To address the fact that participants produced fewer correct responses for Day 8 compared to Day 9 affixes on the plural elicitation generalisation task in Experiment 1, an additional 3AFC generalisation task was administered that ensured the same number of correct responses for both sets of affixes. Each trial began with a black

fixation cross presented for 250ms in the centre of a white background followed by the auditory presentation of a novel singular. The picture of the corresponding single character with the phrase ‘one [novel word_[sing]]’ displayed underneath it were presented in the centre of the screen 300ms after the onset of the spoken word and remained on the screen for 2000ms. A blank screen was then presented for 250ms, which was followed by the presentation of the same character depicted as a triplet. The written form of one possible plural was presented for 1000ms on the left of the screen together with its spoken form. The first word disappeared from the screen and a second possible plural was presented for 1000ms in the centre of the screen together with its spoken form. The second word disappeared from the screen and a third possible plural was presented for 1000ms on the right of the screen together with its spoken form. All three words had the same correct stem but a different affix selected from the gender-appropriate set of plural affixes (e.g. *woathaff*_[fem:plur]; *woathimm*_[fem:plur]; *woathesh*_[fem:plur]; for *woathi*_[fem:sing]). There was a 250ms interval from the last plural disappearing to all three words reappearing on the screen with the phrase “which is the most appropriate?” above the three words. Participants were instructed to respond by pressing the “1”, “2”, or “3” key on the keyboard to indicate that the word on the left, centre, or right, respectively, was the most appropriate plural. There was an inter-trial interval of 250ms. Response time analyses were precluded since participants were cued to respond after a delay. Each singular was presented once and the order of presentation of the trials was randomised for each participant. The same generalisation items as in the plural elicitation task were used. Each affix appeared the same number of times in each location for plurals in each condition (i.e. phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)).

8.2.4. fMRI scanning

8.2.4.1. fMRI procedure

During the fMRI scanning session, participants were presented with three types of trials: spoken word trials, null events, and test trials (Figure 21). Each spoken word trial started with the auditory presentation of a trained word in the 1000ms silent gap between two functional scans and was followed by a single functional scanning

volume acquired in the subsequent 2000ms (i.e. TR = 3000ms). Null events followed the same timeline except that no word was presented in the silent gap. These null events provided an implicit resting baseline against which responses to the plurals could be compared. In occasional test trials, a visual display prompted participants' knowledge of the affix or whole word form for the trained word heard on the previous trial. The visual display remained on screen for 4000ms during which participants' responses were recorded. The entire test trial procedure spanned 6000ms (i.e. 2 TRs). Except for when a test trial visual display was presented, a black fixation cross presented in the centre of a grey background remained on screen for the entire scanning session. Test trials consisted of 4AFC word-to-picture matching trials. For test trials assessing affix knowledge, participants were presented with the pictures of four characters corresponding to words sharing the same stem but having different affixes (e.g. *gleeti*_[fem:sing], *gleetaff*_[fem:plur], *gleetu*_[masc:sing], *gleetopp*_[masc:plur] for a test of *gleetaff*_[fem:plur]). For test trials assessing whole word knowledge, participants were presented with the pictures of four different characters corresponding to words of the same gender and number (e.g. *gleetaff*_[fem:plur], *torthaff*_[fem:plur], *harbaff*_[fem:plur], *hiskesh*_[fem:plur] for a test of *gleetaff*_[fem:plur]). For whole word test trials, each word appeared once as a target and three times as a distractor. For both types of test trials, the four pictures were presented simultaneously along the horizontal axis in the middle of the screen against a grey background. Participants pressed a button with their right index, middle, ring, or little finger to select the picture of the character corresponding to the spoken word they heard on the previous trial. The location of the target was counterbalanced across trials. Since responses to test trials were made after the auditory presentation of the words (i.e. a spoken word trial), the fMRI analyses could be focused on trials where words were presented intact in the absence of a button press response. Response time analyses were precluded since participants were cued to respond after a delay. Participants were given the opportunity to practice the 4AFC task before going in the scanner.

Each of six scanning runs contained 216 trials. Each of the 72 trained spoken words (18 Day 8 plurals, 18 Day 9 plurals, 36 singulars (18 associated with Day 8 plurals; 18 associated with Day 9 plurals)) was presented twice in each scanning run for a total of 144 spoken word trials per run. Null events accounted for 24 trials per

run and test trials made up 48 trials per run (24 test trials spanning two TRs each). Test trials were distributed across the 6 scanning runs such that participants were tested on their affix and whole word knowledge for each trained word (i.e. two tests for each word amounting to 144 test trials for the entire scanning session). Including three initial equilibration scans, each run lasted 11 minutes. The overall scanning session lasted approximately 70 minutes including the acquisition of a structural scan. Stimulus delivery was controlled and behavioural responses were recorded with E-Prime 2.0 software (Psychology Software Tools, Inc.).

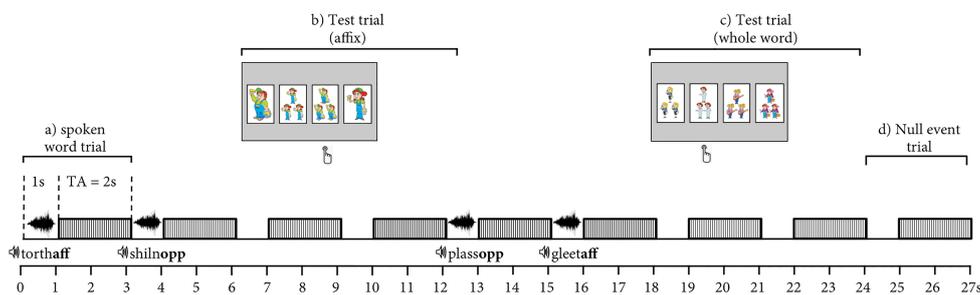


Figure 21. Experiment 2: fMRI scanning procedure

Timeline of the fast sparse fMRI procedure illustrating the rapid alternation of spoken word presentation and single scan volumes and the different types of trials: a) spoken word trial, b) test trial assessing affix knowledge, c) test trial assessing whole word knowledge, d) null event trial providing an implicit baseline.

8.2.4.2. Scanning parameters

8.2.4.2.1. Structural scanning

MRI data were acquired on a 3-Tesla Siemens Tim Trio scanner using a 32-channel head coil. A high-resolution T1-weighted structural scan was acquired for each participant using a three-dimensional MPRAGE sequence (TR = 2250ms, TE = 2.99ms, flip angle = 9°, field of view = 256 x 256 x 192mm, matrix size = 256 x 256 x 192mm, spatial resolution = 1 x 1 x 1 mm).

8.2.4.2.2. Functional imaging

The fMRI session was split into 6 runs of approximately 11 min each. For each participant and each scanning run, 220 echo planar imaging (EPI) volumes comprising 32 slices with a thickness of 3mm were acquired using a continuous, descending acquisition sequence (TR = 3000ms, TA = 2000ms, TE = 30ms, flip angle = 78°, matrix size = 64 x 64, field of view = 192mm x 192mm, in-plane resolution = 3 x

3mm, interslice gap = 25%). The field of view allowed covering most of the brain except for the most superior portion of the parietal lobe. To avoid interfering effects of scanner noise when presenting the spoken words, sparse acquisition was used such that a single volume (TA = 2000ms) was acquired following a silent period (1000ms) during which a single spoken word was presented. Fieldmaps (~5min) were also acquired for use in preprocessing and normalisation. Visual stimuli for the test trials were projected on a screen at the head-end of the scanner table and reflected onto a mirror attached to the head coil above the participants' eyes. Auditory stimuli were presented over Sensimetrics insert headphones (<http://www.sens.com>) and manual responses for test trials were recorded with a MR-compatible button-box.

8.2.5. Data analysis

8.2.5.1. Behavioural analysis

The training, testing, and generalisation tasks were analysed following the same method described for corresponding tasks in Experiment 1 (see Chapter 5, section 5.2.4). For the additional 3AFC generalisation task, data was analysed using loglinear regression with a Poisson distribution applied to the frequency counts of each response type. Accuracy on the 4AFC test trials presented in the scanner was analysed using logistic regression with a binomial distribution.

8.2.5.2. Univariate fMRI analysis

Data were analysed using SPM12 (<http://www.fil.ion.ucl.ac.uk/spm>) in conjunction with automatic analysis (aa) pipelines (version 4.2) (Cusack et al., 2015). The first three volumes of each scanning run were removed to allow for T1 equilibrium effects. Scans were realigned to the first functional image. Functional images were corrected for motion by realigning them to the first functional image. Fieldmaps were used to correct for geometric distortions to the functional images resulting from inhomogeneities in the magnetic field (Cusack, Brett, & Osswald, 2003). Structural scans were coregistered to the mean of the realigned, undistorted functional image. Coregistered structural images were normalised to the 152-subject T1 MNI stereotaxic template brain. The resulting segmentation parameters were used to normalise the functional images, which were resampled to 2mm isotropic voxels. To compensate for residual variability after normalisation and to facilitate inter-subject comparisons,

realigned, normalised functional images were smoothed with a Gaussian kernel of 8mm with full width half maximum.

Data were analysed using a general linear model (GLM) with a 128s high-pass filter and AR1 correction for auto-correlation. The onsets of 22 event types were included in the GLM, each convolved with the canonical SPM haemodynamic response, in addition to 6 movement parameters estimated at realignment included as regressors of no interest. The null events provided an implicit resting baseline. Eight conditions came from specifying the onset of plurals trained on successive days (Day 8, Day 9), crossed with phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular) (i.e. the factorial design embedded in the training stimuli, see section 8.2.2.1) (8 event types). The onset of phonologically diverse plurals trained on Day 8 and Day 9 were also specified (2 event types). The onset of the singulars associated with the plurals just described was also specified (10 event types). The onset of test trials assessing affix knowledge and whole word knowledge were also specified (2 event types). Following parameter estimation of the first-level model and to focus on addressing the research questions and predictions outlined at the outset of the chapter, a second-level analysis consisted of a repeated-measures ANOVA with affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular) conducted on the plurals to assess the main effects and interactions of these factors on blood-oxygen-level dependent (BOLD) activity. All statistical parametric maps and statistics reported in tables were thresholded at $p < .001$ uncorrected at the voxel level and $p < .05$ family-wise-error (FWE) corrected at the cluster level.

In addition to a whole-brain analysis, an ROI analysis focused on the effects of affix training day, phonological consistency, and affix regularity in the left STG, the left HC, and the left IFG. As discussed in Chapter 2, consolidation-related changes in neural activity have been reported in the left STG in the context of lexical (Davis et al., 2009, Takashima et al., 2014) and morphological learning (Leminen et al., 2016). Left superior temporal regions have also been associated with the processing of inflectionally complex words (Tyler et al., 2005; Bozic et al., 2015) and shown to produce greater activity in response to pseudowords compared to real words

(Newman & Twieg, 2001; Kotz, Cappa, von Cramon, & Friederici, 2002). Current neuroanatomical accounts of auditory language processing suggest distinct dorsal and ventral processing streams radiating out of primary auditory cortex (Hickok & Poeppel, 2004, 2007; Davis & Johnsrude, 2007). There is also evidence for differential involvement of the anterior and posterior superior temporal gyrus in grammatical processing (Friederici, 2002, 2011; Humphries, Binder, Medler, & Liebenthal, 2006). In line with these proposals, and for greater anatomical precision, anterior (aSTG) and posterior (pSTG) STG ROIs were defined. The left aSTG and left pSTG ROIs were defined based on two clusters identified in a meta-analysis contrasting neural responses to pseudowords and real words (i.e. pseudowords > real words) (Davis & Gaskell, 2009): an aSTG ROI with a centre of mass of $x = -56$, $y = -6$, $z = -1$ (128 voxels) and a pSTG ROI with a centre of mass of $x = -61$, $y = -28$, $z = 10$ (510 voxels).

Based on the CLS model and previous work implicating the left hippocampus (HC) in word learning and consolidation (Breitenstein et al., 2005, Davis et al., 2009), a left hippocampal ROI was defined anatomically using the Automatic Anatomical Labelling (AAL) atlas (Tzourio-Mazoyer et al., 2002). As discussed in Chapter 6, the anterior HC (aHC) and posterior HC (pHC) have been implicated in separating and integrating novel memories, respectively (Schlichting et al., 2015). The aHC and pHC have also been implicated in encoding and retrieving new information, respectively (Lepage, Habib, & Tulving, 1998; Schacter & Wagner, 1999; Spaniol et al., 2009). In keeping with the suggestion of a functional dichotomy along the anterior-posterior axis of the hippocampus, the left HC ROI was segmented into aHC and pHC based on a recent coordinate-based definition (Poppenk, Evensmoen, Moscovitch, & Nadel, 2013). The left aHC was defined as the portion anterior to $y = -21$ mm within the left HC region of the AAL atlas and the pHC was defined as the portion posterior to $y = -21$ mm. This coordinate incorporates the uncus apex, a landmark for long-axis segmentation of the hippocampus (Poppenk et al., 2013).

Given the parsing functions ascribed to the left IFG by the morpho-phonological decomposition model of morphology and its involvement in the processing of inflectionally complex forms (Tyler et al., 2005, Marslen-Wilson & Tyler, 2007, Bozic et al., 2010, Bozic & Marslen-Wilson, 2010), a left IFG ROI was

defined anatomically by combining pars opercularis (BA44) and pars triangularis (BA45) of the AAL atlas. MarsBaR (<http://marsbar.sourceforge.net/>) was used to extract mean BOLD parameter estimate values within each ROI for each condition and each participant. Repeated-measures ANOVAs with affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular) assessing the main effects and interactions of these factors were fitted to the data extracted.

8.2.5.3. RSA analysis

RSA analyses were performed on realigned data without normalisation or smoothing, generating native space statistical maps. An additional first-level model was constructed for each participant that contained the same regressors as the first-level model described above for the univariate analysis, except that regressors for individual words were used for item-specific modelling. This resulted in 80 regressors per participant per run reflecting the 72 words (18 Day 8 plurals, 18 Day 9 plurals, 36 singulars (18 associated with Day 8 plurals; 18 associated with Day 9 plurals)) and the remaining 8 regressors from the univariate model (2 for the test trials and the 6 motion parameters). For each of the 72 words, single-subject T-statistic maps were estimated for the contrast of word onset compared to the implicit baseline provided by the null events, averaged over the six scanning runs. T-statistic maps were used so that the effect sizes were weighted by their error variance, which reduces the influence of large but highly variable response estimates for multivariate analyses (Misaki, Kim, Bandettini, & Kriegeskorte, 2010).

These resulting T-statistic maps were submitted to RSA analyses (Kriegeskorte et al., 2008) using the RSA toolbox (Nili et al., 2014). As discussed in Chapter 6, RSA involves testing whether the observed similarity in the multivariate brain activity patterns elicited by items belonging to specific conditions (neural RDM) corresponds to the hypothesised similarity among these items (model RDM). Neural RDMs were generated by computing the dissimilarity ($1 - \text{Pearson correlation across voxels}$) of T-statistics between all pairs of plurals. In a searchlight analysis (Kriegeskorte, Goebel, & Bandettini, 2006), data were extracted from native space T-statistic maps masked with a grey-matter mask generated during segmentation in order to restrict the analysis to

grey matter voxels. Data were extracted from spherical searchlights with an 8-mm radius (with a voxel size of 3 x 3 x 3.75mm, i.e. a maximum of 65 voxels per sphere). This was repeated for all searchlight locations in the grey matter mask. The similarity between the neural RDM and each of the model RDMs was computed using a Spearman correlation for each searchlight location, and the resulting correlation coefficient returned to the voxel at the centre of the searchlight. This resulted in a Spearman correlation map for each participant in each grey-matter voxel.

To assess searchlight similarity values across participants at the second-level, the Spearman correlation maps for each participant were Fisher-z-transformed to conform to normality assumptions, normalised to MNI space, and spatially smoothed with a Gaussian kernel of 10mm with full width half maximum. Similarity values from searchlights within the same ROIs as described in the univariate analysis were extracted using MarsBaR. For each ROI and for each participant, one Fisher-z-transformed Spearman correlation value for each model tested was obtained. Second-level analyses used repeated-measures ANOVAs similar to those described for the univariate analysis. A detailed description of the model RDMs tested against the neural RDMs is provided in the RSA results section 8.3.3. Where significant results were obtained, post-hoc one-sample t-tests were performed to test whether the Fisher-z-transformed Spearman correlations were significantly greater than zero.

8.2.5.4. PPI analysis

PPI analyses were performed on the realigned, normalised, and smoothed data after preprocessing as described in the univariate analysis section and run using the CONN toolbox for SPM (version 17.a) (<http://www.nitrc.org/projects/conn>; (Whitfield-Gabrieli & Nieto-Castanon, 2012). The CONN toolbox implements a generalised form of PPI (gPPI; <http://www.nitrc.org/projects/gppi>; McLaren, Ries, Xu, & Johnson, 2012), which allows modelling all condition main effects and interactions simultaneously in a single model compared to the standard implementation in SPM which requires using separate models for each condition tested. Thus, this approach is particularly well-suited for experimental designs with several conditions. Data were analysed using a GLM including the same 22 event types as in the univariate model, each convolved with the canonical SPM haemodynamic response. Nuisance covariates

including cerebrospinal fluid and white-matter signals and their derivatives were estimated and regressed out following the CompCor method (Behzadi, Restom, Liao, & Liu, 2007). A linear detrending term was also added, eliminating the need for global signal regression (Chai, Castañón, Ongür, & Whitfield-Gabrieli, 2012; Murphy, Birn, Handwerker, Jones, & Bandettini, 2009). The six subject-specific motion parameters, the main effect of task condition, and their first order derivatives were also included as potential confounds.

Following parameter estimation of the first level model, an ROI-to-ROI analysis was performed in which functional connectivity between the BOLD signal of two ROIs (one seed ROI and one target ROI) is computed. The ROIs corresponded to the five ROIs used in the univariate and RSA analyses, except that they were defined as 6mm spheres. The spheres were drawn around the centre of mass of each ROI: left aSTG ($x = -56, y = -6, z = -1$), left pSTG ($x = -61, y = -28, z = 10$), left aHC ($x = -28, y = -12, z = -20$), left pHC ($x = -28, y = -24, z = -12$), and left IFG ($x = -47, y = 24, z = 14$). This analysis allowed examining the influence of the experimental factors on the connectivity between neocortical ROIs as well as between hippocampal and neocortical ROIs. Second-level analyses examined functional connectivity between ROIs for the main effects and interactions involving affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular).

To further examine whether each ROI was functionally connected to other brain regions, a seed-to-voxel analysis was performed in which functional connectivity between the BOLD signal in a seed ROI and the BOLD signal in every brain voxel is computed. In a second-level analysis, functional connectivity between each seed ROIs (left aSTG, left pSTG, left aHC, left pHC, left IFG) and every brain voxel was examined for the main effects and interactions involving affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular). To help visualising and interpreting significant effects, the REX toolbox (<http://www.nitrc.org/projects/rex>) was used to extract the mean functional connectivity values for each condition and for each participant within significant clusters. For the ROI-to-ROI analysis, statistics reported in tables were thresholded at

$p < .001$ uncorrected and FDR-corrected for multiple comparisons among the different seeds. For the seed-to-voxel analyses, all statistical parametric maps and statistics reported in tables were thresholded at $p < .001$ uncorrected at the voxel level and $p < .05$ FWE corrected at the cluster level.

8.3. Results

8.3.1. Behavioural results

8.3.1.1. Training

Accuracy on the picture naming tasks presented at the end of each training block in each training session was analysed to examine how participants' mastery of the new words improved over the course of training (Figure 22). Singulars and plurals were analysed separately as there was no plural training on Day 1, there were only 3 training blocks for plurals on Day 8 and Day 9 (compared to 4 training blocks for singulars), and because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix.

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous), training day (Day1, Day 8, Day 9), training block (Block 1, Block 2, Block 3, Block 4), and their interaction was fitted. The three-way interaction between condition, training day, and training block was not significant. The two-way interaction between training day and training block was significant, $\chi^2_{(6)} = 133.66$, $p < .001$. Participants improved significantly across training blocks on Day 1 (Block 2 vs. Block 1: $b = 1.78$, $z = 19.54$, $p < .001$, Block 3 vs. Block 2: $b = 0.81$, $z = 10.43$, $p < .001$, Block 4 vs. Block 3: $b = 0.49$, $z = 5.59$, $p < .001$). On Day 8, participants only improved from the second to the third training block (Block 3 vs. Block 2: $b = 0.92$, $z = 3.14$, $p = .002$). On Day 9, participants only improved from the third to the fourth training block (Block 4 vs. Block 3: $b = 0.64$, $z = 2.17$, $p = .030^+$). The two-way interactions between condition and training block and between condition and training day were not significant. The main effect of condition was significant, $\chi^2_{(3)} = 10.30$, $p = .016$. Participants were more accurate on phonologically consistent (regular) than phonologically diverse singulars ($b = 0.17$, $z = 2.11$, $p = .035^+$), on phonologically consistent (irregular) than phonologically diverse singulars ($b = 0.23$, z

= 2.84, $p = .004$), and on phonologically consistent (irregular) than on phonologically ambiguous singulars ($b = 0.18$, $z = 2.17$, $p = .030^\dagger$). The main effect of training block was significant, $\chi^2_{(3)} = 1426.30$, $p < .001$. Participants became increasingly more accurate over the course of the training blocks (Block 2 vs. Block 1, $b = 1.33$, $z = 17.84$, $p < .001$, Block 3 vs. Block 2, $b = 0.81$, $z = 11.08$, $p < .001$, Block 4 vs. Block 3: $b = 0.52$, $z = 6.32$, $p < .001$). The main effect of training day was significant, $\chi^2_{(2)} = 1654.80$, $p < .001$. Participants were more accurate on Day 9 compared to Day 1 ($b = 2.36$, $z = 24.50$, $p < .001$), on Day 8 compared to Day 1 ($b = 2.87$, $z = 25.50$, $p < .001$) and on Day 8 compared to Day 9 ($b = 0.50$, $z = 3.65$, $p < .001$).

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), training block (Block 2, Block 3, Block 4), and their interaction was fitted. The three-way interaction between condition, training day, and training block was not significant. None of the two-way interactions were significant. The main effect of training block was significant, $\chi^2_{(2)} = 52.59$, $p < .001$. Participants improved across training blocks (Block 3 vs. Block 2: $b = 0.61$, $z = 4.46$, $p < .001$, Block 4 vs. Block 3: $b = 0.44$, $z = 2.64$, $p = .008$). The main effect of condition was significant, $\chi^2_{(4)} = 19.70$, $p < .001$. Participants were significantly more accurate on phonologically diverse than phonologically ambiguous (regular) ($b = 0.52$, $z = 3.04$, $p = .002$) and phonologically ambiguous (irregular) ($b = 0.45$, $z = 2.60$, $p = .009^\dagger$) plurals, on phonologically consistent (regular) than on phonologically ambiguous (regular) ($b = 0.75$, $z = 3.52$, $p < .001$) and phonologically ambiguous (irregular) ($b = 0.68$, $z = 3.17$, $p = .002$) plurals, and on phonologically consistent (irregular) than on phonologically ambiguous (regular) ($b = 0.40$, $z = 2.04$, $p = .041^\dagger$) plurals. The main effect of training day was not significant.

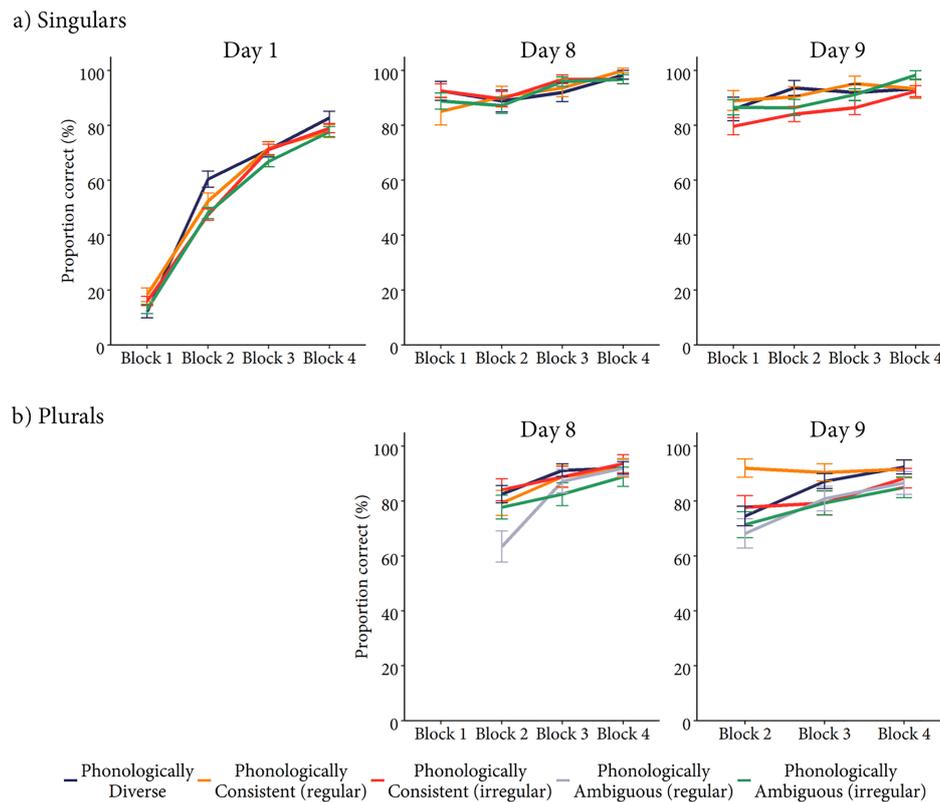


Figure 22. Experiment 2: training accuracy

Mean accuracy on the picture naming task presented at the end of each training block (Block 1, Block 2, Block 3, Block 4) on each training day (Day 1, Day 8, and Day 9) for a) singulars and b) plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (light blue line) and those associated with an irregular (green line) affix. All phonologically ambiguous singulars are depicted by a green line.

8.3.1.2. Testing

Accuracy on the picture naming (Figure 23) and 2AFC recognition memory (Figure 24) tasks presented in a testing block at the end of each training session was analysed to ensure that words with different phonological consistencies and training frequencies (plurals) were learned equally well. Singulars and plurals were analysed separately as there was no plural training on Day 1, there were only 3 training blocks for plurals on Day 8 and Day 9 (compared to 4 blocks for singulars), and because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix.

8.3.1.2.1. Picture naming

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous), training day (Day 1, Day 8, Day 9), and their interaction

was fitted. The interaction between condition and training day was not significant. The main effect of condition was not significant. The main effect of training day was significant, $\chi^2_{(2)} = 78.92$, $p < .001$. Participants were more accurate on Day 8 compared to Day 1 ($b = 2.47$, $z = 5.54$, $p < .001$) and on Day 9 compared to Day 1 ($b = 1.60$, $z = 4.97$, $p < .001$) but performed similarly on Day 8 and Day 9.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction was fitted. The interaction between condition and training day was not significant. The main effect of condition was significant, $\chi^2_{(4)} = 9.85$, $p = .043$. Participants were more accurate on phonologically diverse than on phonologically ambiguous (irregular) ($b = 1.11$, $z = 2.98$, $p = .003$) plurals, on phonologically consistent (irregular) compared to phonologically ambiguous (irregular) ($b = 0.96$, $z = 2.18$, $p = .029^{\dagger}$) plurals, and on phonologically consistent (regular) than on phonologically ambiguous (irregular) ($b = 0.84$, $z = 1.96$, $p = .049^{\dagger}$) plurals. The main effect of training day was not significant.

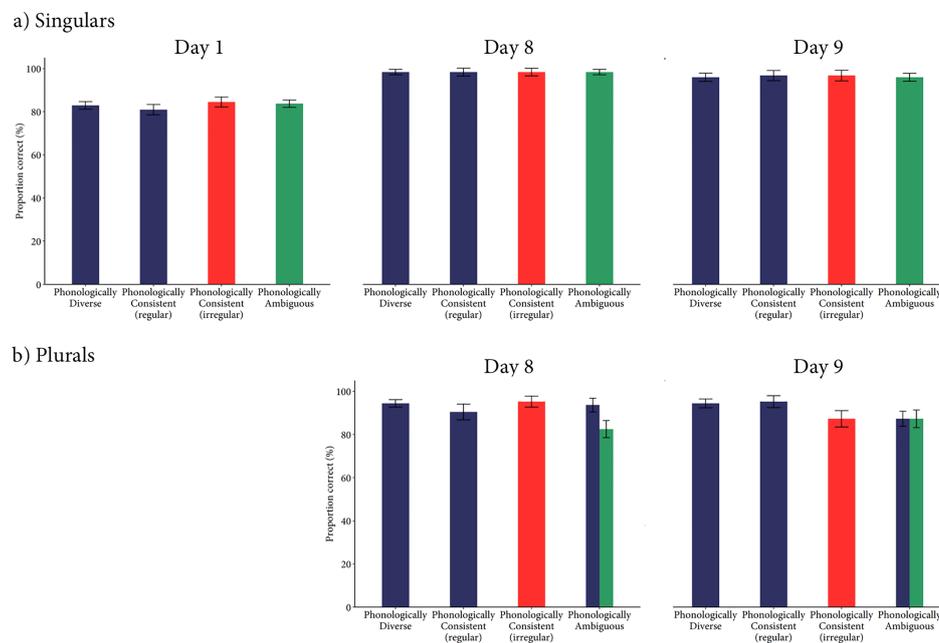


Figure 23. Experiment 2: test accuracy (picture naming)

Mean accuracy on the picture naming task for each training day (Day 1, Day 8, Day 9) for a) singulars and b) plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (green bar) affix.

8.3.1.2.2. 2AFC recognition memory

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous), training day (Day 1, Day 8, Day 9), and their interaction was fitted. Condition and training day did not interact significantly. There was no significant main effect of condition. The main effect of training day was significant, $\chi^2_{(2)} = 6.35$, $p = .042$. Participants were marginally more accurate on Day 8 compared to Day 1 ($b = 0.98$, $z = 1.71$, $p = .088$), more accurate on Day 9 compared to Day 1 ($b = 1.20$, $z = 1.96$, $p = .050^\dagger$) but performed similarly on Day 8 and Day 9.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction was fitted. Condition and training day did not interact significantly. The main effect of condition was significant, $\chi^2_{(4)} = 12.13$, $p = .016$. Participants were more accurate on phonologically diverse than on phonologically ambiguous (regular) ($b = 1.29$, $z = 2.03$, $p = .042^\dagger$) plurals. The main effect of training day was not significant.

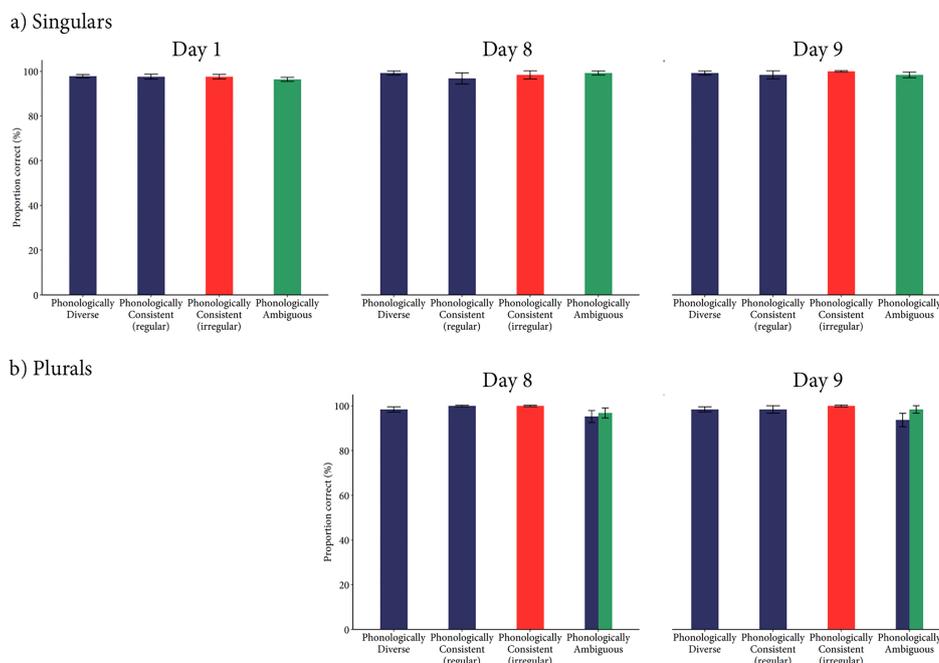


Figure 24. Experiment 2: test accuracy (2AFC)

Mean accuracy on the 2AFC task for each training day (Day 1, Day 8, Day 9) for a) singulars and b) plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (green bar) affix.

8.3.1.3. Generalisation

8.3.1.3.1. Response type analysis

The frequency of each response type (regular generalisations, irregular generalisations (consistent), irregular generalisations (ambiguous), incorrect) for novel phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), and phonologically ambiguous singulars requiring participants to use affixes trained on Day 8 or Day 9 are shown in Figure 25. A loglinear model with phonological consistency (diverse, consistent (regular), consistent (irregular), ambiguous), affix training day (Day 8, Day 9), response type (regular generalisations, irregular generalisations (consistent), irregular generalisations (ambiguous), incorrect) was fitted. The three-way interaction between phonological consistency, affix training day, and response type was significant, $\chi^2_{(9)} = 19.98$, $p = .018$. For phonologically diverse novel singulars, the interaction between affix training day and response type was significant, $\chi^2_{(3)} = 14.93$, $p = .002$. Participants produced more regular generalisations for Day 9 compared to Day 8 affixes ($b = 0.25$, $z = 1.97$, $p = .049^+$) and more incorrect responses for Day 8 compared to Day 9 affixes ($b = 0.80$, $z = 3.17$, $p = .002$). For phonologically consistent (regular) novel singulars, the interaction between affix training day and response type was significant, $\chi^2_{(3)} = 41.31$, $p < .001$. Participants produced more regular generalisations for Day 9 compared to Day 8 affixes ($b = 0.31$, $z = 2.52$, $p = .012$) and more incorrect responses for Day 8 compared to Day 9 affixes ($b = 1.46$, $z = 4.74$, $p < .001$). For phonologically consistent (irregular) novel singulars, the interaction between affix training day and response type was significant, $\chi^2_{(3)} = 22.32$, $p < .001$. Participants produced more irregular generalisations (ambiguous) for Day 9 compared to Day 8 affixes ($b = 1.01$, $z = 2.45$, $p = .014^+$) and more incorrect response for Day 8 compared to Day 9 affixes ($b = 0.93$, $z = 3.60$, $p < .001$). For phonologically ambiguous novel singulars, the interaction between affix training and response type was significant, $\chi^2_{(3)} = 35.25$, $p < .001$. Participants produced more irregular generalisations (consistent) for Day 9 compared to Day 8 affixes ($b = 1.61$, $z = 2.08$, $p = .038^+$) and more incorrect responses for Day 8 compared to Day 9 affixes ($b = 1.41$, $z = 4.54$, $p < .001$).

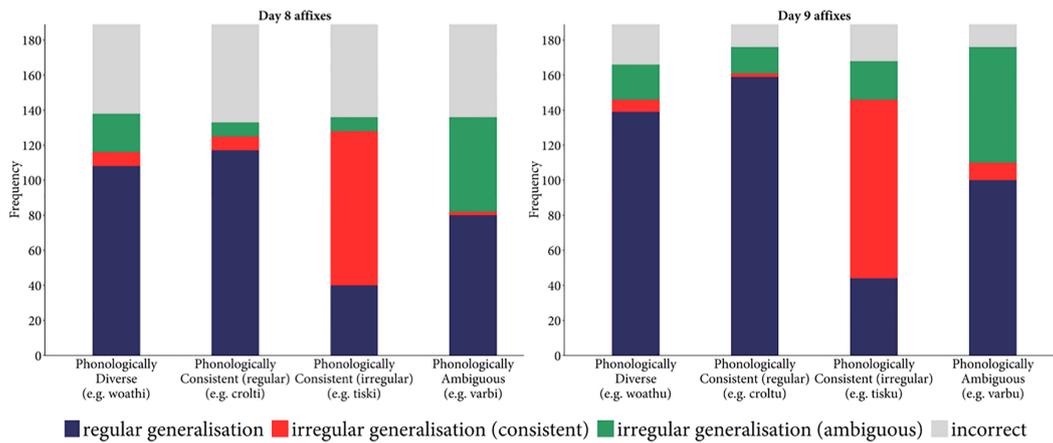


Figure 25. Experiment 2: generalisation results

Frequency of each response type supplied within each phonologically consistency category of novel singulars requiring participants to use affixes trained on Day 8 (left) and on Day 9 (right).

To mitigate potential confounds due to a difference in the number of correct responses for Day 8 compared to Day 9 affixes, the analysis was repeated with a random sub-sample of correct responses for Day 9 affixes to artificially match the number of correct responses produced for Day 8 affixes (Figure 26). The three-way interaction between phonological consistency, affix training day, and response type was significant, $\chi^2_{(6)} = 13.27$, $p = .039$. For phonologically diverse, phonologically consistent (regular), and phonologically ambiguous novel singulars, there was no significant interaction between affix training day and response type. For phonologically consistent (irregular) novel singulars, there was a marginal two-way interaction between affix training day and response type, $\chi^2_{(2)} = 5.12$, $p = .077$. Participants produced more irregular generalisations (ambiguous) for Day 9 compared to Day 8 affixes ($b = 0.87$, $z = 2.05$, $p = .040^\dagger$).

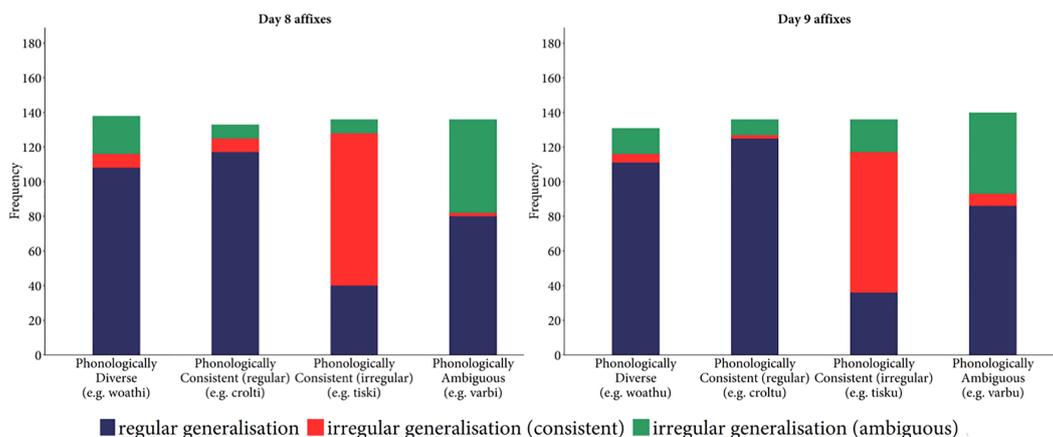


Figure 26. Experiment 2: generalisation results for equalised correct responses

Frequency of each response type supplied within each phonologically consistency category of novel words requiring participants to use affixes of Day 8 (left) and on Day 9 (right).

8.3.1.3.2. Error analysis

A secondary analysis focused on the errors produced by participants. Incorrect responses were subcategorised as a failure to provide any response (no response), the production of a legal but gender-inappropriate affix (gender error), and the production of an illegal affix (illegal affix) (Figure 27). A loglinear model with phonological consistency (diverse, consistent (regular), consistent (irregular), ambiguous), affix training day (Day 8, Day 9), error type (no response, gender error, illegal affix), and their interaction was fitted. There was no significant three-way interaction between phonological consistency, affix training day, and error type. The two-way interaction between affix training day and error type was significant, $\chi^2_{(2)} = 7.84$, $p = .020$. To break down this interaction, differences between the number of each error type were examined for Day 8 and Day 9 affixes separately. For Day 8 affixes, the main effect of error type was significant, $\chi^2_{(2)} = 10.05$, $p = .007$. Participants produced more no response than illegal affix errors ($b = 0.45$, $z = 2.53$, $p = .011$) and more gender than illegal affix errors ($b = 0.52$, $z = 2.90$, $p = .004$) but a similar number of gender and no response errors. For Day 9, the main effect of error type was not significant. There were no significant two-way interactions between affix training day and phonological consistency and between error type and phonological consistency. The main effect of phonological consistency was not significant.

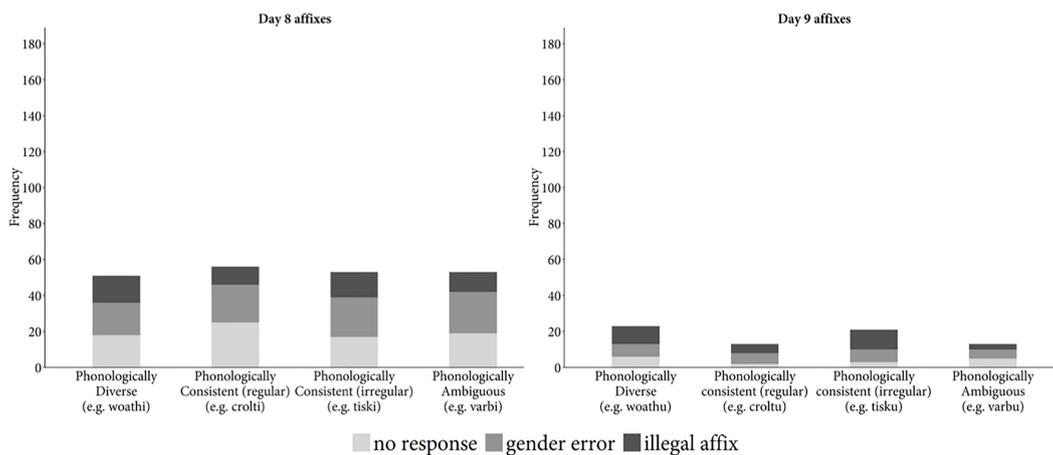


Figure 27. Experiment 2: generalisation errors

Frequency of each error type (no response, gender error, illegal affix) supplied within each phonologically consistency category of novel singulars requiring participants to use affixes trained on Day 8 (left) or on Day 9 (right).

8.3.1.3.3. 3AFC

Due to scheduling problems, 3 participants did not have time to complete the 3AFC generalisation task before the start of their planned scanning session so the analysis reported contains data from 18 participants. The frequency of each response type (regular generalisations, irregular generalisations (consistent), and irregular generalisations (ambiguous), incorrect) for novel phonologically diverse, consistent (regular), consistent (irregular), and ambiguous singulars requiring participants to use affixes trained on Day 8 or Day 9 are shown in Figure 28. A loglinear model with phonological consistency (diverse, consistent (regular), consistent (irregular), ambiguous), affix training day (Day 8, Day 9), response type (regular generalisations, irregular generalisations (consistent), irregular generalisations (ambiguous), incorrect)), and their interaction was fitted. The three-way interaction between phonological consistency, affix training day, and response type was not significant. The two-way interaction between phonological consistency and response type was significant, $\chi^2_{(2)} = 125.33$, $p < .001$. To break down this interaction, the numbers of each response type were compared for each phonological consistency separately. For phonologically diverse novel singulars, participants selected more regular affixes than other affixes (regular vs. irregular (consistent): $b = 2.22$, $z = 11.83$, $p < .001$, regular vs. irregular (ambiguous): $b = 1.94$, $z = 11.07$, $p < .001$). The same pattern was true for phonologically consistent regular novel singulars (regular vs. irregular (consistent): $b = 2.70$, $z = 11.38$, $p < .001$, regular vs. irregular (ambiguous): $b = 2.51$, $z = 11.56$, $p < .001$). For phonologically consistent irregular novel singulars, participants selected more irregular (consistent) affixes than other affixes (irregular (consistent) vs. regular: $b = 1.30$, $z = 9.26$, $p < .001$, irregular (consistent) vs. irregular (ambiguous): $b = 2.33$, $z = 10.66$, $p < .001$). For phonologically ambiguous novel singulars, participants selected a similar number of regular and irregular (ambiguous) affixes. The two-way interactions between affix training day and phonological consistency and between affix training day and response type were not significant.

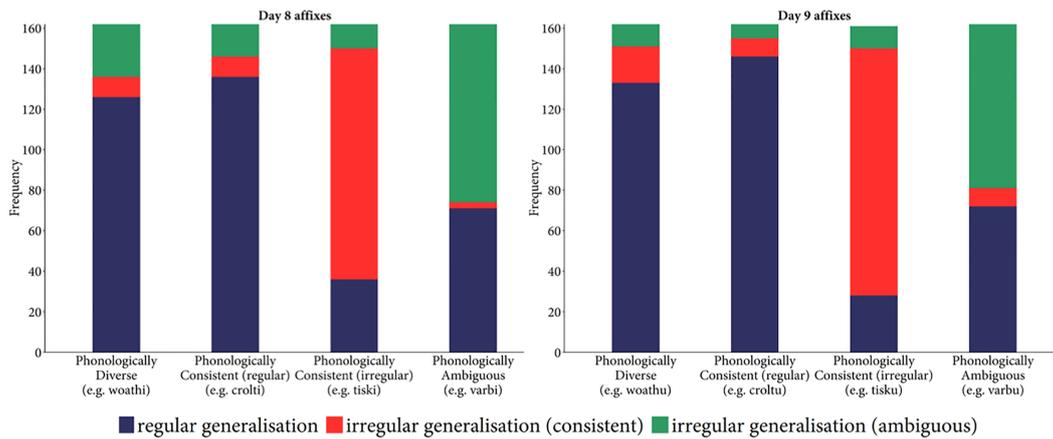


Figure 28. Experiment 2: generalisation 3AFC results

Frequency of each response type supplied within each phonologically consistency category of novel singulars requiring participants to use affixes train on Day 8 (left) and on Day 9 (right).

8.3.1.4. Scanning test trials

Accuracy on the test trials presented in the scanner is shown in Figure 29. Singulars and plurals were analysed separately because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix. For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous), test type (affix knowledge, whole word knowledge), and their interaction was fitted. The interaction between condition and test type was not significant, nor were the main effects.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)), affix training day (Day 8, Day 9), test type (affix knowledge, whole word knowledge), and their interaction was fitted. The three-way interaction between condition, affix training day, and test type was not significant. None of the two-way interactions were significant. The main effect of affix training day was significant, $\chi^2_{(1)} = 11.36$, $p < .001$. Participants were more accurate for Day 8 compared to Day 9 plurals ($b = 0.55$, $z = 3.33$, $p < .001$). The main effects of condition and test type were not significant.

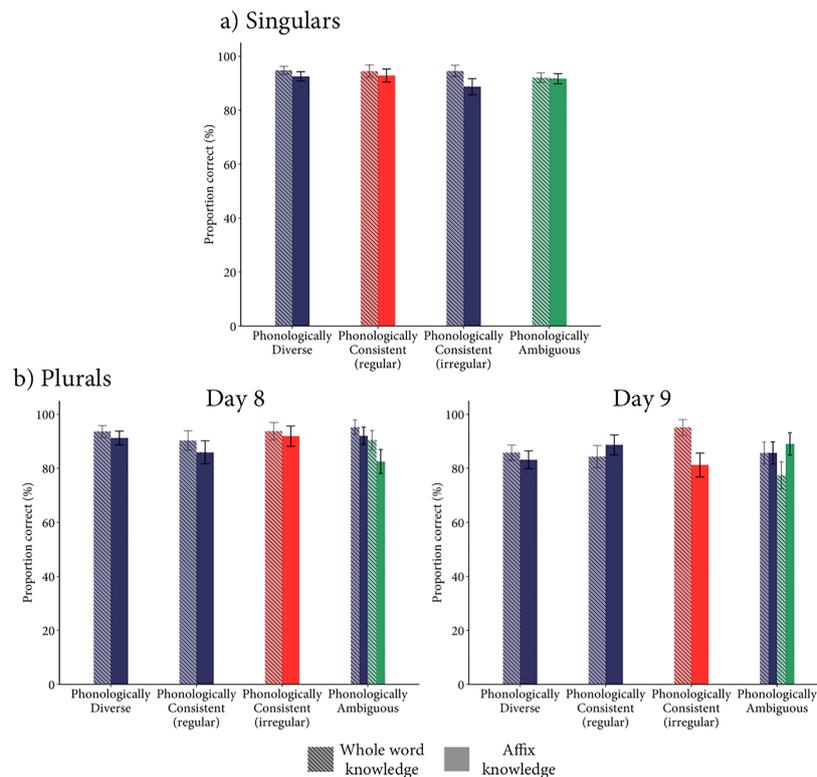


Figure 29. Experiment 2: fMRI test trials accuracy

Accuracy on the test trials presented in the scanner for a) singulars, b) plurals (Day 8, Day 9).

Experimental conditions are along the x axis. Opaque bars represent tests assessing affix knowledge and cross-hatched bars represent tests assessing whole word knowledge. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (green bar) affix.

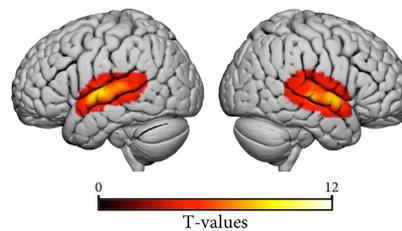
8.3.2. Univariate results

8.3.2.1. Whole-brain analysis

A first contrast tested for BOLD univariate activity associated with the processing of the plurals by subtracting the implicit baseline provided by the null events from all plurals. This contrast revealed activity in bilateral STG (Figure 30a; Appendix 1), largely consistent with the literature on spoken word and pseudoword processing. A three-way repeated-measures ANOVA with factors of affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular) revealed a significant three-way interaction in left ACC and right medial orbitofrontal cortex (Figure 30b, Appendix 1). Mean BOLD parameter estimates extracted from the peak voxel in left ACC ($x = -10, y = 22, z = 42$) showed that for irregular plurals, the interaction between affix training day and phonological consistency was significant, $F(1,20) = 7.22, p = .014$. Day 9 phonologically ambiguous irregulars elicited greater activity than Day 8 phonologically ambiguous irregulars,

$t_{(20)} = 2.73$, $p = .013$. Mean BOLD parameter estimates extracted from the peak voxel in right orbitofrontal cortex ($x = 26$, $y = 42$, $z = 2$) showed that for phonologically consistent plurals, the interaction between affix training day and affix regularity was significant, $F_{(1,20)} = 7.04$, $p = .015$. Day 9 phonologically consistent regulars elicited greater activity than Day 8 phonologically consistent regulars, $t_{(20)} = 2.64$, $p = .016$. For phonologically ambiguous plurals, the interaction between affix training day and affix regularity was also significant, $F_{(1,20)} = 5.40$, $p = .031$. Day 9 phonologically ambiguous irregulars elicited greater activity than Day 8 phonologically ambiguous irregulars, $t_{(20)} = 2.41$, $p = .026$. There were no significant clusters for the two-way interactions or main effects.

a) all plurals > baseline



b) phonological consistency x affix regularity x affix training day interaction

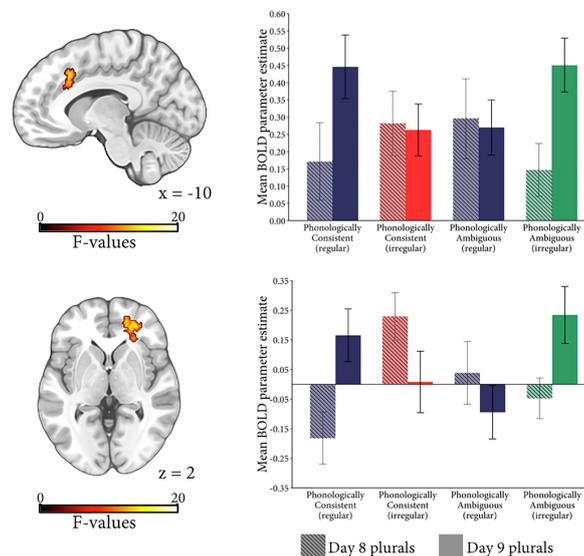


Figure 30. Experiment 2: univariate whole-brain analysis results

a) all plurals > baseline b) mean BOLD parameter estimates extracted from the peak voxel in left ACC ($x = -10$, $y = 22$, $z = 42$) and right medial orbitofrontal cortex ($x = 26$, $y = 42$, $z = 2$). The factorial crossing of phonological consistency and affix regularity is reported along the x axis (phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)). Blue bars correspond to plurals taking a regular affix, red bars correspond to plurals taking an irregular (consistent) affix, and green bars correspond to plurals taking an irregular (ambiguous) affix. Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

8.3.2.1. ROI analysis

A three-way repeated-measures ANOVA with affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular) conducted on the mean BOLD parameter estimates extracted from the left STG, left pSTG, left aHC, left pHC, and left IFG revealed not significant main effects or interactions in any of the ROIs (Figure 31).

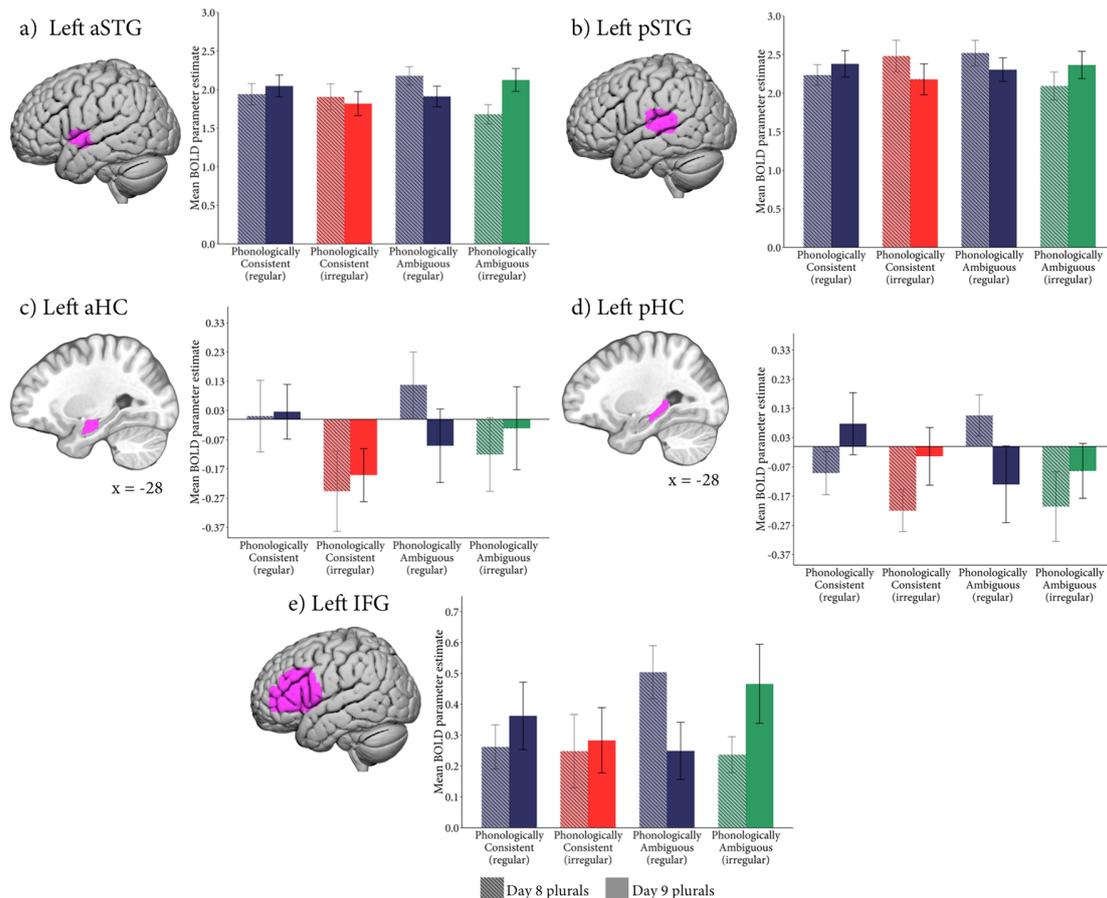


Figure 31. Experiment 2: univariate ROI analysis results

Mean BOLD parameter estimate extracted from the a) left aSTG, b) left pSTG, c) left aHC, d) left pHC, and e) left IFG for each experimental condition. The factorial crossing of phonological consistency and affix regularity is reported as four conditions along the x axis (phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)). Blue bars correspond to plurals taking a regular affix, red bars correspond to plurals taking an irregular affix, and green bars correspond to plurals taking an irregular (ambiguous) affix. Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

8.3.3. RSA results

The univariate results did not reveal any effects of affix training day, phonological consistency, or affix regularity on the neural responses to the plurals in any of the ROIs where overnight changes were expected. As reviewed in Chapter 6, RSA has

increased sensitivity in detecting differences between experimental stimuli than univariate analyses. This is particularly relevant here given the subtle differences between plurals. Thus, RSA was used to perform more specific tests of the similarity structure underlying the plurals. The analyses focused on the same ROIs as the univariate analysis: the left aSTG, left pSTG, left aHC, left pHC, and left IFG.

Eight model RDMs were initially constructed to examine whether phonological consistency and affix regularity influenced the similarity structure of the plurals. Four model RDMs focused on the similarity structure among Day 8 plurals and another four focused on the similarity structure among Day 9 plurals such that effects of overnight consolidation could be assessed. The model RDMs tested for increased similarity structure between pairs of plurals sharing the same phonological consistency and affix regularity (e.g. *skol**taff*** vs. *grol**taff***) compared to pairs of plurals having a different phonological consistency and affix regularity (e.g. *skol**taff*** vs. *tar**bimm***) (Figure 32). More specifically, the model RDMs tested for increased similarity structure among plurals *within* each of the experimental conditions: Day 8 phonological consistent (regular) (e.g. *skol**taff***, *grol**taff***, *wol**taff***), Day 8 phonologically consistent (irregular) (e.g. *hisk**esh***, *lisk**esh***, *fisk**esh***), Day 8 phonologically ambiguous (regular) (e.g. *har**ba**ff*, *yar**ba**ff*, *nar**ba**ff*), Day 8 phonologically ambiguous (irregular) (e.g. *tar**bimm***, *clar**bimm***, *slar**bimm***), Day 9 phonological consistent (regular) (e.g. *zol**topp***, *nol**topp***, *quol**topp***), Day 9 phonologically consistent (irregular) (e.g. *tisk**ull***, *vis**kull***, *yisk**ull***), Day 9 phonologically ambiguous (regular) (e.g. *jar**bo**pp*, *lar**bo**pp*, *glar**bo**pp*), Day 9 phonologically ambiguous (irregular) (e.g. *blar**bo**ot*, *mar**bo**ot*, *far**bo**ot*). The model RDMs also tested for decreased similarity structure *between* experimental conditions (e.g. phonological consistent (regular) (e.g. *skol**taff***, *grol**taff***, *wol**taff***) vs. phonologically ambiguous (irregular) (e.g. *tar**bimm***, *clar**bimm***, *slar**bimm***)). In each model RDM testing for similarity structure among plurals within one experimental condition, every other within-condition similarity comparisons were excluded and replaced with NaN values. The model RDMs also excluded and replaced with NaN values comparisons between plurals having the same plural affix but a different phonological consistency (e.g. *skol**taff*** vs. *har**ba**ff*) and plurals having the same phonological consistency but a different affix (e.g. *har**ba**ff* vs. *tar**bimm***). The model RDMs thus focused on the influence of affix training day, phonological

consistency, and affix regularity on the similarity structure underlying plurals having similar phonological and morphological form but different lexico-semantics.

Analysing similarity of multivariate neural patterns for similar but non-identical plurals allowed merging all six scanning runs into a single analysis which reduces noise in the estimation of T-statistic maps and makes the analysis insensitive to other forms of lexico-semantic similarity that could lead to similar neural representations of identical plurals (Correia, Jansma, & Bonte, 2015, see Blank & Davis, 2016 for a similar approach). The main diagonal of each model RDM coding for similarity between identical plurals was therefore excluded and replaced with NaN values.

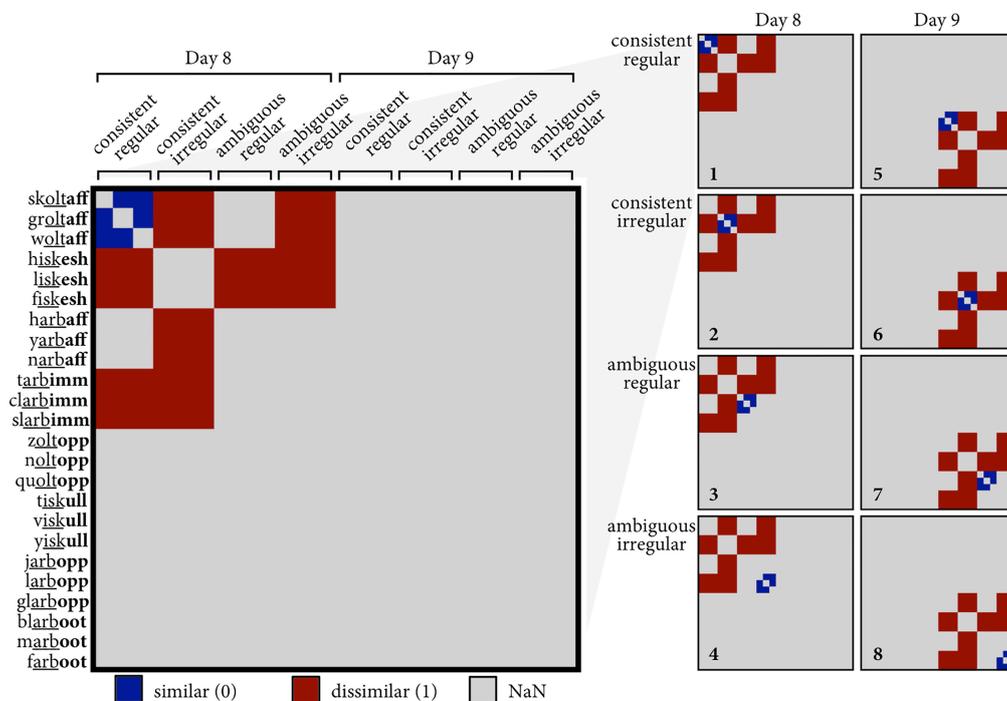


Figure 32. Experiment 2: RSA models RDMs (factorial design)

The eight matrices on the right were used to test similarity between plurals sharing the same phonological consistency and the same affix regularity within each condition (Day 8 consistent regular (model 1), Day 8 consistent irregular (model 2), Day 8 ambiguous regular (model 3), Day 8 ambiguous irregular (model 4), Day 9 consistent regular (model 5), Day 9 consistent irregular (model 6), Day 9 ambiguous regular (model 7), and Day 9 ambiguous irregular (model 8)). Similarity between pairs of identical plurals (the main diagonal) was excluded (“NaN” values in grey). Similarity between plurals sharing the same affix but having a different phonological consistency (i.e. consistent regular vs. ambiguous regular) and between plurals sharing the same consistency but different affixes (i.e. ambiguous regular vs. ambiguous irregular) was also excluded. The similarity among plurals belonging to only one condition was tested in each model so all other within-condition similarity comparisons were excluded (“NaN” values in grey). Within-condition similarity was calculated for each affix training day separately so comparisons of plurals from the opposite day were also excluded (“NaN” values in grey). Similarity structure among plurals sharing the same phonological consistency and affix regularity was predicted in blue (zeroes depicted in blue), but not for plurals having a different phonological consistency and a different affix regularity (ones depicted in red). The 8 models RDMs were correlated with observed multivariate patterns of neural responses to the plurals (i.e. neural RDM).

Fisher-z-transformed correlation coefficients between neural RDMs and each of the 8 model RDMs were extracted from searchlight locations within the left aSTG, left pSTG, left aHC, left pHC, and left IFG ROIs (Figure 33). As each of the 8 model RDMs assessed similarity structure among plurals corresponding to each of the experimental conditions, the mean Fisher-z-transformed correlation coefficients for each model for each ROI were submitted to a repeated-measures ANOVA with affix training day (Day 8 [models 1, 2, 3, 4], Day 9 [models 6, 7, 8, 9]), phonological consistency (consistent [models 1, 2, 5, 6], ambiguous [models 3, 4, 7, 8]), and affix regularity (regular [models 1, 3, 5, 7], irregular [models 2, 4, 6, 8]) (see Figure 32). For the left aSTG, the two-way interaction between affix training day and affix regularity was significant, $F_{(1,20)} = 17.09$, $p = .001$. Collapsing across phonological consistency, there was greater within-condition similarity structure for plurals sharing a regular affix (i.e. phonologically consistent (regular), phonologically ambiguous (regular) [models 1, 3, 5, 7]) for Day 8 compared Day 9 plurals, $t_{(20)} = 3.24$, $p = .004$. There was also greater within-condition similarity structure for plurals sharing an irregular affix (i.e. phonologically consistent (irregular), phonologically ambiguous (irregular) [models 2, 4, 6, 8]) for Day 9 compared to Day 8 plurals, $t_{(20)} = 3.70$, $p = .001$. One-sample t-tests revealed that similarity structure was significantly greater than zero for Day 8 regulars (collapsed across phonological consistency), $t_{(20)} = 1.82$, $p = .042$, and for Day 9 irregulars, $t_{(20)} = 3.69$, $p = .001$. None of the main effects or other interactions were significant. For the left pSTG, the two-way interaction between affix training day and affix regularity was also significant, $F_{(1,20)} = 7.96$, $p = .010$. Collapsing across phonological consistency, there was greater within-condition similarity structure for plurals sharing a regular affix for Day 8 compared to Day 9 plurals, $t_{(20)} = 2.35$, $p = .030$. There was also greater within-condition similarity structure for plurals sharing an irregular affix Day 9 compared to Day 8 plurals, $t_{(20)} = 2.51$, $p = .021$. One-sample t-tests revealed that similarity structure was marginally greater than zero for Day 8 regulars (collapsed across phonological consistency), $t_{(20)} = 1.52$, $p = .073$, and significantly greater than 0 for Day 9 irregulars, $t_{(20)} = 3.66$, $p = .001$. None of the main effects or other interactions were significant. For the left IFG, there was also a significant two-way interaction between affix training day and affix regularity, $F_{(1,20)} = 9.59$, $p = .006$. Collapsing across phonological consistency, there was greater within-

condition similarity structure for plurals sharing an irregular affix for Day 9 compared to Day 8 plurals, $t_{(20)} = 4.71$, $p < .001$. A one-sample t-test showed that similarity structure was significantly greater than 0 for Day 9 irregulars, $t_{(20)} = 4.52$, $p < .001$. For the left aHC, there was a significant interaction between affix training day and affix regularity, $F_{(1,20)} = 6.67$, $p = .008$. Collapsing across phonological consistency, there was greater within-condition similarity structure for plurals sharing a regular affix for Day 8 compared to Day 9 plurals, $t_{(20)} = 2.44$, $p = .024$. There was also greater within-condition similarity structure for plurals sharing an irregular affix for Day 9 compared to Day 8 plurals, $t_{(20)} = 2.29$, $p = .033$. However, one sample t-tests showed that neither the similarity structure for regulars or irregulars was significantly greater than 0. There were no significant main effects or interactions in the left pHC.

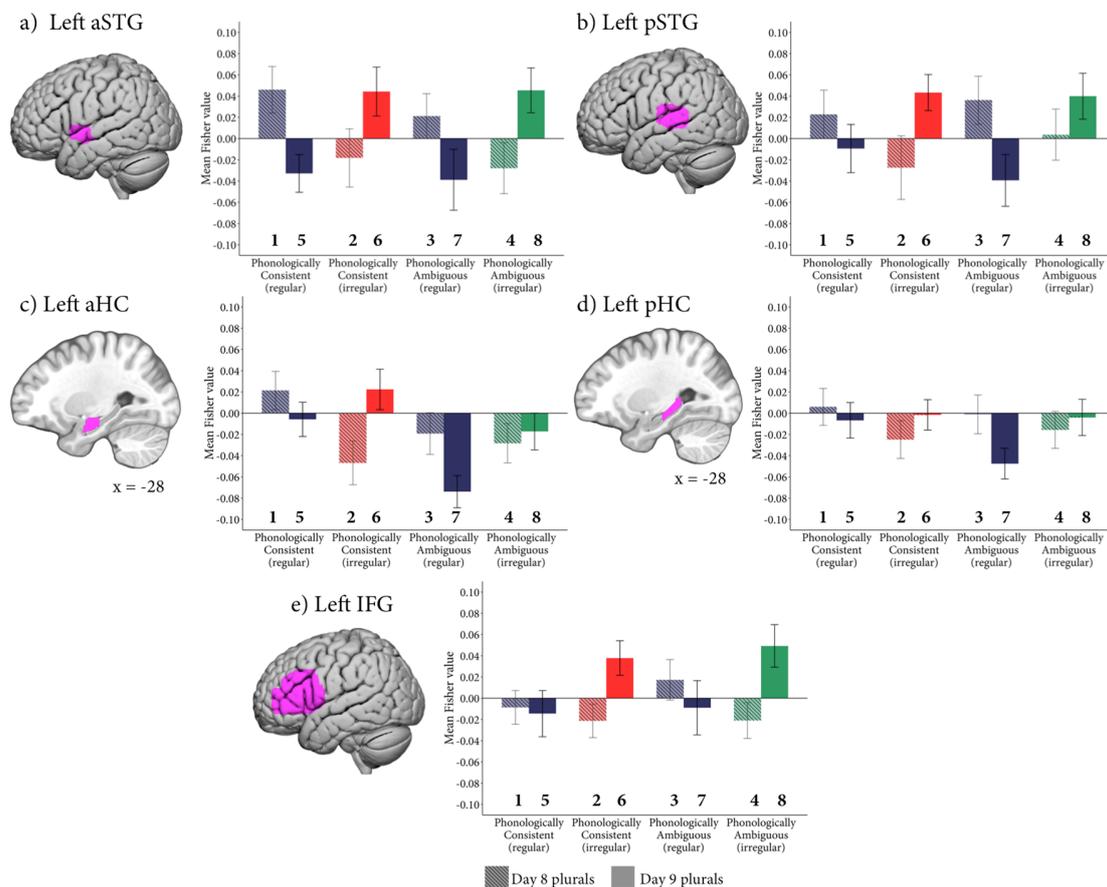


Figure 33. Experiment 2: RSA ROI analysis results (factorial design)

Mean Fisher-z-transformed Spearman correlation coefficient extracted from the a) left aSTG, b) left pSTG, c) left aHC, d) left pHC, and e) left IFG for each model RDM (numbered 1 to 8 on the graphs). The factorial crossing of phonological consistency and affix regularity is depicted as four conditions along the x axis (phonologically consistent (regular), phonologically consistent (irregular), phonologically ambiguous (regular), phonologically ambiguous (irregular)). Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

To follow up on the significant two-way interaction between affix training day and affix regularity found in the left aSTG, left pSTG, and left IFG, two additional model RDMs were constructed to examine whether affix-similarity rather than whole-form-similarity accounted for the increase in similarity structure among plurals sharing a regular affix. In other words, these model RDMs tested for similarity based on type frequency. One model RDM tested for increased similarity between all pairs of high type frequency Day 8 plurals sharing the same regular affix but having phonologically diverse stems (e.g. *gleetaff*, *skoltaff*, *harbaff*) and dissimilarity between pairs of plurals having different affixes and phonological consistencies (e.g. *hiskesh*, *tarbimm*). Another model RDM tested for the same similarity structure among all pairs of high type frequency Day 9 plurals (Figure 33). A significantly better fit of this model RDM for Day 8 compared to Day 9 plurals would be consistent with the increase in similarity structure among plurals sharing a regular affix observed in the initial analysis being driven by affix-similarity rather than whole-form-similarity.

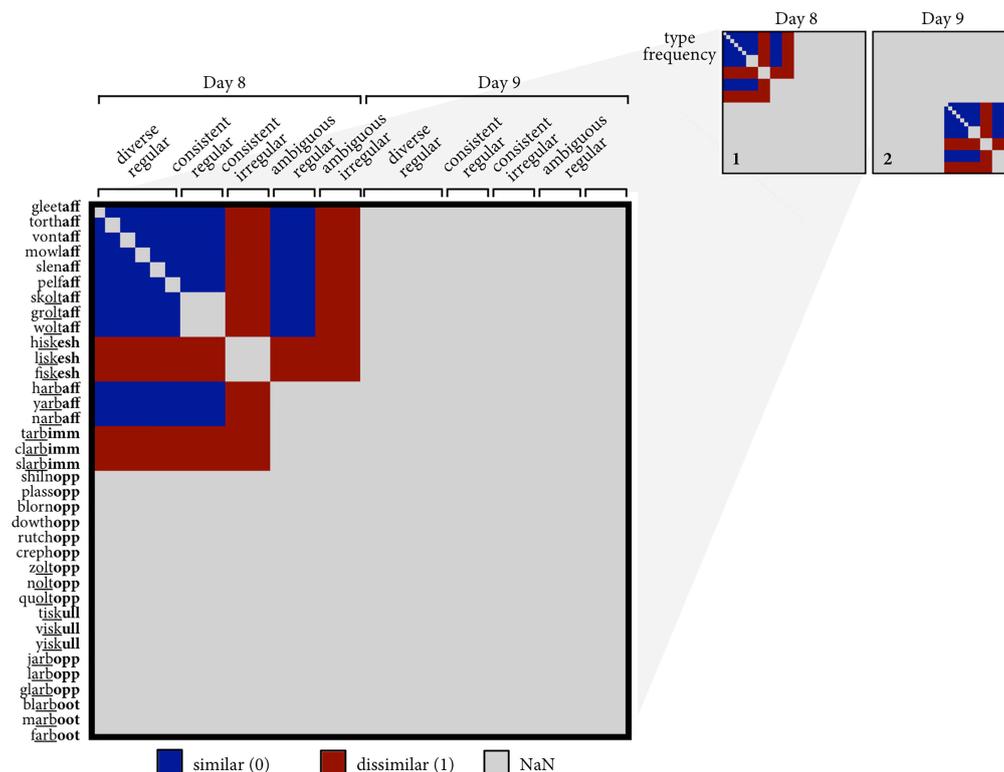


Figure 34. Experiment 2: RSA model RDMs (type frequency)

The two matrices on the right tested similarity between plurals sharing the same regular affix. Similarity between pairs of identical plurals (the main diagonal) was excluded (“NaN” values depicted in grey). Similarity structure among plurals sharing the same regular affix was predicted (zeroes depicted in blue), but not for plurals sharing a different affix (ones depicted in red). The 2 models RDMs were correlated with observed multivariate patterns of neural responses to the plurals (i.e. neural RDM).

For the left aSTG, this analysis showed a significant main effect of affix training day, $F(1,20) = 5.37$, $p = .031$. There was greater similarity structure for high type frequency regulars for Day 8 compared Day 9 plurals, $t(20) = 2.32$, $p = .031$. A one-sample t-test showed that similarity structure was significantly greater than 0 for Day 8 plurals, $t(20) = 1.87$, $p = .038$. There was no significant main effect of affix training day in the left pSTG or in the left IFG (Figure 34).

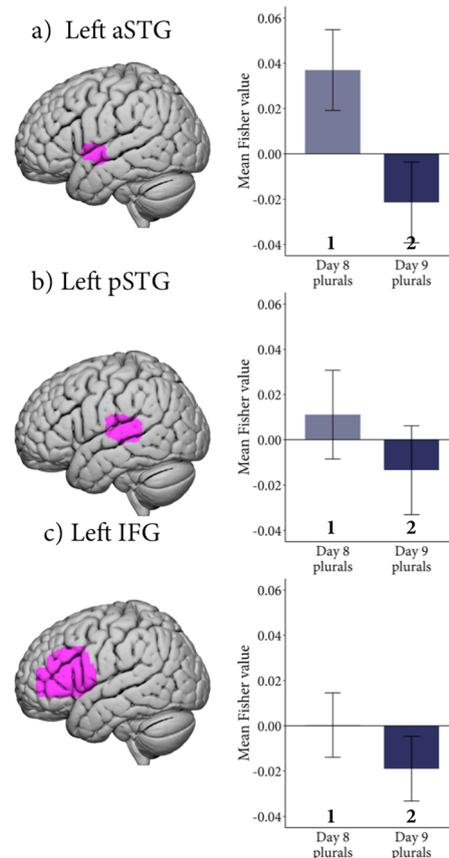


Figure 35. Experiment 2: RSA ROI analysis results (type frequency)

Mean Fisher-z-transformed Spearman correlation coefficient extracted from the a) left aSTG, b) left pSTG, and c) left IFG for models testing for increased similarity structure based on type frequency. Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

As outlined at the outset of the chapter, overnight changes in similarity structure were predicted to be mostly apparent for phonologically ambiguous plurals. This prediction would have led to a three-way interaction between affix training day, phonological consistency, and affix regularity in the analysis of the initial 8 model RDMS (i.e. factorial design, see Figure 32). However, these models excluded comparisons between phonologically ambiguous (regular) (e.g. *harbaff*) and

phonologically ambiguous (irregular) (e.g. *tarbimm*). For completeness and to address the prediction that phonologically ambiguous regular and irregular plurals may become more dissimilar following consolidation, two additional model RDMs were considered. One model RDM tested for increased similarity between pairs of phonologically ambiguous plurals (e.g. *harbaff*, *tarbimm*) and dissimilarity between pairs of plurals sharing a different phonological consistency and affix regularity (e.g. *harbaff* vs. *zoltopp*) (Figure 36). A significantly better fit of this model RDM for Day 9 compared to Day 8 plurals would be consistent with the prediction that dissimilarity between phonologically ambiguous regulars and irregulars may increase with overnight consolidation. This analysis revealed no main effect of affix training day or similarity structure significantly greater than 0 in any ROI.

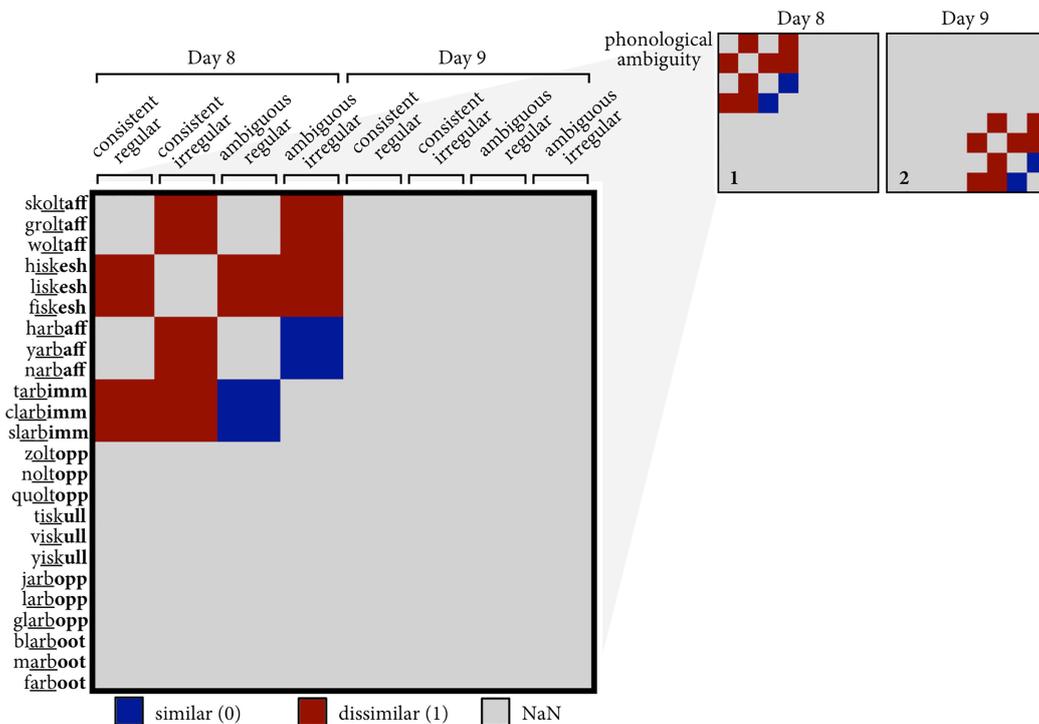


Figure 36. Experiment 2: RSA model RDMS (phonological ambiguity)

The two matrices on the right tested similarity between phonologically ambiguous plurals. Similarity between pairs of identical plurals (the main diagonal) was excluded (“NaN” values depicted in grey). Similarity structure among phonologically ambiguous plurals was predicted (zeroes depicted in blue), but not for plurals having different a different phonological consistency and affix regularity (ones depicted in red). The 2 models RDMS were correlated with observed multivariate patterns of neural responses to the plurals (i.e. neural RDM).

8.3.4. PPI results

8.3.4.1. ROI-to-ROI analysis

An ROI-to-ROI analysis examined functional connectivity between the left aSTG, left pSTG, left aHC, left pHC, and left IFG ROIs (see methods section 8.2.5.4. for seed definition details) for main effects and interactions involving affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular) (Appendix 2). Functional connectivity within the hippocampus was not considered (i.e. between the left aHC and left pHC) since spurious motion-related variance tends to be more similar at nearby voxels than distant voxels (Power, Schlaggar, & Petersen, 2014) and because no clear predictions regarding intra-hippocampal patterns of connectivity could be derived from the emerging research examining issues of intra-regional connectivity. None of the ROIs were functionally connected for the three-way interaction, two-way interactions, or main effects of phonological consistency and affix regularity. For the main effect of affix training day, the left aHC was more functionally connected with the left pSTG for Day 9 than for Day 8 plurals, $t_{(20)} = 3.08$, $p = .003$ (left pSTG – left aHC, $t_{(20)} = 3.05$, $p = .003^{12}$).

In light of the overnight increase in similarity structure among high type frequency regulars observed in the RSA analysis, a further contrast examined the functional connectivity between ROIs for all high type frequency regulars (i.e. phonologically diverse, phonologically consistent (regular), phonological ambiguous (regular)) trained on Day 8 and Day 9. This revealed that the left aSTG was more functionally connected to the left IFG, $t_{(20)} = 3.18$, $p = .003$ (left IFG – left aSTG, $t_{(20)} = 3.14$, $p = .004$) and to the left pHC, $t_{(20)} = 3.19$, $p = .002$ (left pHC – left aSTG, $t_{(20)} = 3.18$, $p = .003$) for Day 8 compared to Day 9 regulars. In addition to being more functionally connected to the left aSTG and the left pHC, the left IFG was more functionally connected to the left pSTG, $t_{(20)} = 1.89$, $p = .036$, (left pSTG – left IFG, $t_{(20)} = 1.89$, $p = .037$) for Day 8 compared to Day 9 regulars.

¹² In PPI/gPPI analyses, the connectivity values between ROIs are non-symmetrical because the underlying models/measures are also non-symmetrical (Whitfield-Gabrieli & Nieto-Castanon, 2012).

8.3.4.2. *Seed-to-voxel analysis*

To determine whether the seed ROIs showed greater functional connectivity with other brain regions, a seed-to-voxel analysis was performed for all main effects and interactions involving affix training day (Day 8, Day 9), phonological consistency (consistent, ambiguous), and affix regularity (regular, irregular). There were no significant effects for the three-way interaction for any seed. There were no significant effects for the two-way interactions between phonological consistency and affix regularity and between affix training day and phonological consistency for any seed.

For the two-way interaction between affix training day and affix regularity, the left aSTG was functionally connected to two clusters of voxels in the left middle frontal gyrus (MFG) (Appendix 3). Mean condition-specific connectivity values extracted within each significant cluster revealed greater functional connectivity between the left aSTG and the left MFG for Day 9 compared to Day 8 regular plurals (cluster 1: $t_{(20)} = 4.27$, $p < 0.001$; cluster 2: $t_{(20)} = 5.83$, $p < .001$) and greater functional connectivity between the left aSTG and the left MFG for Day 8 compared to Day 9 irregular plurals (cluster 1: $t_{(20)} = 3.19$, $p = 0.005$; cluster 2: $t_{(20)} = 2.87$, $p = .009$) (Figure 37). The left pSTG was also functionally connected to a cluster of voxels in the left MFG (Appendix 3). Mean condition-specific connectivity values extracted within the significant cluster revealed greater functional connectivity between the left pSTG and the left MFG for Day 8 compared to Day 9 irregular plurals, $t_{(20)} = 5.10$, $p < 0.001$, but only marginally more functional connectivity between the left pSTG and the left MFG for Day 9 compared to Day 8 regular plurals, $t_{(20)} = 1.84$, $p = 0.081$ (Figure 37). There were no significant effects for the left aHC, left pHC, or left IFG.

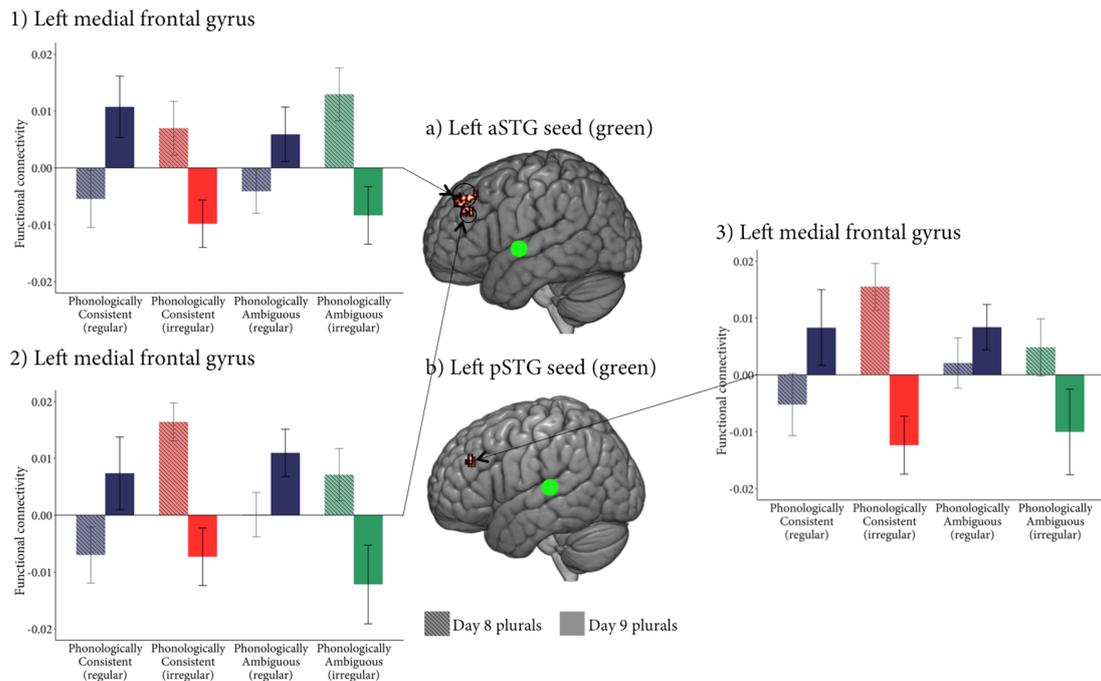


Figure 37. Experiment 2: PPI seed-to-voxel analysis affix regularity x affix training day interaction
 Mean functional connectivity values extracted for each condition reported along the x axis for significant clusters showing an interaction effect between affix training day and affix regularity. a) the left aSTG seed was functionally connected to two clusters of voxels in the left MFG in 1) and 2). b) the left pSTG seed was also connected to a cluster of voxels in left MFG in 3). Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

There were no significant effects for the main effect of affix training day for any seed. For the main effect of phonological consistency, the left aSTG was more functionally connected with two clusters of voxels in the left inferior parietal lobule (IPL) and one cluster of voxels in right inferior temporal gyrus (ITG) for phonologically consistent than phonologically ambiguous plurals (Appendix 3; Figure 38). No significant effects were observed for the opposite contrast (i.e. ambiguous > consistent). The left pSTG was more functionally connected with a cluster of voxels in the left IPL for phonologically consistent than phonologically ambiguous plurals (Appendix 3; Figure 38). No significant effects were observed for the opposite contrast (i.e. ambiguous > consistent). The left aHC was more functionally connected with a cluster of voxels in left IFG for phonologically consistent compared to phonologically ambiguous plurals. The opposite contrast showed that the left aHC was more functionally connected with a cluster of voxels in the right supramarginal gyrus for phonologically ambiguous compared to phonologically consistent plurals (Appendix

3; Figure 38). No significant effects were found for the main effect of phonological consistency for the left pHC and left IFG seeds.

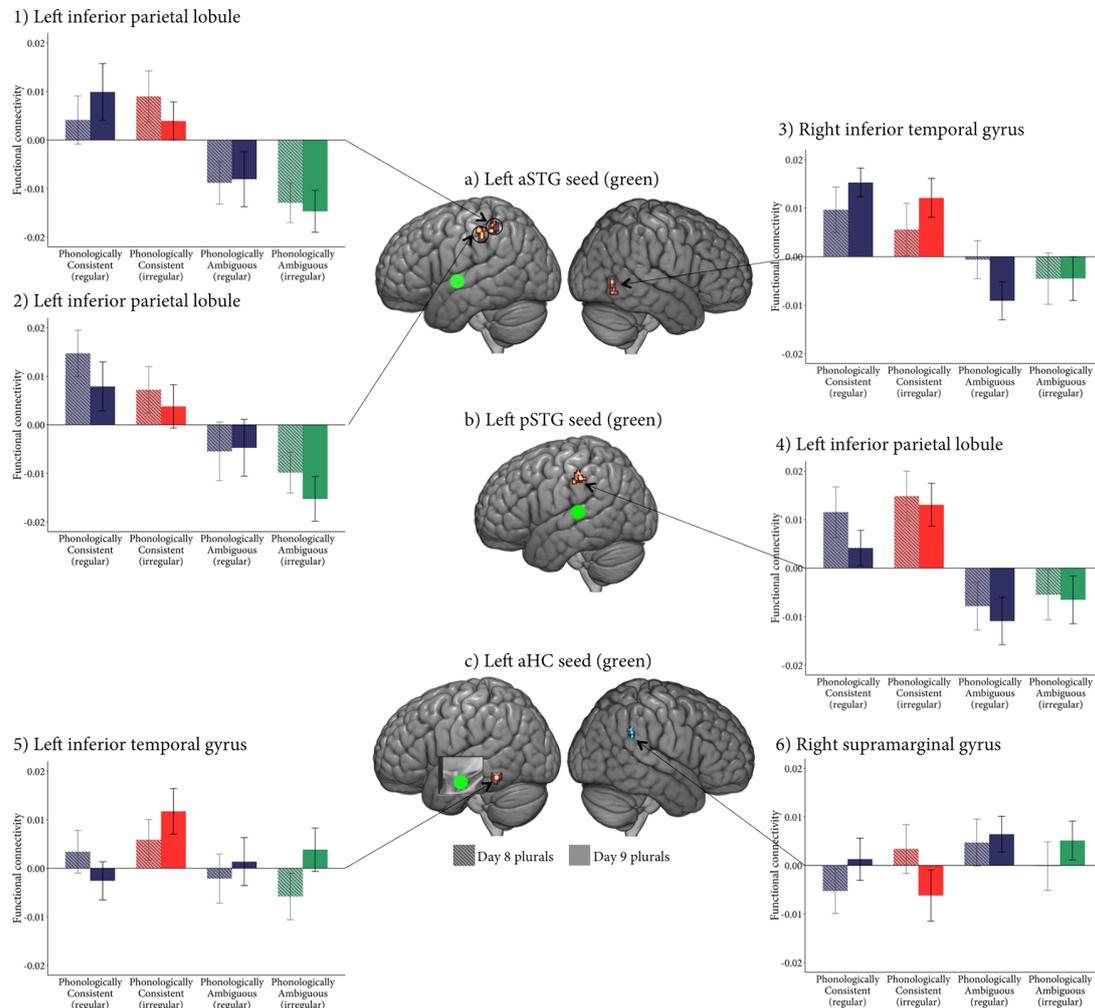


Figure 38. Experiment 2: PPI seed-to-voxel analysis main effect of phonological consistency

Mean functional connectivity values extracted for each condition reported along the x axis for significant clusters showing a main effect of phonological consistency. a) the left aSTG seed was more functionally connected to two clusters of voxels in the left IPL in 1) and 2) and one cluster of voxels in 3) right ITG for phonologically consistent than phonologically ambiguous plurals. b) the left pSTG was more functionally connected to a cluster of voxels in 4) the left IPL for phonologically consistent than phonologically ambiguous plurals. c) the left aHC was more functionally connected to a cluster of voxels in 5) left ITG for phonologically consistent than phonologically ambiguous plurals and more functionally connected to a cluster of voxels in 6) right supramarginal gyrus for phonologically ambiguous than phonologically consistent plurals. Significant clusters for the phonologically consistent > phonologically ambiguous contrasts are shown in red. The significant cluster for the opposite contrast (phonologically ambiguous > phonologically consistent) is shown in blue. Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

For the main effect of affix regularity, the left aHC was more functionally connected to clusters of voxels in left lingual gyrus, left superior occipital gyrus, and right fusiform gyrus for irregular compared to regular plurals (Appendix 3, Figure 39).

No significant effects were observed for the opposite contrast (i.e. regulars > irregulars). The left IFG was more functionally connected to clusters of voxels in the left IPL and right supramarginal gyrus for irregular compared to regular plurals (Appendix 3, Figure 39). No significant effects were observed for the opposite contrast (i.e. regulars > irregulars). No significant effects were found for the left aSTG, left pSTG, and left pHC seeds. Additional contrasts examining all high type frequency regulars trained on Day 8 and Day 9 revealed no significant effects for any seed.

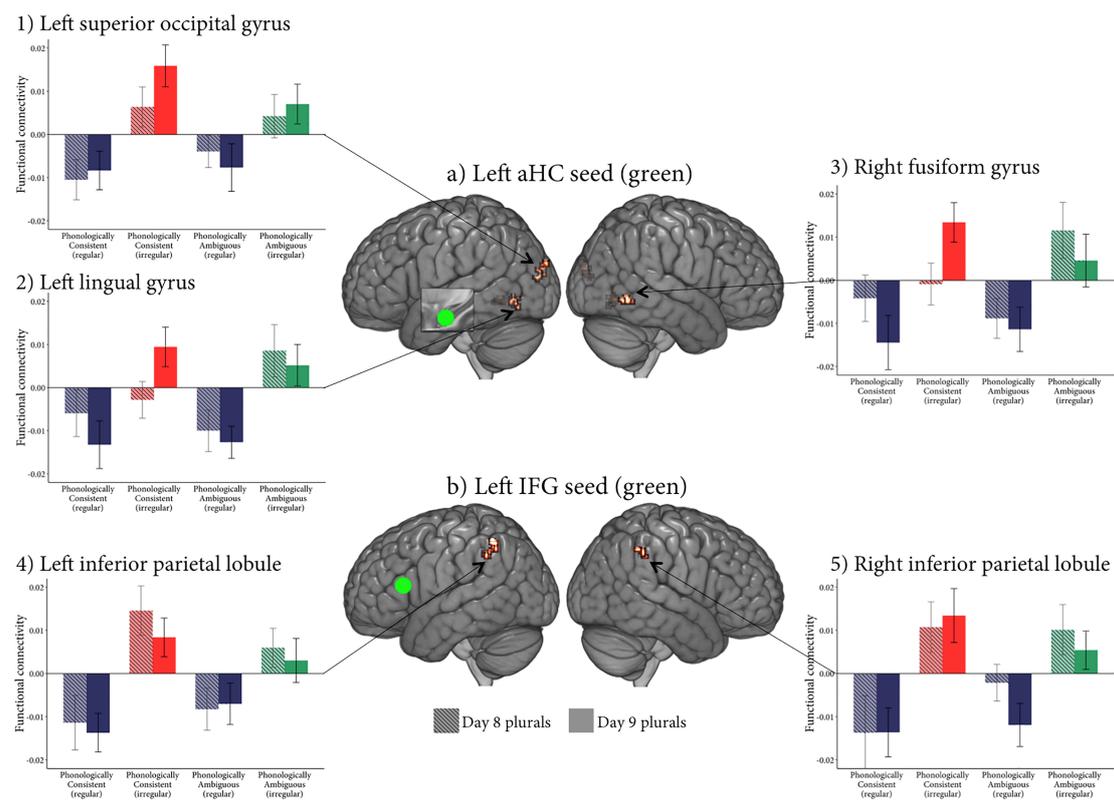


Figure 39. Experiment 2: PPI seed-to-voxel analysis main effect of affix regularity

Mean functional connectivity values extracted for each condition reported along the x axis for significant clusters showing a main effect of affix regularity. a) the left aHC seed was more functionally connected to clusters of voxels in 1) the left superior occipital gyrus, 2) the left lingual gyrus, and 3) the right fusiform gyrus for irregular than regular plurals. b) the left IFG seed was more functionally connected to clusters of voxels in the 4) left IPL and 5) right IPL for irregular than regular plurals. Cross-hatched bars represent Day 8 plurals, opaque bars represent Day 9 plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

8.4. Discussion

8.4.1. Learning and generalisation

Experiment 2 combined the behavioural paradigm developed in Experiment 1 with RSA and PPI analyses of fMRI data to characterise the neural representations and

functional networks underlying the newly-learned plurals. Behavioural results from the training sessions largely replicated those obtained in Experiment 1 and further confirmed the efficacy of the training schedule in ensuring that participants developed high proficiency on the novel morphological system. Across the picture naming and 2AFC recognition memory tasks presented at the end of each training session, participants performed similarly for both sets of plurals as evidenced by a lack of significant main effects or interactions involving affix training day. As such, differences in the generalisation of Day 8 and Day 9 plural affixes to novel singulars could not be attributed to differences in initial learning. Similarly, the neural representations underlying Day 8 and Day 9 plurals could be compared for plurals that were initially learned equally well.

The increase in irregular generalisations (ambiguous) following overnight consolidation for novel phonologically ambiguous singulars found in Experiment 1 and reported by Mirković and Gaskell (in prep.) was not replicated. Instead, participants produced a similar number of regular generalisations and irregular generalisations (ambiguous) for both Day 8 and Day 9 affixes. For novel phonologically consistent (irregular) and phonologically ambiguous singulars, participants produced fewer irregular generalisations (ambiguous) and irregular generalisations (consistent), respectively, for Day 8 compared to Day 9 affixes. These latter two effects could reflect the strengthening of the representations underlying the irregulars following overnight consolidation manifesting in a reduction of inappropriate irregular generalisations. A similar observation was made in Experiment 1, where an increase in irregular generalisations (consistent) and irregular generalisations (ambiguous) was observed for Day 8 compared to Day 9 affixes. It is conceivable that a similar effect manifested here as a reduction of inappropriate irregular generalisations. Further data would be required to substantiate this proposal.

The reduction in irregular generalisations (ambiguous) remained marginally significant for novel phonologically consistent (irregular) singulars after equalising the number of correct responses across the two training days suggesting that the effect was not confounded, at least not entirely, by forgetting of Day 8 affixes or interference from Day 9 affixes. For novel phonologically diverse and phonologically consistent

(regular) singulars, participants produced more regular generalisations for Day 9 compared to Day 8 affixes. However, this result was not replicated when equalising the number of correct responses across the two training days suggesting that forgetting of Day 8 affixes or interference from Day 9 affixes may have, at least partially, confounded the effect.

To address the fact that participants produced fewer correct responses for Day 8 compared to Day 9 affixes in the plural elicitation generalisation task, an additional 3AFC generalisation task was included in Experiment 2. Performance on this additional generalisation task showed that participants were sensitive to the phonological consistency of the novel singulars, supplying more appropriate than inappropriate generalisations given the phonological cue contained in the novel singulars. However, no effects of overnight consolidation were observed. It should be noted that the generalisation items in the 3AFC generalisation task were the same as in the plural elicitation generalisation task such that re-exposure to the generalisation items may have confounded the results. A more detailed consideration of the discrepancy in generalisation behaviour between Experiments 1 and 2 is the focus of Experiment 3, reported in Chapter 9.

Despite the absence of overnight changes in generalisation behaviour, participants were more accurate for Day 8 compared to Day 9 plurals on test trials presented in the scanner, assessing affix and whole-word knowledge. These results are consistent with previous work suggesting that overnight consolidation plays a role in the development of lexicalised representations for newly-learned words and affixes indexed by improved recognition performance (Gaskell & Dumay, 2003; Dumay & Gaskell, 2007, 2012; Tamminen & Gaskell, 2008; Davis et al., 2009; Tamminen et al., 2010; Henderson et al., 2012; Merx et al., 2011; Tamminen et al., 2012, 2015; Leminen et al., 2016). Thus, whilst there were no overnight changes in generalisation behaviour, improved performance on test trials presented in the scanner for Day 8 compared to Day 9 plurals may reflect more 'optimal' representations following overnight consolidation (Stickgold & Walker, 2013).

8.4.2. Neural representations and functional networks

Turning to the fMRI results, the univariate analysis showed that the processing of the plurals engaged bilateral STG largely consistent with the existing literature on the processing of both words and pseudowords (Newman & Twieg, 2001; Kotz, Cappa, von Cramon, & Friederici, 2002). It is particularly noteworthy that the peak activation found in left STG ($x = -56, y = -6, z = 0$) falls squarely on the centre of mass of the anterior STG cluster ($x = -56, y = -6, z = -1$) reported by Davis & Gaskell (2009), in a meta-analysis of 11 neuroimaging studies in which responses to pseudowords were compared to responses to words. It is this aSTG cluster that revealed the most compelling results in RSA and PPI analyses.

The whole-brain univariate analysis revealed a three-way interaction between affix training day, phonological consistency, and affix regularity in the left ACC and right orbitofrontal cortex. This interaction reflected greater activity in response to phonologically ambiguous irregulars for Day 9 compared to Day 8 plurals in the left ACC and right orbitofrontal cortex and greater activity in response to phonologically consistent regulars for Day 9 compared to Day 8 plurals in right orbitofrontal cortex. The ACC and prefrontal cortex have been implicated in conflict monitoring and resolution (Barch, Braver, Sabb, & Noll, 2000; Braver, Barch, Gray, Molfese, & Snyder, 2001; Botvinick, Cohen, & Carter, 2004). The involvement of the ACC in the processing of regular English verbs (Tyler et al., 2005; Stamatakis et al., 2005) has been suggested to reflect increased demands on its monitoring functions as inflectionally complex forms need to be decomposed. The three-way interaction could suggest a shift in the demands put on the monitoring functions of the ACC and prefrontal cortex with overnight consolidation. However, this result does not provide information regarding the underlying representations of the newly-learned plurals or how they may change with overnight consolidation. Furthermore, analyses focused on the response profile of the hippocampus and neocortical representational areas where effects of affix training day, phonological consistency, and affix regularity were expected did not reveal any main effects or interactions.

8.4.1.2. Overnight changes in neural representations

More informative results came from the RSA analysis, which revealed an interaction between affix training day and affix regularity in the left aSTG, left pSTG, and left IFG. In both STG regions, this interaction was characterised by an increase in similarity structure among regulars and a decrease in similarity structure among irregulars following overnight consolidation. Follow-up analyses testing whether this increase in similarity structure was driven by type frequency was shown to fit uniquely in the left aSTG. This finding is in line with the CLS model in which encountering several items sharing a common feature (here, an affix) supports the development of similarity structure (McClelland et al., 1995; O'Reilly & Norman, 2002; McClelland, 2013; O'Reilly et al., 2014; Kumaran et al., 2016). It is also consistent with the purported role of sleep in supporting the development of overlapping representations that underlie schema formation (here, a regular inflectional pattern) (Lewis & Durrant, 2011). A role for type frequency and consolidation in supporting the development of generalisable affix knowledge has previously been reported in behavioural experiments (Tamminen et al., 2012, 2015), but their influence on underlying neural representations was largely unknown. Previous evidence for overnight changes in the neural representations of newly-learned affixes in the left STG came from an MEG study, which lacked the spatial resolution of fMRI and sensitivity of RSA (Leminen et al., 2016). Here, neural evidence for the role of overnight consolidation and type frequency in supporting the development of an affix representation is provided for the first time.

The fact that neural representations consistent with the development of an affix representation were restricted to the left aSTG is significant, since it potentially relates to its role in representing lexical and semantic knowledge and mediating access to these representations (Scott & Johnsrude, 2003; Rogers et al., 2004; Scott, 2005; Patterson, Nestor, & Rogers, 2007; Lambon Ralph, Pobric, & Jefferies, 2009). The present results suggest, in line with previous neural network simulations and behavioural work, that an affix shared by several words (i.e. high type frequency) facilitates the development of an affix representation. They further suggest, in line with the CLS model and previous behavioural work, that overnight consolidation may

be necessary for such a representation to emerge. The role of the aSTG in representing lexical and semantic knowledge remains controversial and a perhaps more established view is that such functions are subserved by posterior middle temporal regions (Hickock & Poeppel, 2004, 2007). Neither account, however, makes specific predictions about the representations underlying newly-learned words or indeed about the role of consolidation in supporting the development of such representations. Thus, future data on the role of anterior temporal regions in supporting morphological learning will be valuable.

Another line of work has implicated anterior temporal regions in sentence (Stowe et al., 1998; Friederici, Meyer, & von Cramon, 2000; Vandenberghe, Nobre, & Price, 2002; Humphries et al., 2006; Rogalsky & Hickok, 2008), phrase (Bemis & Pylkkänen, 2011), and, more recently, morphological (Brooks & Cid de Garcia, 2015) processing leading to the hypothesis that it may be involved in syntactic structure-building (Friederici & von Cramon, 2000; Grodzinsky & Friederici, 2006) or incremental semantic composition (Stowe, Haverkort, & Zwarts, 2005; Vandenberghe et al., 2002). Thus, an alternative explanation for the current results could be that high type frequency regulars engage combinatorial processes more strongly following overnight consolidation. However, model RDMS coding for similarity structure among high type frequency regulars did not show a significant fit in the left IFG, which has consistently been associated with combinatorial processes in the context of inflectional morphology in both univariate and multivariate fMRI analyses (Bozic et al., 2010, 2015). While it is difficult to argue from a null effect, the present results appear to be most clearly in line with the similarity structure in the left aSTG reflecting the development of an affix representation. Chapter 10 returns to and expands on this suggestion.

There was initial similarity structure among Day 9 irregulars, which was not observed for Day 8 irregulars in the left aSTG, left pSTG, and left IFG. This result is surprising and contrary to the prediction that was made on the basis of the results obtained in Experiment 1 and reported by Mirković and Gaskell (in prep.) whereby increases in irregular generalisations following overnight consolidation, particularly in the case of phonologically ambiguous novel singulars, were thought to reflect

increased similarity structure among irregulars. The present results instead point to a decrease in similarity structure among irregulars following overnight consolidation. One speculative explanation for this result is that high token frequency may support the development of whole-word representations. Behavioural findings have shown that high token frequency inflected and derived words yield faster reaction times and more accurate responding than low token frequency inflected and derived words in a range of tasks across different languages (Laine, Niemi, Koivuselkä-Sallinen, & Hyönä, 1995; Sereno & Jongman, 1997; Bertram, Laine, & Karvinen, 1999; Bertram, Hyönä, & Laine, 2000; Bertram, Schreuder, & Baayen, 2000; Ford, Marslen-Wilson, & Davis, 2003; Lehtonen & Laine, 2003; Soveri, Lehtonen, & Laine, 2007). The present results could be seen to be compatible with this work and further suggest a role for overnight consolidation in strengthening whole-word representations for high token frequency irregulars. Of course, this interpretation is entirely predicated on the assumption that the lack of similarity structure among Day 8 irregulars reflects a strengthening of their whole-word representations. Further research will be needed to provide support for this proposal.

This interpretation may initially seem to support dual-route models of inflectional morphology according to which regulars rely on combinatorial processing whilst irregulars rely on whole-form storage. However, behavioural token frequency effects have previously been successfully accounted for in distributed neural networks (Davis, van Casteren, & Marslen-Wilson, 2003). Furthermore, unlike English irregulars, the irregulars used in the present study were all made up of a stem and affix such that dual-route models would predict the engagement of decompositional processes. That is, irregulars were ‘irregular’ simply by virtue of their low type frequency and high token frequency. It is precisely in those conditions that the CLS model would predict a strengthening of item-specific representations and a lack of similarity structure (cf. Chapter 3; McClelland et al., 1995; Kumaran & McClelland, 2012). On this view, it is the similarity structure before consolidation, rather than the lack of similarity structure after consolidation, which is most surprising. The current dataset does not allow distinguishing whether this initial similarity structure was driven by the whole-forms (i.e. *hiskesh*, *liskesh*, *fiskesh* all share *iskesh*) or by the affix alone. Indeed, comparisons isolating affix-based similarity, as were possible for high

type frequency regulars, were not possible for irregulars. Future investigations including high token frequency affixed forms sharing the same affix but having phonologically distinct stems (e.g. *hiskesh*, *tarbesh*) would go some way to addressing this limitation. It is plausible that high token frequency irregulars initially engage combinatorial processes across the left IFG and left STG more strongly as their low type frequency and high token frequency requires more effortful attempts to parse them into their constituent stem and affix. Saliency tags attached to high token frequency irregulars during learning may be subsequently used during sleep to enhance item-specific (i.e. whole-word) representations (Stickgold & Walker, 2013). Chapter 10 further elaborates on this proposal.

8.4.1.3. Overnight changes in functional networks

Functional connectivity between the left aHC and the left pSTG was greater for Day 9 compared to Day 8 plurals. This result is entirely consistent with predictions derived from the CLS model and previous functional connectivity work suggesting greater hippocampal involvement shortly after learning (McClelland et al., 1995; Takashima et al., 2009). The fact that the left aHC showed increased connectivity with the left pSTG is particularly interesting, since it potentially relates to the role of the aHC in memory encoding (Lepage et al., 1998; Schacter & Wagner, 1999; Spaniol et al., 2009). Recall from Chapter 2 that Davis et al. (2009) found reduced activity in the pSTG for consolidated relative to unconsolidated pseudowords. This result was interpreted as indexing more word-like phonological representations following overnight consolidation. The present results suggest that such representations may become independent of the hippocampus following overnight consolidation. This is entirely consistent with the CLS model. Stronger functional connectivity with the aHC may reflect continued encoding of Day 9 plurals.

Focusing on all high type frequency regulars revealed greater functional connectivity between the left aSTG and the left IFG for Day 8 compared to Day 9 plurals. The increase in functional connectivity between the left aSTG and left IFG is consistent with the CLS model and previous work reporting increased neocortical – neocortical functional connectivity following overnight consolidation (Takashima et al., 2009). The overnight increase in similarity structure among Day 8 regulars in the

RSA analysis was argued to reflect the development of an affix representation. Increased functional connectivity between the left aSTG and left IFG following overnight consolidation may reflect more efficient combinatorial processes following overnight consolidation once an affix representation has emerged. Additionally, there was increased functional connectivity between the left aSTG and the left pHC following overnight consolidation. This result seems at odds with the CLS model and previous functional connectivity work predicting decreased hippocampal – neocortical functional connectivity following overnight consolidation (McClelland et al., 1995; Takashima et al., 2009). Given the suggested role of the left pHC in retrieval processes (Lepage et al., 1998; Schacter & Wagner, 1999; Spaniol et al., 2009) this finding may reflect more efficient retrieval of high type frequency regulars following overnight consolidation. Overall, these results point to differential roles for the left aHC and left pHC in the learning and consolidation of novel inflectional morphology. Future work capitalising on advances in high-resolution anatomical and functional scanning and subfield segmentation of the hippocampus will likely open up new opportunities to examine how hippocampal subregions may contribute to consolidation processes.

The seed-to-voxel PPI analysis showed that for the two-way interaction between affix training day and affix regularity, the left aSTG was functionally connected to the left MFG. The interaction reflected a profile that was opposite to that observed in the RSA analysis. Where there was an overnight increase in similarity structure among Day 8 regulars in the left aSTG, there was an overnight decrease in functional connectivity between the left aSTG and the left MFG. Conversely, where there was an overnight decrease in similarity structure among Day 9 irregulars, there was an overnight increase in functional connectivity between the left aSTG and the left MFG. The clusters of voxels in the left MFG were situated in dorsolateral prefrontal cortex (DLPFC), encompassing BA9 and extending inferiorly into BA46. The mPFC has more typically been implicated in memory consolidation and schema formation (van Kesteren et al., 2010a, 2010b, Preston & Eichenbaum, 2013). The DLPFC is thought to subserve inhibitory control processes (Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004), which are engaged in memory retrieval suppression (Anderson, Bunce, & Barbas, 2016). The present results may be interpreted as

illustrating a shift in the types of representations being inhibited or suppressed in the retrieval of the newly-learned plurals. High token frequency irregulars may initially be prioritised as indexed by increased similarity structure. As a result, the retrieval of high type frequency regulars may be inhibited or suppressed as indexed by increased functional connectivity between the aSTG and DLPFC. Overnight consolidation may support the development of similarity structure among high type frequency regulars and enhance item-specific representations for salience-tagged high token frequency irregulars leading to the opposite pattern of similarity structure and functional connectivity with the DLPFC. Chapter 10 elaborates on this proposal.

8.4.3. Summary

In sum, RSA and PPI analyses have shown overnight changes in the representations and functional networks underlying newly-learned plurals. The clearest picture emerged from high type frequency regulars, where results were consistent with overnight consolidation supporting the development of an affix representation, enhancing functional connectivity between putative language representational areas, and reducing functional connectivity to inhibitory prefrontal areas. Together, these results point to a more optimal representation of high type frequency regulars following overnight consolidation. Results were less clear for high token frequency irregulars but a tentative suggestion is that overnight consolidation may act to strengthen whole-word representations. Overnight changes in similarity structure and functional connectivity were not mirrored by overnight changes in generalisation behaviour. Experiment 3, reported in the next chapter, explores one possible reason for the differences in generalisation behaviour in Experiments 1 and 2.

Chapter 9: Experiment 3

9.1. Introduction

Experiment 2 revealed overnight changes in the neural representations and functional networks underlying the newly-learned plurals. These changes were argued to be consistent with the development of an affix representation for high type frequency regulars and, more tentatively, suggested to reflect a strengthening of whole-word representations for high token frequency irregulars. Participants performed better on Day 8 compared to Day 9 plurals on test trials presented in the scanner, which assessed affix and whole-word knowledge. This result was argued to provide some indication that the neural representations and functional networks underlying Day 8 may reflect more ‘optimal’ representations (Stickgold & Walker, 2013).

However, Experiment 2 did not replicate the overnight changes in generalisation observed in Experiment 1. Recall that in Experiment 1 and in the work of Mirković and Gaskell (in prep.), participants produced more irregular generalisations (ambiguous) (e.g. *varbimm*) than regular generalisations (e.g. *varbaff*) for novel phonologically ambiguous singulars (e.g. *varbi*) after a period of overnight consolidation. These phonologically ambiguous singulars contained a phonological cue (e.g. *arb*) that had been associated with both a high type frequency regular affix (e.g. *harbaff*) and a high token frequency irregular affix (e.g. *tarbimm*) during training. Some of the predictions regarding the similarity structure and functional networks underlying the newly-learned plurals were derived from these generalisation results. As noted in Chapter 8, one cognitive marker of morphological knowledge is the ability to generalise to novel forms (Bybee & Slobin, 1982; Bybee & Moder, 1983; Prasada & Pinker, 1993; Tamminen et al., 2015). Accordingly, the work presented in this thesis operates on the assumption that generalisation behaviour can be used to derive predictions about the neural representations underlying morphological knowledge. Thus, the aim of Experiment 3 was to examine a potential factor that could have contributed to differences in generalisation across Experiments 1 and 2.

To examine differences in generalisation across Experiments 1 and 2 more directly, an additional analysis of the plural elicitation generalisation task was

performed adding Experiment (1, 2) as an additional predictor. To this end, generalisations to novel phonologically consistent (regular) singulars were excluded from Experiment 2 (Figure 40). A loglinear model with Experiment (1, 2), phonological consistency (diverse, consistent, ambiguous), affix training day (Day 8, Day 9), response type (regular generalisations, irregular generalisations (consistent), irregular generalisations (ambiguous), incorrect), and their interaction was then fitted. The four-way interaction between Experiment, phonological consistency, affix training day, and response type was significant, $\chi^2_{(6)} = 18.49$, $p = .005$. To break down this interaction, interactions between Experiment, affix training day, and response type were examined for each level of phonological consistency separately. For phonologically diverse novel singulars, the three-way interaction was not significant. For phonologically consistent novel singulars, the three-way interaction between Experiment, affix training day, and response type was significant, $\chi^2_{(3)} = 14.59$, $p = .002$. To further break down this interaction, interactions between Experiment and affix training day were examined for each level of response type separately. For regular generalisations, the interaction between Experiment and affix training day was significant, $\chi^2_{(1)} = 9.75$, $p = .002$. Participants produced more regular generalisations for Day 8 affixes in Experiment 2 compared to Experiment 1 ($b = 0.83$, $z = 2.73$, $p = .006$) but a similar number of regular generalisations for Day 9 affixes across both experiments. For irregular generalisations (consistent) and irregular generalisations (ambiguous), the two-way interactions between Experiment and affix training day were not significant. For incorrect responses, the two-way interaction between Experiment and affix training day was significant, $\chi^2_{(1)} = 4.54$, $p = .033$. Participants produced more incorrect responses for Day 8 affixes in Experiment 1 compared to Experiment 2 ($b = 0.37$, $z = 2.03$, $p = .043^+$) but a similar number of incorrect responses for Day 9 affixes across both experiments. For phonologically ambiguous novel singulars, the three-way interaction between Experiment, affix training day, and response type was significant, $\chi^2_{(3)} = 14.59$, $p = .002$. For regular generalisations, the interaction between Experiment and affix training day was significant, $\chi^2_{(1)} = 7.56$, $p = .006$. Participants produced more regular generalisations for Day 8 affixes in Experiment 2 compared to Experiment 1 ($b = 0.44$, $z = 2.36$, $p = .018$) but a similar number of regular generalisations for Day 9 affixes across both experiments. For

irregular generalisations (ambiguous), the interaction between Experiment and affix training day was significant, $\chi^2_{(1)} = 7.07$, $p = .008$. Participants produced more irregular generalisations (ambiguous) for Day 9 affixes in Experiment 2 compared to Experiment 1 ($b = 0.48$, $z = 2.30$, $p = .022$) but a similar number of irregular generalisations (ambiguous) for Day 8 affixes across both experiments. For irregular generalisations (consistent) and incorrect responses, the two-way interactions between Experiment and affix training day were not significant. These results are summarised in Table 1.

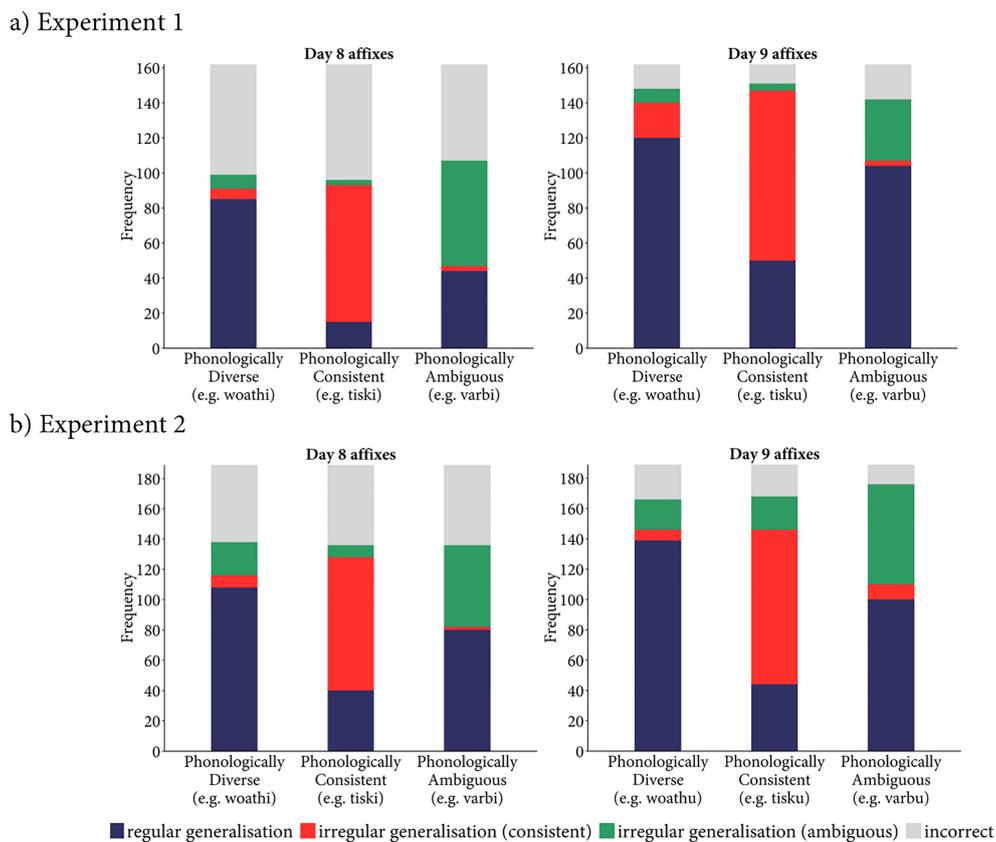


Figure 40. Experiment 3: generalisation results Experiments 1 vs. 2

Frequency of each response type supplied within each phonological consistency category of novel singulars requiring participants to use affixes trained on Day 8 (left) and Day 9 (right) for a) Experiment 1 and b) Experiment 2. To allow direct comparisons between the two experiments, phonologically consistent (regular) novel singulars were excluded from Experiment 2.

To ensure that differences in generalisation across Experiments 1 and 2 were not attributable to differences in the learning of the novel morphological system across the two experiments, accuracy on the picture naming and 2AFC recognition memory tasks presented in the testing blocks at the end of each training session were reanalysed including Experiment (1, 2) as an additional factor (refer to Chapter 5

Effect	p-value
Four-way interaction	
<i>Experiment x phonological consistency x affix training day x response type</i>	.005
Three-way interactions (at each level of phonological consistency)	
1. Phonologically Diverse	
<i>Experiment x affix training day x response type</i>	ns.
2. Phonologically Consistent	
<i>Experiment x affix training day x response type</i>	.002
Two-way interactions (at each level of response type)	
1. Regular generalisations	
<i>Experiment x affix training day</i>	.002
<i>Exp. 2 Day 8 > Exp. 1 Day 8</i>	.006
Exp. 2 Day 9 = Exp. 1 Day 9	ns.
2. Irregular generalisations (consistent)	
<i>Experiment x affix training day</i>	ns.
3. Irregular generalisations (ambiguous)	
<i>Experiment x affix training day</i>	ns.
4. Incorrect responses	
<i>Experiment x affix training day</i>	.033
<i>Exp. 2 Day 8 < Exp. 1 Day 8</i>	.043
Exp. 2 Day 9 = Exp. 1 Day 9	ns.
3. Phonologically Ambiguous	
<i>Experiment x affix training day x response type</i>	.002
Two-way interactions (at each level of response type)	
1. Regular generalisations	
<i>Experiment x affix training day</i>	.002
<i>Exp. 2 Day 8 > Exp. 1 Day 8</i>	.018
Exp. 2 Day 9 = Exp. 1 Day 9	ns.
2. Irregular generalisations (consistent)	
<i>Experiment x affix training day</i>	ns.
3. Irregular generalisations (ambiguous)	
<i>Experiment x affix training day</i>	.008
4. Incorrect responses	
<i>Experiment x affix training day</i>	ns.

Table 1. Experiments 1 vs. 2 analysis summary

Summary of the analysis comparing the plural elicitation generalisation task results in Experiments 1 and 2. The effects tested are listed on the left and p-values reported on the right. Significant effects are in red.

section 5.3.2 and Chapter 8 section 8.3.1.2 for the initial analyses of Experiment 1 and 2, respectively). Phonologically consistent (regular) forms were excluded from Experiment 2 so that both experiments could be included in the same analysis. For singulars, logistic regression models with Experiment (1, 2), condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8, Day 9), and their interaction were fitted to the picture naming and 2AFC recognition memory data separately. There were no main effects or interactions involving Experiment on either task.

For plurals, logistic regression models with Experiment (1, 2), condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction were fitted to the picture naming and 2AFC recognition memory data separately. There were no main effects or interactions involving Experiment on either task. Thus, differences in generalisation between Experiments 1 and 2 could not be attributed to differential learning outcomes. It should also be noted that the two experiments had similar sample sizes (Experiment 1: $n = 18$; Experiment 2: $n = 22$) and participants of similar ages (Experiment 1: mean age = 23, $SD = 4$; Experiment 2: mean age = 23, $SD = 3$) recruited from the same volunteer panel following the same procedures. It is thus unlikely than differences in generalisation reflected sample differences across Experiments 1 and 2.

Together, these observations suggest that differences in generalisation across the two experiments may be attributable to differences in the training stimuli set. Recall that in Experiment 2, 3 phonologically varied stems from Experiment 1 were replaced with stems containing a new phonological cue (e.g. *-olt*; *zolt*, *golt*, *tol*), which was consistently associated with a regular affix (e.g. *zoltaff*, *goltaff*, *tolaff*) analogous to phonologically consistent regular English verbs (e.g. *walked*, *talked*, *stalked*). This change introduced a set of highly phonologically consistent regulars within a set of more phonologically diverse regulars. As discussed in Chapters 1 and 3, phonological consistency effects in the processing and generalisation of English regular past tenses have been reported (Seidenberg & Bruck, 1990; Seidenberg, 1992; Daugherty & Seidenberg, 1992; Albright & Hayes, 2003). It is thus plausible that in

Experiment 2, phonological consistency among a subset of regulars being supported by the overall high type frequency of the regular pattern may have biased participants towards generalising the regular inflectional pattern to a greater extent than in Experiment 1. The analysis comparing generalisation across Experiments 1 and 2 indeed showed that for both phonologically consistent and phonologically ambiguous novel singulars, there were more regular generalisations following overnight consolidation in Experiment 2 compared to Experiment 1. Furthermore, whilst participants produced more irregular generalisations (ambiguous) for phonologically ambiguous singulars for Day 9 affixes in Experiment 2 compared to Experiment 1, there was no increase in irregular generalisations (ambiguous) following overnight consolidation. Together, these results suggest that the changes made to the training stimuli set in Experiment 2 may have supported the preferential consolidation of high type frequency regulars over high token frequency irregulars, at least to a greater extent than in Experiment 1. The RSA analysis showing overnight changes in the similarity structure underlying the neural representations of high type frequency regulars provides some support for this interpretation.

The introduction of phonologically consistent regulars in Experiment 2 may also be argued to have shifted the conditional probabilities (G. A. Miller & Selfridge, 1950) of the training stimuli set. In each training stimuli set, let the probability of a regular or irregular plural affix (B) occurring given a phonologically structured stem (A) be defined as

$$p(B|A) = p(A \cap B)/p(A)$$

In Experiment 1, the conditional probability of a regular affix occurring given a phonologically structured stem (e.g. *har**ba**ff*, *yar**ba**ff*, *nar**ba**ff*) was 0.33 defined as

$$p(\text{regular affix}|\text{phonologically structured stem}) = p(\text{phonologically structured stem} \cap \text{regular affix})/p(\text{phonologically structured stem})$$

or

$$0.33 = (3/18)/(9/18)$$

The conditional probability of an irregular affix occurring given a phonologically structured stem (e.g. *hiskesh*, *liskesh*, *fiskesh*, *tarbimm*, *clarbimm*, *slarbimm*) was 0.67 defined as

$$p(\text{irregular affix}|\text{phonologically structured stem}) = \frac{p(\text{phonologically structured stem} \cap \text{irregular affix})}{p(\text{phonologically structured stem})}$$

or

$$0.67 = (6/18)/(9/18)$$

In Experiment 2, the conditional probability of a regular affix occurring given a phonologically structured stem (e.g. *zoltaff*, *groltaff*, *toltaff*, *harbaff*, *yarbaff*, *narbaff*) was 0.50 defined as

$$p(\text{regular affix}|\text{phonologically structured stem}) = \frac{p(\text{phonologically structured stem} \cap \text{regular affix})}{p(\text{phonologically structured stem})}$$

or

$$0.50 = (6/18)/(12/18)$$

The conditional probability of an irregular affix occurring given a phonologically structured stem (e.g. *hiskesh*, *liskesh*, *fiskesh*, *tarbimm*, *clarbimm*, *slarbimm*) was 0.50 defined as

$$p(\text{irregular affix}|\text{phonologically structured stem}) = \frac{p(\text{phonologically structured stem} \cap \text{irregular affix})}{p(\text{phonologically structured stem})}$$

or

$$0.50 = (6/18)/(12/18)$$

The shift in conditional probabilities in the training stimuli set across Experiments 1 and 2 may have biased the consolidation and generalisation towards high type frequency regulars. The aim of Experiment 3 was to manipulate the phonological consistency and type and token frequency of the plurals during training

to bias the consolidation and generalisation of the plurals towards high token frequency irregulars. It was hypothesised that such manipulations, described in detail in the methods section 9.2.2.1, would increase irregular generalisations following overnight consolidation thereby replicating the results of Experiment 1 and Mirković and Gaskell (in prep.).

9.2. Materials and methods

9.2.1. Participants

Twenty-three participants (7 males) aged between 20 and 35 (mean age = 25, SD = 4) were recruited from the MRC Cognition and Brain Sciences Unit Participant Panel and provided their informed consent to take part in the study. Participants were tested under the approval of the Cambridge Psychology Research Ethics Committee. All were monolingual speakers of British English with little or no knowledge of a second language, no known hearing or language impairments, and no neurological or psychiatric disorders. Participants were paid to take part in the study.

9.2.2. Experimental stimuli

9.2.2.1. Training stimuli

The training stimuli described in Experiment 1 were used with some minor changes. Removing the 3 phonologically consistent (regular) stems (e.g. *zoltaff*, *groltaff*, *toltaff*) introduced in Experiment 2 allowed re-establishing the same conditional probabilities as in Experiment 1. Unlike Experiments 1 and 2, the phonologically ambiguous and phonologically consistent irregulars were associated with the same irregular affix (e.g. *tarbesh*, *hiskesh*) analogous to irregular English verbs having phonologically distinct stems but following the same type of transformation to form their past tenses (e.g. *draw* – *drew*, *blow* – *blew*; *hang* – *hung*, *cling* – *clung*) (Figure 41; refer to Chapter 5 Figure 7 and Chapter 8 Figure 19 for the training stimuli for Experiment 1 and 2, respectively). This modification was made to increase the type frequency of the irregular inflectional pattern. Re-establishing the conditional probabilities of Experiment 1 and increasing the type frequency of the irregular inflectional pattern were expected to bias the consolidation and generalisation towards the irregulars. Increasing the type frequency of the irregular inflectional pattern required reducing the token frequency of the irregulars to 12 presentations across training in order to

keep the affix frequency the same across conditions. The token frequency of the irregulars was nonetheless twice that of the regulars. Eight training lists were created to counterbalance the assignment of plural affixes to regular and irregular conditions and the assignment of phonological cues to consistent and ambiguous conditions. The assignment of each set of plural affixes (*-aff*, *-esh*; *-opp*, *-ull*) to each gender was counterbalanced across participants, as was the order of gender training across Day 8 and Day 9. As in Experiments 1 and 2, the pairing of nouns with specific occupations was varied across different training lists to further control for any unintended associations between the novel nouns and pictures.

a)

Phonological consistency	Type/Token frequency	Training set 1		Training set 2	
		Singular nouns	Plural nouns	Singular nouns	Plural nouns
Diverse	High type/ Low token (regular)	gleeti	gleetaff	shilnu	shilnopp
		torthi	torthaff	plassu	plassopp
		dulti	dultaff	blornu	blornopp
		vonti	vontaff	hulfu	hulfopp
		zolli	zollaff	rutchu	rutchopp
		mowli	mowlaff	grollu	grollopp
		sleni	slenaff	crephu	crehopp
		chiffi	chiffaff	thiltu	thiltopp
		pelfi	pelfaff	dowthu	dowthopp
		Ambiguous	High type/ Low token (regular)	harbi	harbaff
yarbi	yarbaff			larbu	larbopp
narbi	narbaff			glarbu	glarbopp
High token/ Low type (irregular)	tarbi		tarbesh	blarbu	blarbull
	clarbi		clarbesh	marbu	marbull
	slarbi		slarbesh	farbu	farbull
Consistent	High token/ Low type (irregular)	hiski	hiskesh	tisku	tiskull
		liski	liskesh	visku	viskull
		fiski	fiskesh	visku	viskull

b)

Type frequency	Token frequency	Affix frequency
12	6	72
6	12	72

c)

Figure 41. Experiment 3: training stimuli set

a) example training stimuli set. Phonologically diverse and phonologically ambiguous plurals taking a regular affix are in blue, phonologically consistent and phonologically ambiguous plurals taking an irregular affix are in red. b) type, token, and affix (type x token) training frequencies associated with each plural affix. c) example singular and plural nouns associated with familiar occupations for female and male characters.

9.2.2.3. Generalisation stimuli

The generalisation stimuli were the same as for Experiment 1. Nine were phonologically diverse (e.g. *woath*), 9 were phonologically consistent (e.g. *zisk*), and 9 were phonologically ambiguous (e.g. *slarb*). Each novel stem was combined with both singular affixes (e.g. *woathi*_[fem:sing], *woathu*_[masc:sing]) yielding a total of 54 generalisation items. Each new singular noun was paired with the picture of a single female and male character depicting a familiar occupation not used in the training stimuli sets (e.g. *woathi*_[fem:sing] = *teacher*_[fem:sing], *woathu*_[masc:sing] = *teacher*_[masc:sing]).

9.2.3. Experimental design and procedure

The training schedule was the same as for Experiments 1 and 2. On Day 1, participants were trained on the novel singulars in one session lasting approximately 1hr30. On Day 8, participants were trained on the first set of plural affixes in one session lasting approximately 1hr15. On Day 9, approximately 24 hours after learning the first set of plurals, participants were trained on the second set of plurals before completing the plural elicitation and 3AFC generalisation tasks.

9.3. Results

9.3.1. Training

Accuracy on the picture naming tasks presented at the end of each training block in each training session was analysed to examine how participants' mastery of the new words improved over the course of training (Figure 42). Singulars and plurals were analysed separately as there was no plural training on Day 1, there were only 3 training blocks for plurals on Day 8 and Day 9 (compared to 4 blocks for singulars), and because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix.

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8, Day 9), training block (Block 1, Block 2, Block 3, Block 4), and their interaction was fitted. The three-way interaction between condition, training day, and training block was not significant. The two-way interaction between training day and training block was significant, $\chi^2_{(6)} = 77.03$, $p < .001$. Participants improved significantly across training blocks on Day 1 (Block 2 vs. Block 1: $b = 1.77$, $z = 19.03$, $p < .001$, Block 3 vs. Block 2: $b = 0.49$, $z = 6.94$, $p < .001$, Block 4 vs. Block 3: $b = 0.44$, $z = 6.14$, $p < .001$). On Day 8, participants improved from the second to the third block (Block 3 vs. Block 2: $b = 0.58$, $z = 2.36$, $p = .019^+$) and from the third to the fourth block (Block 4 vs. Block 3: $b = 0.93$, $z = 2.64$, $p = .008^+$). On Day 9, participants improved from the first to the second block (Block 2 vs. Block 1: $b = 0.53$, $z = 2.52$, $p = .012^+$) and from the second to the third block (Block 3 vs. Block 2: $b = 0.55$, $z = 2.07$, $p = .038^+$). The two-way interactions between condition and training block and between condition and training day were not significant. The main effect of training block was significant,

$\chi^2_{(3)} = 1222.50$ $p < .001$. Participants became increasingly more accurate over the course of the training blocks (Block 2 vs. Block 1, $b = 1.37$, $z = 18.42$, $p < .001$, Block 3 vs. Block 2, $b = 0.52$, $z = 7.81$, $p < .001$, Block 4 vs. Block 3: $b = 0.46$, $z = 6.71$, $p < .001$). The main effect of training day was significant, $\chi^2_{(2)} = 2668.20$, $p < .001$. Participants were more accurate on Day 9 compared to Day 1 ($b = 2.99$, $z = 30.95$, $p < .001$), on Day 8 compared to Day 1 ($b = 2.97$, $z = 30.96$, $p < .001$) but performed similarly on Day 8 and Day 9.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), training block (Block 2, Block 3, Block 4), and their interaction was fitted. The three-way interaction between phonological consistency, training day, and training block was not significant. The two-way interaction between condition and training day was significant, $\chi^2_{(3)} = 12.25$, $p = .007$. Participants were more accurate on phonologically ambiguous (regular) plurals on Day 9 compared to Day 8 ($b = 0.64$, $z = 2.60$, $p = .009$). The two-way interactions between condition and training block and between condition and training day were not significant. The main effect of training block was significant, $\chi^2_{(2)} = 69.88$ $p < .001$. Participants became increasingly more accurate over the course of the training blocks (Block 3 vs. Block 2, $b = 0.77$, $z = 5.83$, $p < .001$, Block 4 vs. Block 3, $b = 0.31$, $z = 2.01$, $p = .044^\dagger$). The main effect of condition was significant, $\chi^2_{(3)} = 37.74$, $p < .001$. Participants were more accurate on phonologically consistent than on phonologically ambiguous (regular) ($b = 0.82$, $z = 4.12$, $p < .001$) and phonologically ambiguous (irregular) ($b = 0.86$, $z = 4.31$, $p < .001$) plurals. Participants were also more accurate on phonologically diverse than on phonologically ambiguous (regular) ($b = 0.63$, $z = 4.25$, $p < .001$) and phonologically ambiguous (irregular) ($b = 0.67$, $z = 4.51$, $p < .001$) plurals. The main effect of training day was not significant.

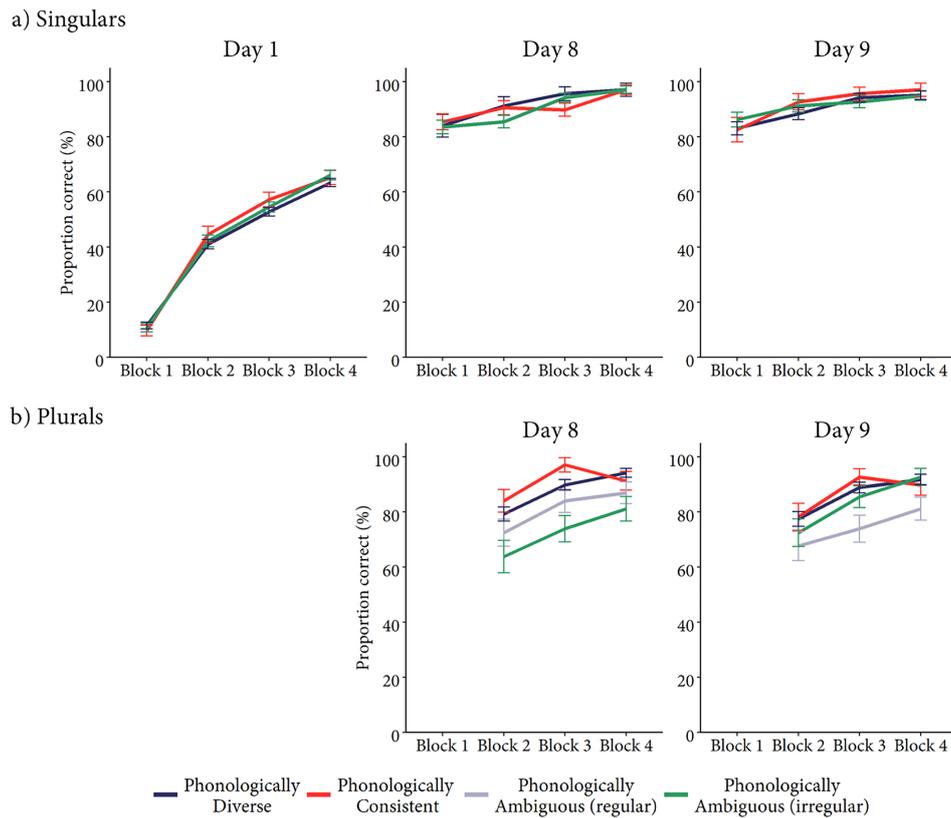


Figure 42. Experiment 3: training accuracy

Mean accuracy on the picture naming task presented at the end of each training block (Block 1, Block 2, Block 3, Block 4) on each training day (Day 1, Day 8, and Day 9) for a) singulars and b) plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014). Phonologically ambiguous plurals are split into those associated with a regular (light blue line) and those associated with an irregular (green line) affix. All phonologically ambiguous singulars are depicted by a green line.

9.3.2. Testing

Accuracy on the picture naming (Figure 43) and 2 AFC recognition memory (Figure 44) presented in a testing block at the end of each training session was analysed to ensure that words with different phonological consistencies and training frequencies (plurals) were learned equally well. Singulars and plurals were analysed separately as there was no plural training on Day 1, there were only 3 training blocks for plurals on Day 8 and Day 9 (compared to 4 blocks for singulars), and because phonologically ambiguous plurals were split according to whether they were associated with a regular or an irregular plural affix.

9.3.2.1. Picture naming

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8, Day 9), and their interaction was fitted. The interaction between condition and

training day was not significant. The main effect of condition was significant, $\chi^2_{(2)} = 8.00$, $p = .018$. Participants were more accurate on phonologically consistent ($b = 0.40$, $z = 2.34$, $p = .020^+$) and phonologically ambiguous ($b = 0.28$, $z = 2.14$, $p = .032^+$) than on phonologically diverse singulars. The main effect of training day was significant, $\chi^2_{(2)} = 170.35$, $p < .001$. Participants were more accurate on Day 8 compared to Day 1 ($b = 2.08$, $z = 7.21$, $p < .001$) and on Day 9 compared to Day 1 ($b = 2.12$, $z = 7.37$, $p < .001$) but performed similarly on Day 8 and Day 9.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction was fitted. The interaction between condition and training day was not significant. Nor the main effect of condition or the main effect of training day was significant.

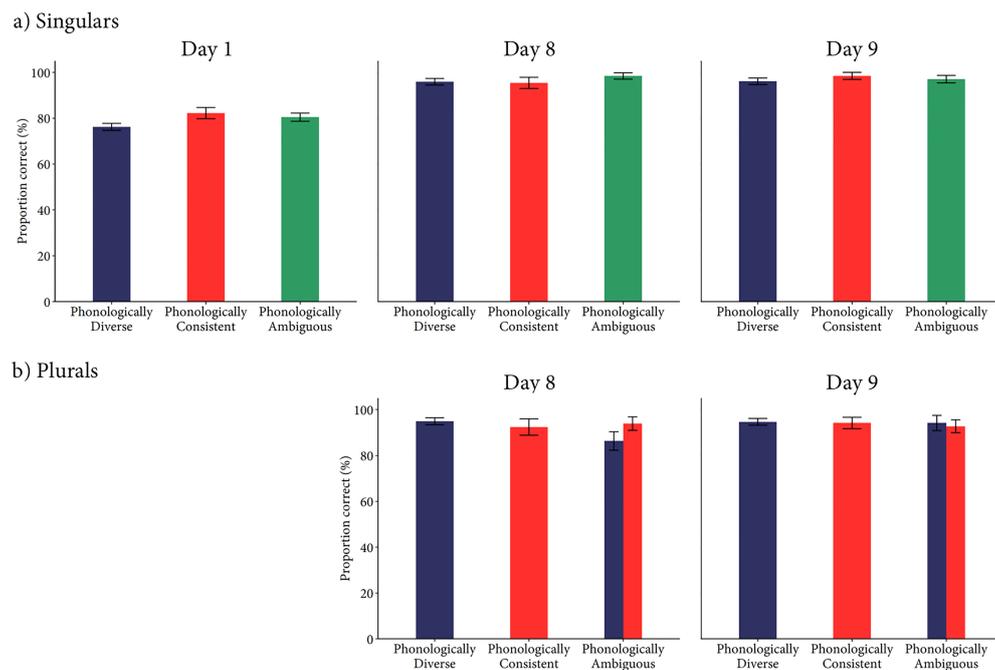


Figure 43. Experiment 3: test accuracy (picture naming)

Mean accuracy on the picture naming task presented in the testing block the end of each training day (Day 1, Day 8, Day 9) for a) singulars and b) plurals. Error bars represent ± 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (red bar) affix.

9.3.2.2. 2AFC recognition memory

For singulars, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous), training day (Day 1, Day 8, Day 9), and their interaction was fitted. The interaction between condition and training day was not significant. The main effect of training day was significant, $\chi^2_{(2)} = 8.41$, $p = .015$. Participants were more accurate on Day 9 compared to Day 1 ($b = 1.13$, $z = 2.41$, $p = .016$) but performed similarly on Day 8 and Day 1 and on Day 8 and Day 9. The main effect of condition was not significant.

For plurals, a logistic regression with condition (phonologically diverse, phonologically consistent, phonologically ambiguous (regular), phonologically ambiguous (irregular)), training day (Day 8, Day 9), and their interaction was fitted. The interaction between condition and training day was not significant. Nor the main effect of condition or training day was significant.

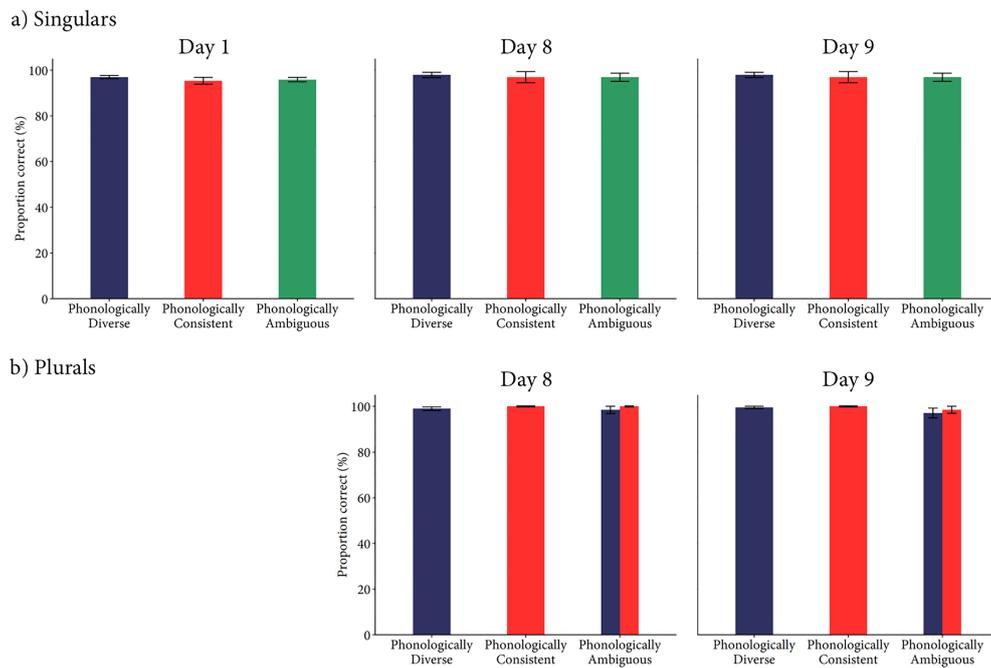


Figure 44. Experiment 3: test accuracy (2AFC)

Mean accuracy on the picture naming task presented in the testing block the end of each training day (Day 1, Day 8, Day 9) for a) singulars and b) plurals. Error bars represent +/- 1 standard error of the mean after removing between-subject variability (O'Brien & Cousineau, 2014).

Phonologically ambiguous plurals are split into those associated with a regular (blue bar) and those associated with an irregular (red bar) affix.

9.3.3. Generalisation

9.3.3.1. Plural elicitation

9.3.3.1.1. Response type analysis

The frequency of each response type (regular generalisations, irregular generalisations, incorrect) for novel phonologically diverse, phonologically consistent, and phonologically ambiguous novel singulars requiring participants to use plural affixes trained on Day 8 or Day 9 are shown in Figure 45. A loglinear model with phonological consistency (diverse, consistent, ambiguous), affix training day (Day 8, Day 9), response type (regular generalisations, irregular generalisations, incorrect), and their interaction was fitted. The three-way interaction between phonological consistency, affix training day, and response type was not significant. The two-way interaction between phonological consistency and response type was significant, $\chi^2_{(4)} = 95.28$, $p < .001$. For novel phonologically diverse singulars, participants produced more regular than irregular generalisations ($b = 0.61$, $z = 5.33$, $p < .001$), and more correct responses than incorrect responses (regular generalisations vs. incorrect: $b = 1.14$, $z = 8.20$, $p < .001$; irregular generalisations vs. incorrect: $b = 0.52$, $z = 3.44$, $p < .001$). For novel phonologically consistent singulars, participants produced more irregular than regular generalisations ($b = 0.91$, $z = 7.41$, $p < .001$) and incorrect responses ($b = 1.17$, $z = 8.64$, $p < .001$) but a similar number of regular generalisations and incorrect responses. For novel phonologically ambiguous singulars, participants produced a similar number of regular and irregular generalisations and produced more correct than incorrect responses (regular generalisations vs. incorrect: $b = 0.89$, $z = 6.41$, $p < .001$; irregular generalisations vs. incorrect: $b = 0.68$, $z = 4.78$, $p < .001$).

The two-way interaction between affix training day and response type was significant, $\chi^2_{(2)} = 99.41$, $p < .001$. Participants produced a similar number of regular generalisations for Day 8 and Day 9 affixes but produced more irregular generalisations for Day 9 compared to 8 affixes ($b = 0.50$, $z = 5.40$, $p < .001$) and more incorrect responses for Day 8 compared to Day 9 affixes ($b = 1.24$, $z = 7.53$, $p < .001$).

The two-way interaction between affix training day and phonological consistency was not significant. The main effect of response type was significant, $\chi^2_{(2)} = 141.24$, $p < .001$. Participants produced a similar number of regular and irregular generalisations and produced more correct than incorrect responses (regular generalisations vs.

incorrect: $b = 0.82$, $z = 9.98$, $p < .001$; irregular generalisations vs. incorrect: $b = 0.83$, $z = 10.18$, $p < .001$).

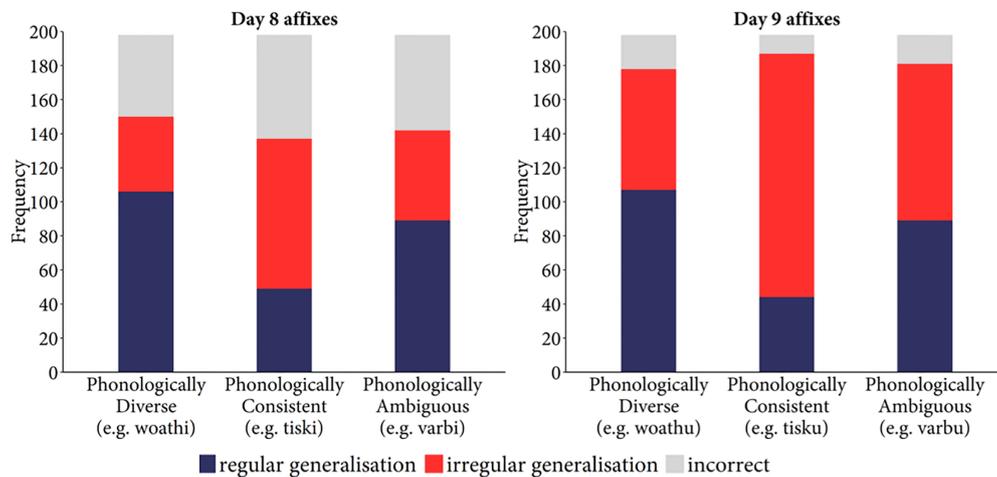


Figure 45. Experiment 3: generalisation results

Frequency of each response type supplied within each phonological consistency category of novel singulars requiring participants to use affixes trained on Day 8 (left) and on Day 9 (right). Regular generalisations are in blue (e.g. *woathaff*) and irregular generalisations are in red (e.g. *varbesh*).

To mitigate potential confounds due to a difference in the number of correct responses for Day 8 compared to Day 9 affixes, the analysis was repeated with a random sub-sample of correct responses for Day 9 affixes to artificially match the number of correct responses produced for Day 8 affixes (Figure 46). The three-way interaction between phonological consistency, affix training day, and response type was not significant. The two-way interaction between phonological consistency and response type remained significant, $\chi^2_{(2)} = 70.40$, $p < .001$. For novel phonologically diverse singulars, participants produced more regular than irregular generalisations ($b = 0.55$, $z = 4.63$, $p < .001$). For novel phonologically consistent singulars, participants produced more irregular than regular generalisations ($b = 0.82$, $z = 6.25$, $p < .001$). For novel phonologically ambiguous singulars, sub-sampling the data revealed significant more regular generalisations than irregular generalisations ($b = 0.24$, $z = 2.03$, $p = .043$). The two-way interaction between affix training day and response type remained significant, $\chi^2_{(1)} = 15.35$, $p < .001$. Sub-sampling the data revealed that participants produced more regular generalisations for Day 8 compared to Day 9 affixes ($b = 0.25$, $z = 2.59$, $p = .010$) and more irregular generalisations for Day 9 compared to Day 8 affixes ($b = 0.26$, $z = 2.62$, $p = .009$). None of the main effects were significant.

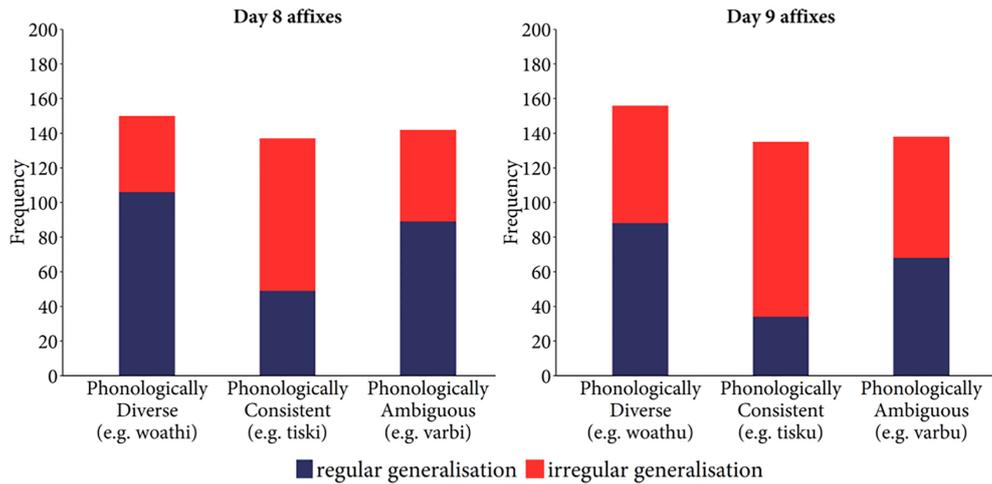


Figure 46. Experiment 3: generalisation results for equalised correct responses

Frequency of each response type supplied within each phonological consistency category of novel words requiring participants to use affixes of Day 8 (left) and on Day 9 (right).

9.3.3.1.1. Error analysis

A secondary analysis focused on the errors produced by participants. Incorrect responses were subcategorised as a failure to provide any response (no response), the production of a legal but gender-inappropriate affix (gender error), and the production of an illegal affix (Figure 47). A loglinear model with phonological consistency (diverse, consistent, ambiguous), affix training day (Day 8, Day 9), error type (no response, gender error, illegal affix), and their interaction was fitted. There was no significant three-way interaction between phonological consistency, affix training day, and error type. The two-way interaction between affix training day and error type was significant, $\chi^2_{(2)} = 6.80$, $p = .033$. To break down this interaction, differences between the numbers of each error type were examined for Day 8 and Day 9 affixes separately. For Day 8 affixes, the main effect of error type was not significant. For Day 9 affixes, the main effect of error type was significant, $\chi^2_{(2)} = 12.18$, $p = .002$. Participants produced more illegal affix errors than no response ($b = 0.83$, $z = 2.20$, $p = .028^+$) and gender errors ($b = 1.12$, $z = 3.38$, $p < .001$) but a similar number of gender errors and no response errors. The two-way interactions between affix training day and phonological consistency and between phonological consistency and error type were not significant. The main effect of affix training day was significant, $\chi^2_{(1)} = 67.97$, $p < .001$. Participants produced more incorrect responses for Day 8 compared to Day 9 affixes ($b = 1.24$, $z = 7.53$, $p < .001$). The main effect of error type was significant, $\chi^2_{(2)} = 9.41$, $p = .009$. Participants produced more illegal affix errors than no response

($b = 0.56$, $z = 2.88$, $p = .004$) and gender errors ($b = 0.36$, $z = 2.36$, $p = .018^+$) but a similar number of gender errors and no response errors. The main effect of phonological consistency was not significant.

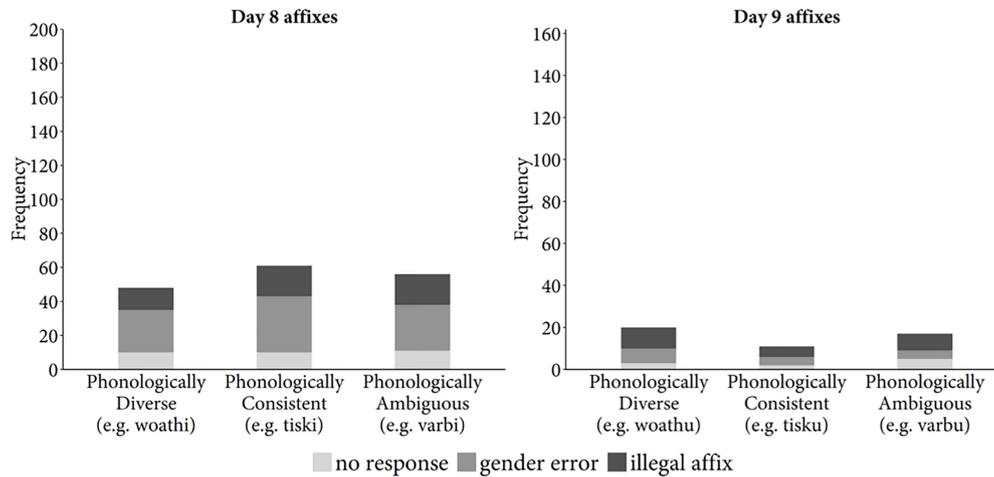


Figure 47. Experiment 3: generalisation errors

Frequency of each error type (no response, gender error, illegal affix) supplied within each phonological consistency category of novel singulars requiring participants to use affixes trained on Day 8 (left) or on Day 9 (right).

9.3.3.1. 3AFC

The frequency count of each response type (regular generalisations, irregular generalisations) for novel phonologically diverse, consistent, and ambiguous singulars requiring participants to select affixes trained on Day 8 or Day 9 are shown in Figure 48. A loglinear model with phonological consistency (diverse, consistent, ambiguous), affix training day (Day 8, Day 9), response type (regular generalisation, irregular generalisation), and their interaction was fitted. The three-way interaction between phonological consistency, affix training day, and response type was not significant. The two-way interaction between phonological consistency and response type was significant, $\chi^2_{(2)} = 125.33$, $p < .001$. To break down this interaction, the numbers of each response type were compared for each phonological consistency separately. For phonological diverse novel singulars, participants selected more regular than irregular generalisations ($b = 0.77$, $z = 7.16$, $p < .001$). For phonological consistent novel singulars, participants selected more irregular than regular generalisations ($b = 0.86$, $z = 7.80$, $p < .001$). For phonologically ambiguous novel singulars, participants selected more regular than irregular generalisations ($b = 0.20$, $z = 2.01$, $p = .045^+$). The main effect of response type was not significant.

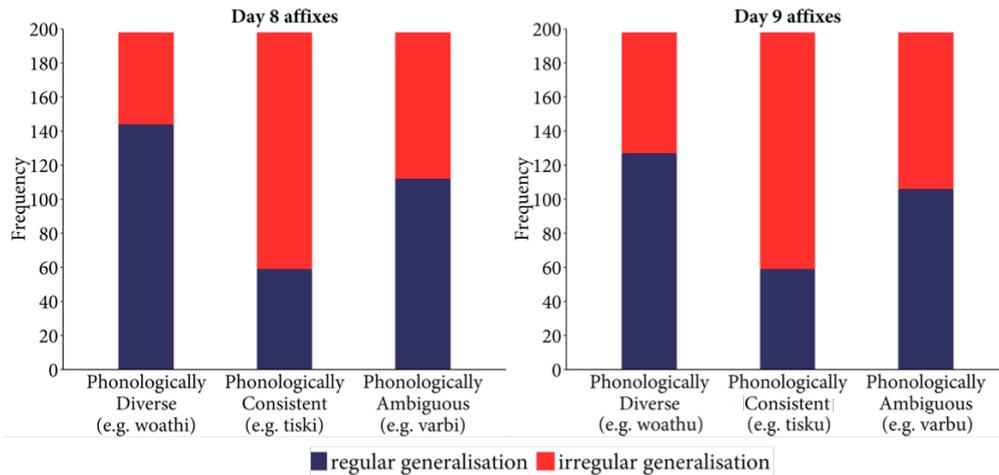


Figure 48. Experiment 3: generalisation 3AFC results

Frequency of each response type supplied within each phonological consistency category of novel singulars requiring participants to use affixes train on Day 8 (left) and on Day 9 (right).

9.4. Discussion

Experiment 3 sought to address the discrepancy in generalisation results between Experiments 1 and 2. Recall that in both Experiment 1 and the work of Mirković and Gaskell (in prep.), participants produced more irregular generalisations (ambiguous) (e.g. *varbimm*) than regular generalisations (e.g. *varbaff*) for novel phonologically ambiguous singulars (e.g. *varbi*) after a period of overnight consolidation. However, this particular overnight change in generalisation was not replicated in Experiment 2. It was hypothesised that such differences in generalisation results may be attributable to changes made to the training stimuli set across the first two experiments. It was argued at the outset of this chapter that these changes may have biased the consolidation and generalisation towards high type frequency regulars. Experiment 3 examined whether modifying the phonological consistency and type and token frequency of the training stimuli could bias consolidation and generalisation towards high token frequency irregulars thereby replicating the pattern of results obtained in Experiment 1 and reported by Mirković and Gaskell (in prep.). Behavioural results from the training sessions largely replicated those obtained in Experiments 1 and 2. Across the picture naming and 2AFC recognition memory task presented at the end of each training session, participants performed similarly for both sets of plurals as evidenced by a lack of significant main effects or interactions involving affix training day. As such, differences in the generalisation of Day 8 and Day 9 plural affixes to novel singulars could not be attributed to differences in initial learning.

The increase in irregular generalisations (ambiguous) following overnight consolidation for novel phonologically ambiguous singulars found in Experiment 1 and reported by Mirković and Gaskell (in prep.) was again not replicated. Instead, participants produced a similar number of regular and irregular generalisations for both Day 8 and Day 9 affixes. However, equalising the number of correct responses across the two training days revealed more regular generalisations. This result was further borne out in the 3AFC generalisation task. These results again suggest that forgetting of Day 8 affixes or interference from Day 9 affixes may have confounded the initial analysis. Indeed, as in Experiments 1 and 2, affix training day had an effect on overall performance on the plural elicitation task as evidenced by participants producing significantly more incorrect responses for Day 8 compared to Day 9 affixes. It is noteworthy that, unlike in Experiments 1 and 2, there was no three-way interaction between affix training day, phonological consistency, and response type on the plural elicitation generalisation task. It thus appears that the modifications made to the training stimuli were ineffective in boosting the consolidation and generalisation of high token frequency irregulars. Increasing the type frequency of the irregulars required lowering their token frequency to ensure affix frequency was kept the same across conditions. In both Experiments 1 and 2, an argument was made for the role of high token frequency in potentially tagging new memories as being particularly relevant to future behaviour leading to their preferential consolidation. In Experiment 1, this was argued to be particularly important in ambiguous cases where different mappings are competing for consolidation. This was evidenced by an increase in irregular generalisations (ambiguous) for novel phonologically ambiguous singulars. There was also evidence of a reduction in regular generalisations for novel phonologically consistent singulars, which was argued to reflect a reduction in inappropriate generalisations supported by a strengthening of the representation of high token frequency consistent irregulars. Whilst the generalisation results in Experiment 2 were less straightforward than those of Experiment 1, there was nonetheless evidence of more optimal generalisations for both phonologically consistent and phonologically ambiguous novel singulars. In both cases, there was an overnight decrease in inappropriate irregular generalisations. This reduction in inappropriate generalisations was argued to reflect an overnight strengthening of the

whole-word representations for high token frequency irregulars. The RSA analysis pointed to an overnight reduction in similarity structure among high token frequency irregulars, which was argued to reflect a strengthening of their whole-word representations. This argument remains tentative and awaits further empirical support. However, whilst it is difficult to argue from a null effect, it is conceivable that a reduction in the token frequency of the irregulars in Experiment 3 explains, at least in part, the lack of overnight changes in generalisation characterised by a non-significant three-way interaction between affix training day, phonological consistency, and response type. Hypothetically, a certain threshold level of token frequency may be required to strengthen whole-word representations.

Alternatively, the lack of consolidation-related changes in generalisation in Experiment 3 might reflect a ceiling effect in the sense that overnight consolidation may not enhance associations that are already strong at encoding. Higher type frequency for the irregulars may have supported stronger encoding by facilitating the emergence of relationships among them. The suggestion that weakly encoded newly-learned associations benefit more from periods of offline consolidation including sleep than more strongly or overlearned associations has been discussed previously (Empson & Clarke, 1970; Ekstrand, 1977; Tilley & Empson, 1978; Cipolli, 1995). Previous work has reported greater consolidation benefits from sleep on the recall of weakly encoded compared to strongly encoded lists of word pairs (Drosopoulos et al., 2006). Stronger encoding has also been suggested to account for the failure of previous study to reveal sleep-related improvements in the retention of material after certain learning tasks (Smith, 2001; Paller & Voss, 2004; Rauchs, Desgranges, Foret, & Eustache, 2005). It is thus conceivable that similar effects may extend here to the development of representations that can support generalisation. While further data will be required to substantiate this proposal, it seems that the distributional statistics of the material to be learned may play a key role in its subsequent consolidation trajectory and generalisation.

In sum, the results obtained in Experiment 3 did not replicate the results obtained in Experiment 1 or those reported by Mirković and Gaskell (in prep.). The aim of the experiment was to examine whether manipulating the phonological

consistency and type and token frequency of the training stimuli would allow boosting the consolidation and generalisation of high token frequency irregulars. Unlike Experiments 1 and 2, the results provided no evidence of consolidation-related changes in generalisation. These findings highlight the fact that overnight consolidation may be highly sensitive to the distributional statistics of novel material to be remembered such that seemingly small changes to the training stimuli can alter measurable overnight changes in generalisation. The next chapter concludes this thesis by summarising its key findings and relating them to consolidation and generalisation theory more broadly.

Chapter 10: Conclusions and future directions

10.1. Thesis summary

The experiments reported in this thesis aimed to characterise the cognitive and neural representations and mechanisms supporting morphological learning and generalisation within a CLS framework. To this end, the work presented was uniquely situated at the intersection of three literatures pertaining to morphological learning, memory consolidation, and generalisation. Experiment 1 examined the extent to which the contribution of overnight consolidation to morphological generalisation varies with the type and token frequency and phonological consistency of newly-learned inflected words. Experiment 2 employed RSA and PPI analyses of fMRI data to characterise the neural representations and functional networks underlying newly-learned inflected words. Experiment 3 addressed differences in generalisation results between the first two experiments.

As reviewed in Chapters 1 and 3, child acquisition data, neural network simulations, diachronic change data, and more limited behavioural experiments have shown type and token frequency as well as phonological consistency to play important roles in morphological learning and generalisation (Seidenberg & Bruck, 1990; Daugherty & Seidenberg, 1992; Plunkett & Marchman, 1991, 1993; Bybee & Newman, 1995; Hare & Elman, 1995; Hare et al., 1995; Ellis & Schmidt, 1998; Albright & Hayes, 2003). These chapters also highlighted the lack of empirical evidence for the influence of these factors on the neural representations and mechanisms underlying morphological learning and generalisation. The view was adopted that neural network simulations and the single-mechanism theories of morphology they most clearly align with best capture the influence of type and token frequency and phonological consistency on morphological learning and generalisation. Chapter 2 introduced the CLS model of learning and memory and emphasised that its formulation is grounded in the observation that neural networks can learn and generalise by gradually adjusting connections in overlapping, distributed networks of neurons. Importantly, the CLS model suggests a central role for memory consolidation in supporting the development of such overlapping, distributed representations. Thus, the CLS model and computational accounts of morphology could be fruitfully combined to derive

predictions regarding the neural representations and mechanisms supporting morphological learning and generalisation, how they might be influenced by type and token frequency and phonological consistency, and change with overnight consolidation.

Building on the work of Mirković and Gaskell (in prep.), Experiment 1 presented in Chapter 5 tackled the issue of whether overnight consolidation changes the influence of type and token frequency and phonological consistency on the generalisation of novel inflectional affixes. Participants were trained on a novel morphological system in which novel stems (e.g. *gleet*; *shiln*) combined with novel plural affixes (e.g. *-aff*, *-opp*) to form new plural nouns (e.g. *gleetaff*; *shilnopp*) referring to the occupation of male and female characters (e.g. *gleetaff*_[fem:plur] = doctors_[fem:plur]; *shilnopp*_[masc:plur] = painters_[masc:plur]). A morphological learning paradigm was developed in which participants learned two sets of novel plural affixes, distinguished by gender, on two successive days. Extending previous work looking at word learning and consolidation (Davis et al., 2009), this paradigm allowed comparing effects of initial learning and overnight consolidation on the generalisation of the newly-learned plural affixes. The new morphological system was designed to mimic the type and token frequency and phonological consistency characteristics of the English past tense such that the influence of overnight consolidation in modulating their influence on generalisation could be assessed. Replicating the results of Mirković and Gaskell (in prep.), participants produced more irregular generalisations and less regular generalisations for novel phonologically ambiguous singulars after a period of overnight consolidation. These novel phonologically ambiguous singulars contained an ambiguous cue, which had been associated with both a high type frequency regular plural affix and a high token frequency irregular plural affix during training, analogous to phonologically ambiguous English verbs (e.g. ‘*flow – flowed*’; ‘*blow – blew*’).

Having replicated the results of Mirković and Gaskell (in prep.) and established overnight changes in the influence of type and token frequency and phonological consistency on the generalisation of newly-learned plural affixes, Experiment 2, discussed in Chapter 8, focused on characterising their underlying

neural representations and functional networks. As reviewed in the introductory chapters, whilst the neural representations and mechanisms underlying morphological processing have been the object of extensive investigation, those underlying morphological learning remain largely unexplored. Chapters 6 and 7 introduced RSA and PPI analyses and emphasised their suitability to test predictions derived from the CLS model. During fMRI scanning, participants listened to the plurals trained immediately before and thus having remained largely unconsolidated and to those trained on the previous day and hence having had a chance to undergo overnight consolidation. The RSA analysis showed increased similarity structure among high type frequency regulars in the left aSTG and decreased similarity structure among high token frequency irregulars in the left aSTG, left pSTG, and left IFG following overnight consolidation. The PPI analysis revealed greater functional connectivity between left aSTG and the left DLPFC for high type frequency regulars before overnight consolidation and greater functional connectivity between the same two regions for high token frequency irregulars after overnight consolidation. These overnight changes in neural representations and functional networks were not directly mirrored by changes in generalisation behaviour and the generalisation results of Experiment 1 were not replicated.

Experiment 3, reported in Chapter 9, sought to examine whether the small changes made to the training stimuli between Experiments 1 and 2 could account for the differences in generalisation results. To this end, the training stimuli was modified with the aim of biasing consolidation and generalisation towards high token frequency irregulars. As in Experiments 1 and 2, participants showed sensitivity to the phonological consistency of the novel singulars in generalising the newly-learned plural affixes but there was no evidence of consolidation-related changes. The next section relates the main contributions of this thesis to the CLS model as well as to consolidation and generalisation theory more broadly.

10.2. Main contributions

As reviewed in Chapters 2 and 3, the CLS model provides a detailed computational and neurobiological framework to characterise the representations and mechanisms supporting learning and generalisation. One motivation for the formulation of the

CLS model comes from the observation of catastrophic interference (McCloskey & Cohen, 1989; Ratcliff, 1990; McClelland et al., 1995; French, 1999) in connectionist networks whereby, if learned too quickly, new information disrupts or erases existing information. A proposed solution assumes a functional and anatomical separation between a hippocampal system that supports the rapid acquisition of information about individual items and a neocortical system that uses slower, interleaved learning to generate distributed, overlapping representations that capture similarities across items (McClelland et al., 1995; Kumaran et al., 2016). Within the CLS model, interactions between the hippocampal and neocortical system play an important role in supporting consolidation. Patterns of hippocampal activity that characterise the encoding of new information drive the gradual integration of this new knowledge in the neocortex. This is thought to occur when the hippocampus replays the contents of a new item or event back to the neocortex during offline periods such as sleep (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996; Peigneux et al., 2004; Rasch et al., 2007).

The CLS model also emphasises different representational schemes in the hippocampal and neocortical system. While the hippocampus uses non-overlapping, pattern-separated representations to emphasise differences between items or events (Bakker et al., 2008; Lacy et al., 2011), the neocortical system employs overlapping, similarity-based representations (McNaughton, 2010; Yamins et al., 2014) that emphasises commonalities across them. It is these overlapping representations that are thought to support generalisation. As highlighted in Chapter 3, however, the developmental course of such overlapping representations, the factors that influence their development, and how they support generalisation remains unclear.

10.2.1. Overnight changes in generalisation

Experiment 1 demonstrated changes in the generalisation of novel plural affixes as a consequence of learning and overnight consolidation. More specifically, it suggested a role for overnight consolidation in modulating the influence of type and token frequency and phonological consistency on generalisation. Recall that participants produced more irregular generalisations (ambiguous) and less regular generalisations for phonologically ambiguous novel singulars after a period of overnight

consolidation. Previous work considering the influence of these factors on generalisation has either focused on proficient speakers generalising to novel forms (Prasada & Pinker, 1993; Albright & Hayes, 2003) or learners generalising to novel forms immediately after training (Bybee & Newman, 1995). The work presented here adds to this body of work by suggesting a role for overnight consolidation in modulating the influence of type and token frequency and phonological consistency on generalisation. The overnight changes in generalisation partly replicate those of Mirković and Gaskell (in prep.) obtained in a between-subject design in which participants remained awake or slept for an equivalent period of time suggesting a role for sleep, rather than the mere passage of time, in accounting for the results.

As discussed in Chapter 5, Mirković and Gaskell (2016) argue for a graded contribution of sleep-dependent consolidation to morphological learning such that arbitrary mappings may show stronger consolidation effects than systematic mappings. By this account, arbitrary mappings may be more dependent on hippocampal pattern separation and subsequently show greater effects of sleep through hippocampal replay whereas systematic mappings may be more swiftly integrated into neocortical networks such that effects of sleep may not be as apparent. The results obtained by Mirković and Gaskell (in prep.) and in Experiment 1 go some way to supporting this proposal. For arbitrary mappings between stems and affixes, as is the case for phonologically ambiguous plurals, the increase in irregular generalisations and the decrease in regular generalisations, may reflect greater dependence on hippocampal pattern separation to keep the two mappings distinct during encoding and lead to greater subsequent effects of overnight consolidation.

However, these results do not necessarily need to imply greater reliance on hippocampal pattern separation. On the one hand, the hippocampal system is not only involved in pattern separation but also in pattern completion, which has been argued to lead to integration in some circumstances (McClelland et al., 1995; O'Reilly & Rudy, 2001; Guzowski et al., 2004; Leutgeb et al., 2004; Schlichting et al., 2015). On the other hand, there have now been several empirical demonstrations that the hippocampus is involved in, and even necessary, for rapid statistical learning (Strange, Duggins, Penny, Dolan, & Friston, 2005; Harrison, Duggins, & Friston, 2006; Turk-

Browne, Scholl, Chun, & Johnson, 2009; Turk-Browne, Scholl, Johnson, & Chun, 2010; Schapiro, Kustner, & Turk-Browne, 2012; Schapiro, Turk-Browne, Norman, & Botvinick, 2016) leading to revisions of the CLS model (Kumaran & McClelland, 2012; Kumaran et al., 2016; Schapiro, Turk-Browne, Botvinick, & Norman, 2017) that outline a role for the hippocampus in representing similarities across items or episodes and in supporting generalisation in some cases.

Furthermore, recall that in Experiment 1 participants produced fewer irregular generalisations (consistent) for novel phonologically diverse singulars after a period of overnight consolidation as well as fewer regular generalisations for novel phonologically consistent singulars. Both of these effects were interpreted as reflecting a strengthening of the regular and irregular mappings manifesting in a reduction of inappropriate generalisations reflecting more ‘optimal’ representations (Stickgold & Walker, 2013). Whilst less straightforwardly than in Experiment 1, the generalisation results in Experiment 2 also showed a reduction in inappropriate responses with overnight consolidation. Thus, overnight changes in generalisation performance reported in this thesis add to the work of Mirković and Gaskell (2016) and suggest that the picture is likely more complex than consolidation simply prioritising arbitrary over systematic mappings.

A more cautious, and arguably more biologically plausible, suggestion is that consolidation processes are ecologically guided by the features of the to-be-remembered items evoked at encoding or in the peri-encoding period (Stickgold & Walker, 2013). While arbitrariness (here, phonological ambiguity) may constitute one such feature, the results obtained in Experiments 1 and 2 suggest that frequency measures are also relevant. Thus, different features likely contribute to determining the consolidation trajectory of newly-learned information. By this account, hippocampal replay may preserve features evoked at encoding or in the peri-encoding period rather than indiscriminately replay all items or events or prioritising arbitrary over systematic mappings.

10.2.2. Overnight changes in neural representations

Using RSA, Experiment 2 demonstrated overnight changes in the neural representations of newly-learned inflected words for the first time. Recall that there

was increased similarity structure among plurals sharing the same high type frequency regular affix in the left aSTG and decreased similarity structure among plurals sharing the same high token irregular affix in the left aSTG, left pSTG, and left IFG following overnight consolidation. The increase in similarity structure among high type frequency regulars is consistent with previous behavioural work suggesting that the formation of context-independent derivational affix representations occurs after a period of offline consolidation (Merkx et al., 2011; Tamminen et al., 2012, 2015) and benefits from high type frequency (Tamminen et al., 2015). It is also consistent with single-mechanism theories of morphology which consider type frequency to be a determining factor of productivity and previous network simulations (Plunkett & Marchman, 1991, 1993; Plunkett & Juola, 1999) showing that high type frequency is needed for a network to shift from memorising specific mappings to discovering regularities shared across several mappings.

Furthermore, as noted in Chapter 3, the CLS model predicts that generalisation emerges from stored knowledge of the shared features across items encoded in memory by overlapping representations. On this view, affixes with high type frequency should benefit more from the overlap in neocortical representations and consequently be more likely to form representations that are independent of the stem contexts in which they were learned. In their MEG study, Leminen et al. (2016) reported increased responsiveness of the left STG to real stems combined with trained derivational affixes following overnight consolidation. The authors argued that this result could be consistent with the development of an affix representation following overnight consolidation. As acknowledged by the authors, however, their result could not distinguish whether the development of such an affix representation was driven by just the affix, the whole affixed form, or both. Leminen et al. (2016) also used the same stimuli in two MEG recording sessions such that effects of overnight consolidation could not be disentangled from effects of re-exposure to the trained affixes.

In contrast, Experiment 2 employed a single scanning session following two days of training with different plurals. Changes in the underlying representation of the plurals could therefore be more confidently attributed to consolidation processes

rather than additional learning taking place from re-exposure. Furthermore, the combination of fMRI and RSA used in Experiment 2 augment the results of Leminen et al. (2016) by providing evidence for similarity structure among newly-learned plurals being driven by inflected forms sharing the same regular affix. The results reported in this thesis thus make a unique contribution to the existing body by providing the first neural evidence consistent with an affix representation emerging as a consequence of learning and overnight consolidation and that high type frequency plays a key role in supporting the development of such representation. Leminen et al. (2016) also suggested that their results are consistent with the left STG playing a key role in the representation of lexico-semantic knowledge. The results obtained in Experiment 2 showed that the best fit for a model testing for an affix representation was restricted to the left aSTG. As discussed in Chapter 8, the characterisation of the anterior portion of the STG remains elusive and future work will be required to gain a better understanding of its roles in morphological learning, consolidation, and generalisation. However, the results of Experiment 2 suggest that it may be implicated in the representation of newly-acquired morphological knowledge.

The development of similarity structure among high type frequency regulars consistent with the development of an affix representation following overnight consolidation also bears upon the existing literature on the role of sleep in schema formation. Recall from Chapter 3 that Lewis and Durrant's (2011) model of schema formation suggests that more strongly potentiated connections for overlapping elements between memories have a higher chance of withstanding synaptic downscaling during sleep and thus facilitate the emergence of a schema. Whilst further work assessing the role of sleep in the learning, consolidation, and generalisation of novel morphology will be required, the emergence of similarity structure among high type frequency regulars observed in Experiment 2 would be consistent with this account. Furthermore, previous studies (Ellenbogen et al., 2007; Durrant et al., 2011; Lau et al., 2011; Coutanche et al., 2013) have reported enhanced generalisation performance following a single period of sleep consistent with the rapid development of overlapping neocortical representations that can support generalisation.

Despite not observing increases in regular generalisations following overnight consolidation in the work presented in this thesis, the increase in similarity structure among high type frequency regulars is nonetheless consistent with the suggestion that sleep may play a role in schema formation. These results are significant since they provide evidence for similarity structure emerging in neocortical representational areas after a 24hr period of offline consolidation containing sleep. As reviewed in Chapter 3, the earliest formulation of the CLS model (McClelland et al., 1995) considered the development of similarity structure as being a protracted process spanning several days to months. The evidence provided here suggests that a single night's sleep may be sufficient to support the development of similarity structure and, as such, aligns more clearly with Lewis and Durrant's (2011) model of schema formation during sleep.

Another important, and perhaps more intriguing, finding concerns the decrease in similarity structure among high token frequency irregulars in neocortical representational areas following overnight consolidation. As discussed in Chapter 3, the CLS model predicts that many repetitions of the same item will lead to changes in the synaptic connections involved in the representation of that specific item but not support the development of similarity structure (McClelland et al., 1995; Kumaran & McClelland, 2012). The decrease in similarity structure among high token frequency irregulars could thus be interpreted as strengthening their idiosyncratic representations. As already noted, Lewis and Durrant's (2011) schema formation model predicts a role for high type frequency in supporting the development of overlapping representations during sleep. Could the model also account for the role of token frequency in strengthening idiosyncratic representations? In line with Lewis and Durrant's (2011) model, Stickgold and Walker (2013) note that when replay of a recently formed memory is accompanied by the parallel activation of a larger set of recently formed memories sharing common features, schema formation may occur. They also note that when a single item memory is reactivated in conjunction with an existing network of connections or schema, item integration may occur. This would, for example, account for the emergence of lexical competition between new (e.g. *cathedruke*) and existing words (e.g. *cathedral*) with consolidation (Gaskell & Dumay, 2003; Dumay & Gaskell, 2007, 2012). Of particular relevance here, Stickgold and

Walker (2013) argue that if neither schema formation nor item integration occurs, sleep may simply enhance the representation of recently encoded item-specific memories. Thus, the decrease in similarity structure among high token frequency regulars following overnight consolidation may reflect the fact that without sufficiently high type frequency to support the development of a new schema (here, an affix representation) or similarity with existing words to support integration, overnight consolidation may instead work to strengthen their idiosyncratic, whole-word representations.

This kind of characterisation of the consolidation process may be seen as similar to Lewis and Durrant's (2011) model but emphasises the relevance of type and token frequency in determining the consolidation trajectory of newly-learned items (Figure 49). As noted in Chapter 8, the initial similarity structure among high token frequency irregulars remains intriguing and more difficult to relate to existing theoretical frameworks. Speculatively, it may reflect their greater salience. As the consolidation process ensues, low type frequency may be insufficient to support the development of overlapping representations. Initial salience tags attached to high token frequency during learning may subsequently be used during sleep to enhance item-specific representations. Of course, this account is predicated on the assumption that similarity structure observed before (among high token frequency irregulars) and after (among high type frequency regulars) a period of overnight consolidation may index different kinds of representations (i.e. one based on salience, the other reflecting the formation of a schema). Thus, questions relating to the factors determining the consolidation process for different types of information as well as the possibility of wakeful tagging and its underlying neural mechanisms remain important unresolved issues for future research.

It should be noted that the changes in similarity structure observed in Experiment 2 are also largely consistent with Bybee's usage-based model (1985, 1988, 1995, 2001) introduced in Chapter 3. By this account, when several inflected words share the same affix (i.e. high type frequency), these inflected words will be related to each other and the representation of the affix will emerge. The increase in similarity structure among high type frequency regulars in Experiment 2 provides neural

evidence that is consistent with this account and further suggests a role for overnight consolidation in supporting the development of an affix representation. Bybee's usage-based model (1985, 1988, 1995, 2001) also proposes that inflected words with high individual token frequency are less likely to be acquired by forming relations with other inflected words. The higher the token frequency of an inflected word the stronger its lexical representation. On this view, the decrease in similarity structure among high token frequency irregulars observed in Experiment 2 would be consistent with a strengthening of their whole-word representations.

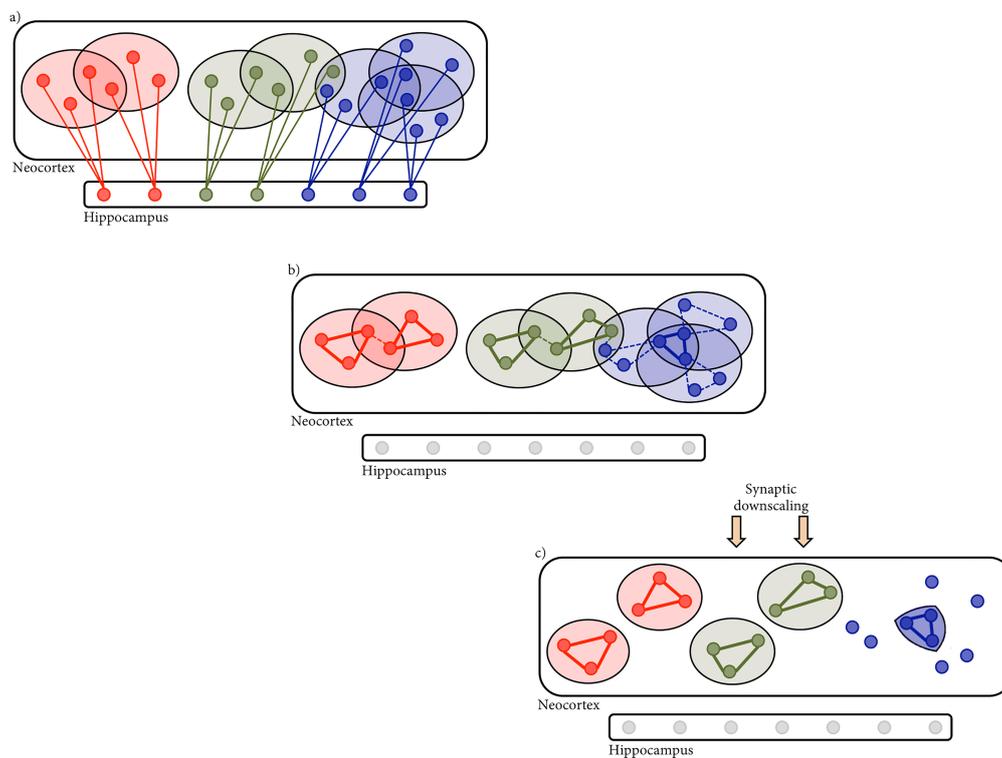


Figure 49. Overnight consolidation of type and token frequency

Adaptation of Lewis and Durrant (2011) model of schema formation to account for the development of neocortical representations for high token and high type frequency inflected plurals through hippocampal replay during overnight consolidation. a) hippocampal memories for two high token frequency irregular consistent plurals (red), two high token frequency ambiguous plurals (green), and three high type frequency regular plurals (blue) are reactivated during sleep. This reactivation includes some neocortical neurons unique to each memory, and some that are shared by memories. b) for high token frequency plurals (red and green) unique neurons are potentiated more strongly and as a result develop stronger connection to each other (solid lines) than overlapping neurons (dotted lines). For high type frequency plurals (blue) shared neurons are potentiated more strongly and as a result develop stronger connections to each other (solid lines) than neurons not involved in this overlap (dotted lines). c) following synaptic downscaling, only these strong connections remain intact. This set of neuron now represents abstraction from the high type frequency regulars, encoding their shared features only (i.e. affix representation) and item-specific representations for high token frequency irregulars.

10.2.3. Overnight changes in functional networks

Using PPI, Experiment 2 showed overnight changes in the functional networks underlying newly-learned inflected words for the first time. Recall that there was greater functional connectivity between the left aSTG and the left DLPFC for high type frequency regulars before overnight consolidation and greater functional connectivity between the same two regions for high token frequency irregulars after overnight consolidation. These results are consistent with the suggestion outlined in the previous section that different types of information may undergo different consolidation trajectories. Two lines of evidence, each related to a different purported role of the DLPFC in memory processing, are relevant. As discussed in Chapter 8, one line of evidence suggests that prefrontal regions contribute to memory through control over retrieval processes within other brain areas (Moscovitch, 1992; Buckner & Wheeler, 2001; Miller & Cohen, 2001; Dobbins, Foley, Schacter, & Wagner, 2002). The DLPFC is thought to control memory retrieval by selecting memories relevant to the current context and suppressing irrelevant memories (Anderson et al., 2016). Much of the work on the role of the DLPFC in memory suppression has thus far focused on hippocampally-dependent memories. However, as noted by Anderson et al. (2016), it is likely that over time, memory consolidation leads to an integrated, established representation in the neocortex. The results obtained in Experiment 2 suggest that the DLPFC may functionally interact with neocortical representational areas to orchestrate the retrieval of different mappings that are competing for consolidation.

Relatedly, previous work has suggested an important role for the DLPFC in resolving competition between memories at retrieval (Badre, Poldrack, Paré-Blagoev, Insler, & Wagner, 2005; Thompson-Schill, Bedny, & Goldberg, 2005). On this view, high token frequency irregulars may initially be prioritised given their greater representational structure. As a result, the DLPFC may work to inhibit or suppress the retrieval of high type frequency regulars as indexed by increased functional connectivity between the aSTG and DLPFC. As similarity structure among high type frequency regulars develops with overnight consolidation, the DLPFC may subsequently work to inhibit or suppress the retrieval of high token frequency regulars

as indexed by the opposite pattern of functional connectivity. This interpretation is, of course, limited by the lack of previous data available and poorly specified nature of the interactions between PFC and neocortical representational areas in memory inhibition and suppression or how they may be implicated in the consolidation process. Such questions remain open for future research.

A second line of evidence suggests that prefrontal regions contribute to memory through control over encoding processes within other brain regions. According to the Hemispheric Encoding Retrieval Asymmetry model (Tulving, Kapur, Craik, Moscovitch, & Houle, 1994; Habib, Nyberg, & Tulving, 2003) the left PFC is more involved in episodic memory encoding, whereas the right PFC is more involved in memory retrieval. In line with this model, repetitive transcranial magnetic stimulation studies have provided evidence for a causal role for the left DLPFC in the encoding of novel verbal stimuli (Sandrini, Cappa, Rossi, Rossini, & Miniussi, 2003). At encoding, new stimuli are initially processed by representational neocortical areas before being transmitted to MTL structures. At this stage, PFC is thought to play a role in providing top-down control over the encoding process by guiding, modifying, and elaborating the representations in representational neocortical areas and MTL structures on the basis of current goals and task demands, and by ensuring the representations are amenable to long-term storage (Simons & Spiers, 2003). The DLPFC has been suggested to play a key role in such organisational functions (Fletcher, Shallice, & Dolan, 1998).

By this account, the results reported in Experiment 2 may be interpreted as suggesting the differential involvement of the left DLPFC top-down control processes depending on the type and token frequency and training day of the novel inflected words. The lack of similarity structure among high type frequency regulars on Day 9 may make them less amenable to long-term storage. As a result, greater functional connectivity between the left aSTG and the left DLPFC may reflect greater demands on the organisational functions of the left DLPFC to ensure successful subsequent consolidation and long-term storage. By contrast, top-down control processes from the left DLPFC may not be required for high token frequency irregulars that have

initial similarity structure and thus may already be amenable to consolidation and long-term storage.

The CLS model proposes that new memories are initially dependent on the hippocampus and that over the course of consolidation they become represented neocortically. Other accounts (see Preston & Eichenbaum, 2013) suggest that context-guided retrieval of appropriate memory networks for new memories, supported by the PFC, plays a key role in the consolidation of new memories. While these studies typically focus on the mPFC as playing an important role in supporting schema formation, the results obtained here suggest that DLPFC may also play a role in memory consolidation. Strategic memory processing by the prefrontal cortex with regards to either encoding or retrieval may contribute to memory consolidation and underpin different consolidation trajectories for different mappings in a complex learning environment. Indeed, both RSA and PPI results obtained in Experiment 2 may reflect different consolidation trajectories for high type and high token frequency inflected words.

Several questions remain open for future research. Most work thus far, including the work presented in this thesis, has examined brain processes related to consolidation either immediately or at relatively short time periods after learning. The suggestion that different types of information may undergo different consolidation trajectories would require further research at longer timescales. Another challenge moving forward will be to disentangle the roles of memory encoding and retrieval processes in the consolidation, and representation of new memories. It is particularly difficult to design experiments that unequivocally control for memory encoding and retrieval processes and it is likely that both occurred in Experiment 2 as participants listened to the trained plurals. Acquiring fMRI data as participants learn new inflected words may provide a means to isolate encoding processes. Similarly, recall tasks (e.g. picture naming) in the scanner may allow to better target retrieval processes.

10.3. Limitations

Experiment 2 showed differences in the neural representations and functional networks underlying inflected words learned on the same day as scanning and those

learned on the day before in a within-subject design. A similar comparison between groups of participants that sleep or remain awake following training has yet to be conducted. Therefore, the possibility that some of the findings might be due to interference between novel inflected words learned first and second, forgetting of the words trained first, or due to the mere passage of time in the absence of sleep cannot be ruled out. Results from the plural elicitation generalisation tasks in all three experiments showed that affix training day had an effect on overall performance with participants producing more incorrect responses for Day 8 compared to Day 9 affixes. As noted in Chapter 5, this could be due to forgetting of Day 8 affixes or interference from Day 9 affixes, two possibilities that cannot be distinguished on the basis of present data (Mensink & Raaijmakers, 1988; Anderson, 2003). Future studies that include between-group assessment of inflected words learned and tested on the same day or subsequent days would be helpful in addressing this limitation. It seems likely, however, that several processes including strengthening of item-specific memories, abstraction of features shared across several memories, reduction of interference between memories, and forgetting of irrelevant memories work synergistically to support the development and generalisation of new knowledge. A challenge for future research will be to gain a better understanding of how these processes contribute to learning, consolidation, and generalisation. As such, the limitations of the work presented in this thesis are testament to the complexities and intricacies of the field of learning and memory as a whole. It is difficult to design experiments that unequivocally target one process and this may be particularly true for complex, multi-componential stimuli such as inflectional morphology.

The generalisation tasks employed in the three experiments were non-speeded. As discussed in Chapter 3, Tamminen et al. (2012) suggested that the extent to which generalisation tasks require rapid access to newly-learned information may be an important factor to consider. Recall that Tamminen et al. (2012) found that participants were able to generalise new morphological knowledge in a non-speeded task immediately after learning, but in a speeded task generalisation was only present at a delayed test two days after training. Tamminen et al. (2012) suggested that immediate generalisation may reflect the use of context-dependent representations that may not be sufficient for online linguistic processing on the speeded task. The overnight

changes in neural representations and functional networks observed in Experiment 2 do suggest that different representations can underly very similar patterns of generalisation. Future studies that include both speeded and non-speeded generalisation tests may be helpful in establishing how different types of representations may underpin generalisation. Furthermore, future imaging data acquired whilst participants perform generalisation tasks would allow to characterise the kinds of representations and mechanisms that are recruited during generalisation.

10.4. Future directions

The main contributions and limitations outlined in the sections above raise a number of questions for future research. One of the major questions raised by the work presented in this thesis concerns the precise role of sleep in the consolidation and generalisation of novel morphology. As reviewed in Chapter 2, whilst a growing body of evidence suggests a central role for sleep in memory consolidation, a great deal of uncertainty remains about which stages of the complex sleep architecture are relevant. Using polysomnography, future work may examine whether particular stages or indices of sleep are associated with morphological learning and generalisation. Future work may also determine whether morphological learning and generalisation can be biased or enhanced experimentally by inducing replay of newly-learned morphological knowledge during sleep. A series of recent studies have shown that the presentation of memory cues associated with previously learned material during SWS benefits the consolidation of both declarative and procedural memories (Rasch et al., 2007; Rudoy, Voss, Westerberg, & Paller, 2009; Antony, Gobel, O'Hare, Reber, & Paller, 2012; Bendor & Wilson, 2012; Fuentemilla et al., 2013; Rihm, Diekelmann, Born, & Rasch, 2014; Schreiner & Rasch, 2015; Creery, Oudiette, Antony, & Paller, 2015; Cairney, Lindsay, Sobczak, Paller, & Gaskell, 2016). For example, as reviewed in Chapter 2, Rasch et al. (2007) have shown improvements in memory for object locations after sleep when odour cues associated with individual objects during learning are presented during SWS.

Together, these cueing techniques are known as targeted memory reactivation (TMR; Oudiette & Paller, 2013) and have recently begun to be applied to language learning paradigms (see Schreiner & Rasch, 2017 for a review). For example, Schreiner

and Rasch (2015) have shown that auditory re-exposure to previously trained Dutch words during SWS improves later memory for their German translations in native German speakers. More recently, Batterink and Paller (2017) have shown that auditory cueing of a previously trained artificial grammar containing dependency rules between nonwords during SWS led to modest but significant improvements in generalisation of the rules to novel sequences of nonwords. These initial findings suggest that TMR may function not only to strengthen item-specific memory but may also support the development of more abstract knowledge that can support generalisation. In the context of the work presented in this thesis, future work may look at the influence of TMR on the learning and generalisation of inflectional patterns trained with varying type and token frequency and phonological consistency. Such manipulations might allow delineating the potential effects of sleep on modulating the influence of these factors on consolidation and generalisation. For example, findings showing that the influence of type and token frequency and phonological consistency on generalisation might be differentially sensitive to reactivation during sleep might imply different mechanisms contributing to their consolidation. The overnight changes in the neural representations and functional networks underlying newly-learned inflected words reported in this thesis together with the emergence of TMR in language learning research provide a strong basis for future work to clarify the role of sleep in morphological learning and generalisation.

An important challenge for future research will be to determine which types of information may be preferentially consolidated when complex, multi-componential stimuli are encountered. The possibility of pre-sleep tagging of novel information as being particularly relevant to future behaviour has been considered in interpreting the results of Experiments 1 and 2 and offered as a potentially explanatory mechanism for the preferential consolidation of some inflectional patterns over others. Future work will need to characterise the neural mechanisms that might initially be responsible for creating such tags. Rauchs et al. (2011) have shown greater hippocampal activation during the initial encoding of novel items cued for remembering compared to those cued for forgetting. Importantly, this increase in hippocampal activity at encoding predicted subsequent recognition of the novel items after a night's sleep. Rauchs et al. (2011) argue that the hippocampal activity at encoding may therefore reflect the

tagging of items for subsequent sleep-dependent consolidation. This interpretation is consistent with reports of wake-dependent hippocampal replay in rodents (Foster & Wilson, 2006) and enhanced post-learning hippocampal-neocortical functional connectivity following a task with high subsequent memory in humans (Tambini, Ketz, & Davachi, 2010). In the context of the work presented in this thesis, future work may include an encoding fMRI session such that hippocampal engagement, patterns of representational similarity, and functional connectivity during the initial learning of the novel morphological system may be related to subsequent changes in representation and generalisation following overnight consolidation. Such work would also provide an opportunity to gain further insights into the reciprocity between learning and consolidation processes.

Yet another direction for future research will be to clarify the conditions under which consolidation may strengthen item-specific representations or abstract features common to several items to generate new schemas. The evidence reported in Experiment 2 suggests that it is likely that this is not an either or process and that consolidation may support both. However, the precise features of learning episodes, the potential tagging processes that may occur at encoding or in the peri-encoding period, and the physiological sleep mechanisms that may lead to one type of consolidation over the other remain to be established. Clearly, the questions of how such representations may then support generalisation, the precise conditions under which consolidation may or may not benefit generalisation, and the underlying neural representations and mechanism of memory consolidation remain important unresolved issues.

In sum, a wealth of exciting new findings can be expected from future research aimed at uncovering the relationships between morphological learning, memory consolidation, and generalisation. Success will require an alliance between these traditionally separate research endeavours and the use of complementary research methods including behavioural experiments, neuroimaging, and polysomnography. Theoretical models and hypotheses about the role of memory consolidation in language learning and generalisation specifically but also in learning and generalisation more broadly will likely be transformed considerably in the process.

The work presented in this thesis has made a unique contribution to the existing body of knowledge by demonstrating overnight changes in the neural representations and functional networks underlying newly-learned inflectional morphology. It has also shown that behavioural paradigms coupled with advanced multivariate analyses of fMRI data can be successfully combined to test predictions derived from models of learning, consolidation, and generalisation. Finally, it has provided a strong basis and generated a number of questions for future work to clarify the role of consolidation in the learning and generalisation of morphological knowledge.

Appendix 1

Peak voxels (maximum of three peaks/clusters, separated by > 8mm) showing significant effects in the univariate analysis

Brain region (AAL)	x	y	z	Cluster size (voxels)	Z-value
Plurals > baseline					
Left superior temporal gyrus	-56	-6	0	3948	6.44
	-44	-20	6		6.31
	-60	-26	6		5.99
Right superior temporal gyrus	54	-16	6	4121	6.41
	48	-32	12		6.27
	40	-30	8		5.99
Affix training day x phonological consistency x affix regularity interaction					
Right medial orbitofrontal cortex	26	42	2	162	4.10
	18	48	-6		3.93
	26	34	0		3.64
Left anterior cingulate cortex	-10	22	42	144	4.07
	-16	24	32		3.99
	-10	22	26		3.47

Results thresholded at $p < .001$ uncorrected at the voxel level and $p < .05$ FWE corrected at the cluster level

Appendix 2

Significant effects in the ROI-to-ROI functional connectivity analysis

Seed	Target	T-value	p-unc	p-corr
Main effect day				
day 9 > day 8				
Left aHC	Left pSTG	3.08	0.003	0.012
Left pSTG	Left aHC	3.05	0.003	0.013
Main effect day (type frequency)				
day 8 > day 9				
left aSTG	Left IFG	3.18	0.002	0.006
	Left pHC	3.11	0.003	0.006
left IFG	Left aSTG	3.14	0.003	0.010
	Left pHC	1.89	0.036	0.049
	left pSTG	1.89	0.036	0.049

Results thresholded at $p < .001$ uncorrected (p-unc) and FDR corrected (p-corr) for multiple comparisons

Appendix 3

Significant effects in the seed-to-voxel functional connectivity analysis

Seed ROI	Functionally connected brain region (AAL)	x	y	z	cluster size (voxels)	p-FWE (cluster-level)
Affix training day x affix regularity						
Left aSTG	Left middle frontal gyrus	-42	36	42	31	0.001
	Left middle frontal gyrus	-42	40	30	17	0.045
Left pSTG	Left middle frontal gyrus	-46	38	30	18	0.033
Main effect phonological consistency						
consistent > ambiguous						
Left aSTG	Left inferior parietal lobule	-46	-42	50	36	< 0.001
	Left inferior parietal lobule	-46	-28	40	27	0.002
	Right inferior temporal gyrus	48	-64	-8	20	0.015
Left pSTG	Left inferior parietal lobule	-50	-34	42	21	0.012
Left aHC	Left inferior temporal gyrus	-40	-46	-16	27	0.002
ambiguous > consistent						
Left aHC	Right supramarginal gyrus	56	-42	30	46	< 0.001
Main effect affix regularity						
irregular > regular						
Left aHC	Left lingual gyrus	-22	-72	-4	34	< 0.001
	Left superior occipital gyrus	-14	-94	20	26	0.003
	Right fusiform gyrus	28	-60	-8	17	0.046
Left IFG	Left inferior parietal lobule	-34	-52	50	30	0.001
	Right inferior parietal lobule	40	-40	52	17	0.049

Results thresholded at $p < .001$ uncorrected at the voxel level and $p < .05$ FWE corrected at the cluster level

References

- Adams, J. P., & Dudek, S. M. (2005). Late-phase long-term potentiation: getting to the nucleus. *Nature Reviews Neuroscience*, *6*(9), 737–743.
- Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, *45*(7), 1363–1377.
- Aeschbach, D., Cutler, A. J., & Ronda, J. M. (2008). A role for non-rapid-eye-movement sleep homeostasis in perceptual learning. *Journal of Neuroscience*, *28*(11), 2766–2772.
- Agresti, A. (2002). *Categorical data analysis*. Hoboken, New Jersey: John Wiley & Sons.
- Albright, A., & Hayes, B. (2003). Rules vs . analogy in English past tenses : a computational/experimental study. *Cognition*, *90*(2), 119–161.
- Alegre, M., & Gordon, P. (1999). Frequency effects and the representational status of regular inflections. *Journal of Memory and Language*, *40*(1), 41–61.
- Alvarez, P., & Squire, L. R. (1994). Memory consolidation and the medial temporal lobe : A simple network model. *Proceedings of the National Academy of Sciences*, *91*(15), 7041–7045.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409–429.
- Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review*, *8*(4), 629–647.
- Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, *49*(4), 415–445.
- Anderson, M. C., Bunce, J. G., & Barbas, H. (2016). Prefrontal-hippocampal pathways underlying inhibitory control over memory. *Neurobiology of Learning and Memory*, *134*(A), 145–161.
- Antony, J. W., Gobel, E. W., O'Hare, J. K., Reber, P. J., & Paller, K. A. (2012). Cued memory reactivation during sleep influences skill learning. *Nature Neuroscience*, *15*(8), 1114–1116.
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442–481.
- Ashby, F. G., & Ell, S. n. W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences*, *5*(5), 204–210.
- Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and

- decision bound models of categorization. *Journal of Mathematical Psychology*, 37(3), 372–400.
- Ashby, F. G., & Maddox, W. T. (1998). Stimulus categorization. In *Handbook of Perception and Cognition: Measurement, Judgment, and Decision Making* (pp. 251–301). New York: Academic.
- Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, 56, 149–78.
- Ashby, F. G., & Maddox, W. T. (2010). Human category learning 2.0. *Annals of the New York Academy of Sciences*, 1224, 147–161.
- Atienza, M., & Cantero, J. L. (2008). Modulatory effects of emotion and sleep on recollection and familiarity. *Journal of Sleep Research*, 17(3), 285–294.
- Badre, D., Poldrack, R. A., Paré-Blagoev, E. J., Insler, R. Z., & Wagner, A. D. (2005). Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron*, 47(6), 907–918.
- Bakker, A., Kirwan, C. B., Miller, M., & Stark, C. E. L. (2008). Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science*, 319(5870), 1640–1642.
- Barch, D. M., Braver, T. S., Sabb, F. W., & Noll, D. C. (2000). Anterior Cingulate and the Monitoring of Response Conflict : Evidence from an fMRI Study of Overt Verb Generation. *Journal of Cognitive Neuroscience*, 12(2), 298–309.
- Barnes, C. A., McNaughton, B. L., Mizumori, S. J. Y., Leonard, B. W., & Lin, L.-H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, 83, 287–300.
- Barnes, J. M., & Underwood, B. J. (1959). “Fate” of first-list associations in transfer theory. *Journal of Experimental Psychology*, 58(2), 97–105.
- 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.
- Barrett, T. R., & Ekstrand, B. R. (1972). Effect of sleep on memory: III. Controlling for time-of-day effects. *Journal of Experimental Psychology*, 96(2), 321–327.
- Barsalou, L. (2003). Situated simulation in the human conceptual system. *Language and Cognitive Processes*, 18(5–6), 513–562.
- Battaglia, F. P., Sutherland, G. R., & McNaughton, B. L. (2004). Hippocampal sharp wave bursts coincide with neocortical “up-state” transitions. *Learning & Memory*, 11(6), 697–704.
- Batterink, L. J., & Paller, K. A. (2017). Sleep-based memory processing facilitates grammatical generalization: Evidence from targeted memory reactivation. *Brain*

- and Language*, 167, 83–93.
- Bayley, P. J., Hopkins, R. O., & Squire, L. R. (2003). Successful recollection of remote autobiographical memories by amnesic patients with medial temporal lobe lesions. *Neuron*, 38(1), 135–144.
- Bayley, P. J., & Squire, L. R. (2005). Failure to Acquire New Semantic Knowledge in Patients With Large Medial Temporal Lobe Lesions Peter. *Hippocampus*, 15(2), 273–280.
- Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A Component Based Noise Correction Method (CompCor) for BOLD and Perfusion Based fMRI. *NeuroImage*, 37(1), 90–101.
- Bemis, D. K., & Pylkkänen, L. (2011). Simple composition: a magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience*, 31(8), 2801–2814.
- Bendor, D., & Wilson, M. A. (2012). Biasing the content of hippocampal replay during sleep. *Nature Neuroscience*, 15, 1439–1444.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828.
- Beretta, A., Campbell, C., Carr, T. H., Huang, J., Schmitt, L. M., Christianson, K., & Cao, Y. (2003). An ER-fMRI investigation of morphological inflection in German reveals that the brain makes a distinction between regular and irregular forms. *Brain and Language*, 85(1), 67–92.
- Berko, J. (1958). The child's learning of English morphology. *Word*, 14(2–3), 150–177.
- Bertram, R., Hyönä, J., & Laine, M. (2000). The role of context in morphological processing: Evidence from Finnish. *Language and Cognitive Processes*, 15(4–5), 367–388.
- Bertram, R., Laine, M., & Karvinen, K. (1999). The Interplay of Word Formation Type, Affixal Homonymy, and Productivity in Lexical Processing: Evidence from a Morphologically Rich Language. *Journal of Psycholinguistic Research*, 28(3), 213–226.
- Bertram, R., Schreuder, R., & Baayen, R. H. (2000). The balance of storage and computation in morphological processing: The role of word formation type, affixal homonymy and productivity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 26(2), 489–511.
- Bird, H., Lambon Ralph, M. A., Seidenberg, M. S., McClelland, J. L., & Patterson, K. (2003). Deficits in phonology and past-tense morphology: What's the connection? *Journal of Memory and Language*, 48(3), 502–526.

- Biswal, B. B., Van Kylen, J., & Hyde, J. S. (1997). Simultaneous Assessment of Flow and BOLD Signals in Resting-State Functional Connectivity Maps. *NMR in Biomedicine*, *10*(4–5), 165–170.
- Blank, H., & Davis, M. H. (2016). Prediction Errors but Not Sharpened Signals Simulate Multivoxel fMRI Patterns during Speech Perception. *PLoS Biology*, *14*(11), 1–32.
- Bontempi, B., Laurent-Demir, C., Destrade, C., & Jaffard, R. (1999). Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature*, *400*(6745), 671–675.
- Borg, I., & Groenen, P. J. F. (2005). *Modern Multidimensional Scaling – Theory and Applications*. New York, NY: Springer New York.
- Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends in Cognitive Sciences*, *8*(12), 539–546.
- Boudelaa, S., & Gaskell, M. G. (2002). A re-examination of the default system for Arabic plurals. *Language & Cognitive Processes*, *17*(3), 321–343.
- Bowers, J. S., Davis, C. J., & Hanley, D. A. (2005). Interfering neighbours: The impact of novel word learning on the identification of visually similar words. *Cognition*, *97*(3), 45–54.
- Bozic, M., Fonteneau, E., Su, L., & Marslen-Wilson, W. D. (2015). Grammatical analysis as a distributed neurobiological function. *Human Brain Mapping*, *36*(3), 1190–1201.
- Bozic, M., Tyler, L. K., Su, L., Wingfield, C., & Marslen-Wilson, W. (2013). Neurobiological systems for lexical representation and analysis in English. *Journal of Cognitive Neuroscience*, *25*(10), 1678–1691.
- Braver, T. S., Barch, D. M., Gray, J. R., Molfese, D. L., & Snyder, A. (2001). Anterior Cingulate Cortex and Response Conflict : Effects of Frequency, Inhibition and Errors. *Cerebral Cortex*, *11*(9), 825–836.
- Breitenstein, C., Jansen, A., Deppe, M., Foerster, A. F., Sommer, J., Wolbers, T., & Knecht, S. (2005). Hippocampus activity differentiates good from poor learners of a novel lexicon. *NeuroImage*, *25*(3), 958–968.
- Broadbent, D. E. (1967). Word-Frequency effect and response bias. *Psychological Review*, *74*(1), 1–15.
- Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.
- Brooks, T. L., & Cid de Garcia, D. (2015). Evidence for morphological composition in compound words using MEG. *Frontiers in Human Neuroscience*, *9*(215), 1–8.

- Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.
- Buckner, R. L., & Wheeler, M. E. (2001). The cognitive neuroscience of remembering. *Nature Reviews Neuroscience*, 2(9), 624–634.
- Bunsey, M., & Eichenbaum, H. (1996). Conservation of hippocampal memory function in rats and humans. *Nature*, 379(6562), 255–257.
- Burnham, W. C. (1903). Retroactive amnesia: Illustrative cases and a tentative explanation. *American Journal of Psychology*, 14(3–4), 383–396.
- Bybee, J. L. (1985). *Morphology: A study of the relation between meaning and form*. Philadelphia, PA: John Benjamins.
- Bybee, J. L. (1988). Morphology as lexical organization. In M. Hammond & M. Noonan (Eds.), *Theoretical morphology*. San Diego, CA: Academic press.
- Bybee, J. L. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, 10(5), 425–455.
- Bybee, J. L. (2001). *Phonology and language use*. Cambridge, UK: Cambridge University Press.
- Bybee, J. L., & Hopper, P. (2001). *Frequency and the emergence of linguistic structure*. Amsterdam: John Benjamins.
- Bybee, J. L., & Moder, C. L. (1983). Morphological classes as natural categories. *Language*, 59(2), 251–270.
- Bybee, J. L., & Newman, J. E. (1995). Are stem changes as natural as affixes? *Linguistics*, 33, 633–654.
- Bybee, J. L., & Slobin, D. I. (1982). Rules and Schemas in the Development and Use of the English past Tense. *Language*, 58(2), 265–289.
- Cairney, S. A., Lindsay, S., Sobczak, J. M., Paller, K. A., & Gaskell, M. G. (2016). The benefits of targeted memory reactivation for consolidation in sleep are contingent on memory accuracy and direct cue-memory associations. *Sleep*, 39(5), 1139–1150.
- Carlson, T. A., Schrater, P., & He, S. (2003). Patterns of activity in the categorical representations of objects. *Journal of Cognitive Neuroscience*, 15(5), 704–71.
- Carota, F., Bozic, M., & Marslen-Wilson, W. D. (2013). Decompositional Representation of Morphological Complexity: Multivariate fMRI Evidence from Italian. *Journal of Cognitive Neuroscience*, 26(3), 194–198.
- Carpenter, G. A., & Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37(1), 54–115.

- Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, *14*(2), 147–153.
- Casadio, A., Martin, K. C., Giustetto, M., Zhu, H., Chen, M., Bartsch, D., ... Kandel, E. R. (1999). A transient, neuron-wide form of CREB-mediated long-term facilitation can be stabilized at specific synapses by local protein synthesis. *Cell*, *99*(2), 221–237.
- Cattell, J. M. (1886). The time it takes to see and name objects. *Mind*, *11*(41), 63–65.
- Chai, X. J., Castañón, A. N., Ongür, D., & Whitfield-Gabrieli, S. (2012). Anticorrelations in resting state networks without global signal regression. *NeuroImage*, *59*(2), 1420–1428.
- Chernigovskaya, T., & Gor, K. (2000). The complexity of paradigm and input frequencies in native and second language verbal processing: Evidence from Russian. *Language and Language Behavior*, *3*, 20–37.
- Chomsky, N. (1956). Three models for the description of language. *IRE Transactions on Information Theory*, *2*(3), 113–124.
- Chomsky, N. (1959). On Certain Formal Properties of Grammars. *Information and Control*, *2*(2), 137–167.
- Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York, NY: Harper and Row.
- Cipolli, C. (1995). Sleep, Dreams and memory: an overview. *Journal of Sleep Research*, *4*(1), 2–9.
- Clahsen, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences*, *22*(6), 991–1013.
- Clahsen, H., & Rothweiler, M. (1992). Inflectional rules in children's grammars: Evidence from the development of participles in German. *Morphology Yearbook*, 1–34.
- Clahsen, H., Rothweiler, M., Woest, A., & Marcus, G. F. (1992). Regular and irregular inflection in the acquisition of German noun plurals. *Cognition*, *45*(3), 225–255.
- Clarke, a., & Tyler, L. K. (2014). Object-Specific Semantic Coding in Human Perirhinal Cortex. *Journal of Neuroscience*, *34*(14), 4766–4775.
- Clopath, C. (2012). Synaptic consolidation: An approach to long-term learning. *Cognitive Neurodynamics*, *6*(3), 251–257.
- Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. Cambridge, MA: MIT Press.
- Correia, J. M., Jansma, B. M. B., & Bonte, M. (2015). Decoding Articulatory Features

- from fMRI Responses in Dorsal Speech Regions. *The Journal of Neuroscience*, 35(45), 15015–15025.
- Cottrell, G. W., & Plunkett, K. (1991). Learning the past tense in a recurrent network: Acquiring the mapping from meaning to sounds. In *13th Annual Conference of the Cognitive Science Society* (pp. 328–333). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Coutanche, M. N., Gianessi, C. A., Chanales, A. J., Willison, K. W., & Thompson-Schill, S. L. (2013). The role of sleep in forming a memory representation of a two-dimensional space. *Hippocampus*, 23(12), 1189–1197.
- Cox, D. D., & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage*, 19(2), 261–270.
- Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology. General*, 132(2), 163–201.
- Creery, J. D., Oudiette, D., Antony, J. W., & Paller, K. A. (2015). Targeted Memory Reactivation during Sleep Depends on Prior Learning. *Sleep*, 38(5), 755–763.
- Cusack, R., Brett, M., & Osswald, K. (2003). An Evaluation of the Use of Magnetic Field Maps to Undistort Echo-Planar Images. *NeuroImage*, 18(1), 127–142.
- Cusack, R., Vicente-grabovetsky, A., Mitchell, D. J., Wild, C. J., Auer, T., Linke, A., & Peelle, J. E. (2015). Automatic analysis (aa): efficient neuroimaging workflows and parallel processing using Matlab and XML. *Frontiers in Neuroinformatics*, 8(90), 1–13.
- Daugherty, K., & Hare, M. (1993). What’s in a rule? The past tense by some other name might be called a connectionist net. In *Proceedings of the 1993 connectionist models summer school* (pp. 149–156).
- Daugherty, K., & Seidenberg, M. S. (1992). Rules or connections? The past tense revisited. In *14th Annual Meeting of the Cognitive Science Society* (pp. 259–264). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology*, 16(6), 693–700.
- Davatzikos, C., Ruparel, K., Fan, Y., Shen, D. G., Acharyya, M., W., L. J., ... Langleben, D. D. (2005). Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *NeuroImage*, 28(3), 663–668.
- Davis, M. H., Di Betta, A. M., Macdonald, M. J. E., & Gaskell, M. G. (2009). Learning and consolidation of novel spoken words. *Journal of Cognitive Neuroscience*,

- 21(4), 803–820.
- Davis, M. H., & Gaskell, M. G. (2009). A complementary systems account of word learning: neural and behavioural evidence. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 364(1536), 3773–800.
- Davis, M. H., & Johnsrude, I. S. (2007). Hearing speech sounds: top-down influences on the interface between audition and speech perception. *Hearing Research*, 229(1–2), 132–47.
- Davis, M. H., van Casteren, M., & Marslen-Wilson, W. D. (2003). Frequency Effects in Processing Inflected Dutch Nouns: A Distributed Connectionist Account. *Morphological Structure in Language Processing*, 1–36.
- Davis, T., Love, B. C., & Preston, A. R. (2012). Learning the Exception to the Rule : Model-Based fMRI Reveals Specialized Representations for Surprising Category Members. *Cerebral Cortex*, 22(2), 260–273.
- De Diego Balaguer, R., Rodríguez-Fornells, A., Rotte, M., Bahlmann, J., Heinze, H. J., & Münte, T. F. (2006). Neural circuits subserving the retrieval of stems and grammatical features in regular and irregular verbs. *Human Brain Mapping*, 27(11), 874–888.
- Desai, R., Conant, L. L., Waldron, E., & Binder, J. R. (2006). fMRI of Past Tense Processing: The Effects of Phonological Complexity and Task Difficulty. *Journal of Cognitive Neuroscience*, 18(2), 278–297.
- Devereux, B. J., Clarke, A., Marouchos, A., & Tyler, L. K. (2013). Representational similarity analysis reveals commonalities and differences in the semantic processing of words and objects. *The Journal of Neuroscience*, 33(48), 18906–18916.
- Dhond, R. P., Marinkovic, K., Dale, A. M., Witzel, T., & Halgren, E. (2003). Spatiotemporal maps of past-tense verb inflection. *NeuroImage*, 19(1), 91–100.
- Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends in Cognitive Sciences*, 11(9), 379–386.
- Diedrichsen, J., Provost, S., & Zareamoghaddam, H. (2016). On the distribution of cross-validated Mahalanobis distances. *arXiv Preprint arXiv:1607.01371*, 1–24. Retrieved from <http://arxiv.org/abs/1607.01371>
- Diedrichsen, J., Wiestler, T., & Krakauer, J. W. (2013). Two distinct ipsilateral cortical representations for individuated finger movements. *Cerebral Cortex*, 23(6), 1362–1377.
- Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews. Neuroscience*, 11(2), 114–26.

- Dobbins, I. G., Foley, H., Schacter, D. L., & Wagner, A. D. (2002). Executive control during episodic retrieval: multiple prefrontal processes subserve source memory. *Neuron, 35*(5), 989–996.
- Drosopoulos, S., Schulze, C., Fischer, S., & Born, J. (2007). Sleep 's Function in the Spontaneous Recovery and Consolidation of Memories. *Journal of Experimental Psychology. General, 136*(2), 169–183.
- Dudai, Y. (2002). *Memory from A to Z: Keywords, Concepts, and Beyond*. Oxford: Oxford University Press.
- Dudai, Y. (2004). The Neurobiology of Consolidations, Or, How Stable is the Engram? *Annual Review of Psychology, 55*(1), 51–86.
- Dudai, Y., Karni, A., & Born, J. (2015). The Consolidation and Transformation of Memory. *Neuron, 88*(1), 20–32.
- Dudai, Y., & Morris, R. G. (2000). To consolidate or not to consolidate: what are the questions? In J. J. Bulhuis (Ed.), *Brain, Perception, Memory. Advances in Cognitive Sciences*. (Oxford Uni, pp. 149–162). Oxford.
- Dumay, N., & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words. *Psychological Science, 18*(1), 35–39.
- Dumay, N., & Gaskell, M. G. (2012). Overnight lexical consolidation revealed by speech segmentation. *Cognition, 123*(1), 119–132.
- Durrant, S. J., Taylor, C., Cairney, S., & Lewis, P. A. (2011). Sleep-dependent consolidation of statistical learning. *Neuropsychologia, 49*(5), 1322–1331.
- Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences, 94*(13), 7109–7114.
- Earle, F. S., & Myers, E. B. (2015). Sleep and native language interference affect non-native speech sound learning. *Journal of Experimental Psychology: Human Perception and Performance, 41*(6), 1680–1695.
- Edelman, S., Grill-Spector, K., Kushnir, T., & Malach, R. (1998). Toward direct visualization of the internal shape representational space by fMRI. *Psychobiology, 26*(4), 309–321.
- Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus, 20*(1), 1–10.
- Eichenbaum, H. (2017). Memory: Organization and Control. *Annual Review of Psychology, 3*(68), 19–45.
- Eichenbaum, H., & Cohen, N. J. (2001). *From conditioning to conscious recollection: Memory systems of the brain*. New York, NY: Oxford University Press.

- Eisner, F., & McQueen, J. M. (2006). Perceptual learning in speech: stability over time. *The Journal of the Acoustical Society of America*, *119*(4), 1950–1953.
- Ekstrand, B. R. (1977). The effect of sleep on human long-term memory. In R. R. Drucker-Colin & J. L. McGaugh (Eds.), *Neurobiology of Sleep and Memory* (pp. 419–438). New York, NY: Academic press.
- Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D., & Walker, M. P. (2007). Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences*, *104*(18), 7723–7728.
- Ellis, N. C. (2002). Frequency effects in language processing: A review with implications for theories of implicit and explicit language acquisition. *Studies in Second Language Acquisition*, *24*(2), 143–188.
- Ellis, N. C. (2009). Optimizing the input: Frequency and sampling in usage-based and form-focused learning. In M. H. Long & C. Doughty (Eds.), *The Handbook of Language Teaching* (pp. 139–158). Chichester, UK: Wiley-Blackwell.
- Ellis, N. C., & Schmidt, R. (1998). Rules or Associations in the Acquisition of Morphology? The Frequency by Regularity Interaction in Human and PDP Learning of Morphosyntax. *Language and Cognitive Processes*, *13*(2–3), 307–336.
- Empson, J. A., & Clarke, P. R. (1970). Rapid eye movements and remembering. *Nature*, *227*(5255), 287–288.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and Exemplars in Category Learning. *Journal of Experimental Psychology: General*, *127*(2), 107–140.
- Ervin, S. M. (1964). Imitation and structural change in children's language. In E. Lenneberg (Ed.), *New Directions in the Study of Language* (pp. 163–189). Cambridge, MA: MIT Press.
- Ervin, S. M., & Miller, W. R. (1963). Language development. In H. W. Stevenson (Ed.), *The sixty-second yearbook of the National Society for the Study of Education, Part I* (pp. 108–143). Chicago: University of Chicago Press.
- Estes, W. K. (1986). Array models for category learning. *Cognitive Psychology*, *18*(4), 500–49.
- Estes, W. K. (1994). *Classification and cognition*. Oxford: University Press.
- Evans, S., & Davis, M. H. (2015). Hierarchical Organization of Auditory and Motor Representations in Speech Perception: Evidence from Searchlight Similarity Analysis. *Cerebral Cortex*, *25*(12), 4772–88.
- Fenn, K. M., Margoliash, D., & Nusbaum, H. C. (2013). Sleep restores loss of generalized but not rote learning of synthetic speech. *Cognition*, *128*(3), 280–286.
- Fenn, K. M., Nusbaum, H. C., & Margoliash, D. (2003). Consolidation during sleep of perceptual learning of spoken language, *425*(October), 614–616.

- Ficca, G., & Salzarulo, P. (2004). What is sleep is for memory. *Sleep Medicine*, 5(3), 225–230.
- Fischer, S., & Born, J. (2009). Anticipated reward enhances offline learning during sleep. *2009*, 35(6), 1586–1593.
- Fischer, S., Drosopoulos, S., Tsen, J., & Born, J. (2006). Implicit learning -- explicit knowing: a role for sleep in memory system interaction. *Journal of Cognitive Neuroscience*, 18(3), 311–319. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16602193>
- Fischer, S., Hallschmid, M., Elsner, A. L., & Born J. (2002). Sleep forms memory for finger skills. *Proceedings of the National Academy of Sciences*, 99(18), 11987–11991.
- Fletcher, P. C., Shallice, T., & Dolan, R. J. (1998). The functional roles of prefrontal cortex in episodic memory. I. Encoding. *Brain*, 121(Pt 7), 1239–1248.
- Fodor, J. A. (1968). *Psychological Explanation: An Introduction to the Philosophy of Psychology*. New York, NY: Random House.
- Fodor, J. A. (1975). *The Language of Thought*. Hassocks, Sussex: Harvester Press.
- Fogel, S. M., Smith, C. T., & Cote, K. A. (2007). Dissociable learning-dependent changes in REM and non-REM sleep in declarative and procedural memory systems. *Behavioural Brain Research*, 180(1), 48–61.
- Ford, M., Marslen-Wilson, W. D., & Davis, M. H. (2003). Morphology and frequency: contrasting methodologies. In R. H. Baayen & R. Schreuder (Eds.), *Morphological Structure in Language Processing* (pp. 89–124). Berlin: Mouton de Gruyter.
- Formisano, E., De Martino, F., Bonte, M., & Goebel, R. (2008). “Who” is saying “what”? Brain-based decoding of human voice and speech. *Science*, 322(5903), 970–973.
- Forster, K. I., & Chambers, S. M. (1973). Lexical access and naming time. *Journal of Verbal Learning and Verbal Behavior*, 12(6), 627–635.
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084), 680–683.
- Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Natural Reviews Neuroscience*, 8(9), 700–711.
- Franks, J. J., & Bransford, J. D. (1971). Abstraction of visual patterns. *Journal of Experimental Psychology*, 90(1), 65–74.
- French, R. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4), 128–135. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10322466>

- Frey, U., & Morris, R. G. (1997). Synaptic tagging and long-term potentiation. *Nature*, 385(6616), 533–6.
- Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences*, 6(2), 78–84.
- Friederici, A. D. (2011). The brain basis of language processing: From structure to function. *Physiological Reviews*, 91(4), 1357–92.
- Friederici, A. D., Meyer, M., & von Cramon, Y. D. (2000). Auditory Language Comprehension: An Event-Related fMRI Study on the Processing of Syntactic and Lexical Information. *Brain and Language*, 74(2), 289–300.
- Friederici, A. D., & von Cramon, Y. D. (2000). Syntax in the brain: Linguistic versus neuroanatomical specificity. *Behavioral and Brain Sciences*, 23(1), 32–33.
- Friston, K. J. (2011). Functional and effective connectivity: a review. *Brain Connectivity*, 1(1), 13–36.
- Friston, K. J., Buechel, C., Fink, G. R., Morris, J., Rolls, E., & Dolan, R. J. (1997). Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage*, 6(3), 218–229.
- Friston, K. J., Jezzard, P., & Turner, R. (1994). Analysis of functional MRI time-series. *Human Brain Mapping*, 1(2), 153–171.
- Fuentemilla, L., Miró, J., Ripollés, P., Vilà-Balló, A., Juncadella, M., Castañer, S., ... Rodríguez-Fornells, A. (2013). Hippocampus-dependent strengthening of targeted memories via reactivation during sleep in humans. *Current Biology*, 23(18), 1769–1775.
- Gais, S., Albouy, G., Boly, M., Dang-Vu, T. T., Darsaud, A., Desseilles, M., ... Peigneux, P. (2007). Sleep transforms the cerebral trace of declarative memories. *Proceedings of the National Academy of Sciences of the United States of America*, 104(47), 18778–18783.
- Gais, S., Mölle, M., Helms, K., & Born, J. (2002). Learning-dependent increases in sleep spindle density. *Journal of Neuroscience*, 22(15), 6830–6834.
- Gais, S., Plihal, W., Wagner, U., & Born, J. (2000). Early sleep triggers memory for early visual discrimination skills. *Nature Neuroscience*, 3(12), 1335–1339.
- Gardiner, J. M., Brandt, K. R., Baddeley, A. D., Vargha-Khadem, F., & Mishkin, M. (2008). Developmental amnesia. *Journal of Contemporary Psychotherapy*, 46(11), 2865–2868.
- Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition*, 89(2), 105–132.
- Gaskell, M. G., & Marslen-Wilson, W. D. (1997). Integrating Form and Meaning: A Distributed Model of Speech Perception. *Language and Cognitive Processes*,

- 12(5–6), 613–656.
- Geranmayeh, F., Wise, R. J. S., Mehta, A., & Leech, R. (2014). Overlapping Networks Engaged during Spoken Language Production and Its Cognitive Control. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 34(26), 8728–40.
- Giordano, B. L., McAdams, S., Zatorre, R. J., Kriegeskorte, N., & Belin, P. (2013). Abstract encoding of auditory objects in cortical activity patterns. *Cerebral Cortex*, 23(9), 2025–2037.
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10), 1222–1223.
- Giuditta, A., Ambrosini, M. V., Montagnese, P., Mandile, P., Cotugno, M., Grassi Zucconi, G., & Vescia, S. (1995). The sequential hypothesis of the function of sleep. *Behavioural Brain Research*, 69(1–2), 157–166.
- Goldman-Rakic, P. S. (1988). Topography of cognition: parallel distributed networks in primate association cortex. *Annual Review of Neuroscience*, 11, 137–156.
- Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science*, 13(5), 431–6. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12219809>
- Gómez, R. L., & Maye, J. (2005). The Developmental Trajectory of Nonadjacent Dependency Learning. *Infancy*, 7(2), 183–206.
- Gor, K., & Chernigovskaya, T. (2003). Mental Lexicon Structure in L1 and L2 Acquisition: Russian Evidence. *The Slavic and East European Language Resource Center*.
- Gor, K., & Chernigovskaya, T. (2003). Generation of Complex Verbal Morphology in First and Second Language Acquisition: Evidence from Russian. *19th Scandinavian Conference of Linguistics*, 6, 819–833.
- Gordon Hayman, C. A., Macdonald, C. A., & Tulving, E. (1993). The role of repetition and associative interference in new semantic learning in amnesia: a case experiment. *Journal of Cognitive Neuroscience*, 5(4), 375–389.
- Grodzinsky, Y., & Friederici, A. D. (2006). Neuroimaging of syntax and syntactic processing. *Current Opinion in Neurobiology*, 16(2), 240–246.
- Grossman, M., Carvell, S., Stern, M. B., Gollomp, S., & Hurtig, H. I. (1992). Sentence comprehension in Parkinson's disease: the role of attention and memory. *Brain and Language*, 42(4), 347–84.
- Gupta, A. S., van der Meer, M. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, 65(5), 695–

- 705.
- Guye, M., Bartolomei, F., & Ranjeva, J. P. (2008). Imaging structural and functional connectivity: towards a unified definition of human brain organization? *Current Opinion in Neurology*, 21(4), 393–403.
- Guyonneau, R., VanRullen, R., & Thorpe, S. J. (2004). Temporal codes and sparse representations: A key to understanding rapid processing in the visual system. *Journal of Physiology-Paris*, 98(4), 487–497.
- Guzowski, J. F., Knierim, J. J., & Moser, E. I. (2004). No TitleEnsemble dynamics of hippocampal regions CA3 and CA1. *Neuron*, 44(4), 581–584.
- Habib, R., Nyberg, L., & Tulving, E. (2003). Hemispheric asymmetries of memory: the HERA model revisited. *Trends in Cognitive Sciences*, 7(6), 241–245.
- Hahn, U., & Nakisa, R. C. (2000). German inflection: single route or dual route? *Cognitive Psychology*, 41(4), 313–60.
- Haist, F., Bowden, G. J., & Mao, H. (2001). Consolidation of human memory over decades revealed by functional magnetic resonance imaging. *Nature Neuroscience Neurosci*, 4(11), 1139–1145.
- Halle, M., & Mohanan, K. P. (1985). segmental phonology of Modern English. *Linguistic Inquiry*, 16(1), 57–116.
- Hanson, S. J., Matsuka, T., & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a “face” area? *NeuroImage*, 23(1), 156–166.
- Hare, M., & Elman, J. (1992). A connectionist account of English inflectional morphology: evidence from language change. In *14th Annual Conference of the Cognitive Science Society* (pp. 265–270). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hare, M., & Elman, J. L. (1995). Learning and morphological change. *Cognition*, 56(1), 61–98. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3756909&tool=pmc&entrez&rendertype=abstract>
- Hare, M., Elman, J. L., & Daugherty, K. G. (1995). Default generalization in connectionist networks. *Language and Cognitive Processes*, 10(6), 601–630. Retrieved from file:///Users/Domingo/Documents/Literature/E_Lit/3023.pdf
- Hare, M., Ford, M., & Marslen-Wilson, W. D. (2001). Ambiguity and frequency effects in regular verb inflection. In J. L. Bybee & P. Hopper (Eds.), *Frequency and the emergence on linguistic structure* (pp. 181–200). Amsterdam: John Benjamins.
- Harrison, L. M., Duggins, A., & Friston, K. J. (2006). Encoding uncertainty in the

- hippocampus. *Neural Networks*, 19(5), 535–546.
- Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, 458(7238), 632–635.
- Hasher, L., Goldstein, D. L., & May, C. (2005). It's about time: Circadian rhythms, memory and aging. In C. Izawa & N. Ohta (Eds.), *Human learning and memory: Advances in theory and application* (pp. 199–218). Mahwah, NJ: Lawrence Erlbaum Associates.
- Hawkins, E. A., & Rastle, K. (2016). How does the provision of semantic information influence the lexicalization of new spoken words? *Quarterly Journal of Experimental Psychology*, 69(7), 1322–1339.
- Haxby, J. V, Connolly, A. C., & Guntupalli, J. S. (2014). Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. *Annual Review of Neuroscience*, 37, 435–456.
- Haxby, J. V, Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430.
- Haynes, J.-D., & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*, 8(5), 686–691.
- Haynes, J.-D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, 7, 523–534.
- Haynes, J.-D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading hidden intentions in the human brain. *Current Biology*, 17(4), 323–328.
- Hebb, D. O. (1949). *The Organization of Behavior*. New York, NY: Wiley.
- Heckers, S., Zalesak, M., Weiss, A. P., Ditman, T., & Titone, D. (2004). Hippocampal activation during transitive inference in humans. *Hippocampus*, 14(2), 153–162.
- Henderson, L., Weighall, A. R., Brown, H., & Gaskell, M. G. (2012). Consolidation of vocabulary is associated with sleep in children. *Developmental Science*, 15(5), 674–687.
- Hickok, G., & Poeppel, D. (2004). Dorsal and ventral streams: a framework for understanding aspects of the functional anatomy of language. *Cognition*, 92(1–2), 67–99.
- Hickok, G., & Poeppel, D. (2007). The cortical organization of speech processing. *Nature Reviews Neuroscience*, 8, 393–402.
- Hinton, G. (1986). Learning distributed representations of concepts. In *Proceedings of the Eight Annual Conference of the Cognitive Science Society* (pp. 1–12).
- Hintzman, D. L. (1986). Schema abstraction in a multiple-trace memory model.

- Psychological Review*, 93(4), 411–428.
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95(4), 528–51.
- Hockett, C. F. (1954). Problems of morphemic analysis. *Language*, 23(4), 321–343.
- Hodges, J. R., & Patterson, K. (1995). Is semantic memory consistently impaired early in the course of Alzheimer's disease? Neuroanatomical and diagnostic implications. *Neuropsychologia*, 33(4), 441–459.
- Homa, D. (1984). On the nature of categories. *Psychology of Learning & Motivation*, 18, 49–94.
- Homa, D., Sterling, S., & Trepel, L. (1981). Evaluation of exemplar-based generalization and the abstraction of categorical information Limitations of Exemplar-Based Generalization and the Abstraction of Categorical Information. *Journal of Experimental Psychology*, 7(6), 418–39.
- Homa, D., & Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning and Memory*, 2(3), 322–330.
- Hooper, J. B. (1976). *Introduction to natural generative phonology*. New York, NY: Academic press.
- Hothorn, T., Bretz, F., & Westfall, P. (2008). Simultaneous Inference in General Parametric Models. *Biometrical Journal*, 50(3), 346–363.
- Howes, D. H., & Solomon, R. L. (1951). Visual duration threshold as a function of word-probability. *Journal of Experimental Psychology*, 41(6), 401–410.
- Hu, P., Stylos-Allan, M., & Walker, M. P. (2006). Sleep facilitates consolidation of emotional declarative memory. *Psychological Science*, 17(10), 891–898.
- Huber, K. M., Kayser, M. S., & Bear, M. F. (2000). Role for rapid dendritic protein synthesis in hippocampal mGluR-dependent long-term depression. *Science*, 288(5469), 1254–1257.
- Huber, R., Ghilardi, M. F., Massimini, M., & Tononi, G. (2004). Local sleep and learning. *Nature*, 430, 78–81.
- Humphreys, G., & Forde, E. (2001). Hierarchies, similarity, and interactivity in object recognition: “Category-specific” neuropsychological deficits. *Behavioral and Brain Sciences*, 24(3), 453–76.
- Humphries, C., Binder, J. R., Medler, D. A., & Liebenthal, E. (2006). Syntactic and semantic modulation of neural activity during auditory sentence comprehension. *Journal of Cognitive Neuroscience*, 18(4), 665–679.
- Indefrey, P. (2011). The Spatial and Temporal Signatures of Word Production

- Components: A Critical Update. *Frontiers in Psychology*, 2(255), 1–16.
- Indefrey, P., Brown, C. M., Hagoort, P., Herzog, H., Sach, M., & Seitz, R. J. (1997). A PET Study of Cerebral Activation Patterns Induced by Verb Inflection. *Neuroimage*, 5, 548.
- Jaeger, J. J., Lockwood, A. H., Kemmerer, D. L., van Valin, R. D., Murphy, B. W., & Khalak, H. G. (1996). A Positron Emission Tomographic Study of Regular and Irregular Verb Morphology in English. *Language*, 72(3), 451–497.
- Jaeger, T. F. (2008). Categorical Data Analysis: Away from ANOVAs (transformation or not) and towards Logit Mixed Models. *Journal of Memory and Language*, 59(4), 434–446.
- Jenkins, J. B., & Dallenbach, K. M. (1924). Oblivescence during sleep and waking. *American Journal of Psychology*, 35, 605–612.
- Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: a connectionist model. *Proceedings of the National Academy of Sciences of the United States of America*, 96(13), 7592–7597.
- Joanisse, M. F., & Seidenberg, M. S. (2005). Imaging the past: neural activation in frontal and temporal regions during regular and irregular past-tense processing. *Cognitive, Affective & Behavioral Neuroscience*, 5(3), 282–296.
- Joseph, J. E. (2001). Functional neuroimaging studies of category specificity in object recognition: a critical review and meta-analysis. *Cognitive, Affective & Behavioral Neuroscience*, 1(2), 119–136.
- Juola, P., & Plunkett, K. (2000). Why double dissociations don't mean much. In G. Cohen, R. A. Johnston, & K. Plunkett (Eds.), *Exploring cognition: Damaged brains and neural networks: Readings in cognitive neuropsychology and connectionist modelling* (pp. 319–327). Hove, UK: Psychology Press.
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8(5), 679–685.
- Kamitani, Y., & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology*, 16(11), 1096–1102.
- Khaligh-Razavi, S. M., & Kriegeskorte, N. (2014). Deep Supervised, but Not Unsupervised, Models May Explain IT Cortical Representation. *PLoS Computational Biology*, 10(11), 1–29.
- Kiani, R., Esteky, H., Mirpour, K., & Tanaka, K. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of Neurophysiology*, 97(6), 4296–4309.
- Kireev, M., Slioussar, N., Korotkov, A. D., Chernigovskaya, T. V., & Medvedev, S. V. (2015). Changes in functional connectivity within the fronto-temporal brain

- network induced by regular and irregular Russian verb production. *Frontiers in Human Neuroscience*, 9(36), 1–10.
- Klimovich-Gray, A., Bozic, M., & Marslen-Wilson, W. D. (2017). Domain-specific and Domain-general Processing in Left Perisylvian Cortex: Evidence from Russian. *Journal of Cognitive Neuroscience*, 29(2), 382–397.
- Kloosterman, F., van Haften, T., & Lopes da Silva, F. H. (2004). Two reentrant pathways in the hippocampal-entorhinal system. *Hippocampus*, 14(8), 1026–39.
- Knierim, J. J., Lee, I., & Hargreaves, E. L. (2006). Hippocampal place cells: Parallel input streams, subregional processing, and implications for episodic memory. *Hippocampus*, 16(9), 755–764.
- Korman, M., Doyon, J., Doljansky, J., Carrier, J., Dagan, Y., & Karni, A. (2007). Daytime sleep condenses the time course of motor memory consolidation. *Nature Neuroscience*, 10(9), 1206–1213.
- Kotz, S. A., Cappa, S. F., von Cramon, D. Y., & Friederici, A. D. (2002). Modulation of the Lexical–Semantic Network by Auditory Semantic Priming: An Event-Related Functional MRI Study. *NeuroImage*, 17(4), 1761–1772.
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, 17(8), 401–412.
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2(4), 1–28.
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... Bandettini, P. A. (2008). Matching Categorical Object Representations in Inferior Temporal Cortex of Man and Monkey. *Neuron*, 60(6), 1126–1141.
- Kruskal, J. B., & Wish, M. (1978). *Multidimensional Scaling*. Beverly Hills, CA: Sage University Series.
- Kuczaj, S. A. I. (1977). The acquisition of regular and irregular past tense forms. *Journal of Verbal Learning and Verbal Behavior*, 16(5), 589–600.
- Kumaran, D. (2012). What representations and computations underpin the contribution of the hippocampus to generalization and inference? *Frontiers in Human Neuroscience*, 6(157), 1–11.
- Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated. *Trends in Cognitive Sciences*, 20(7), 512–534.
- Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: a model of the hippocampal system.

- Psychological Review*, 119(3), 573–616.
- Kumaran, D., Summerfield, J. J., Hassabis, D., & Maguire, E. A. (2009). Tracking the Emergence of Conceptual Knowledge during Human Decision Making. *Neuron*, 63(6), 889–901.
- Lachter, J., & Bever, T. G. (1988). The relation between linguistic structure and associative theories of language learning — A constructive critique of some connectionist learning models. *Cognition*, 28(1–2), 195–247.
- Lacy, J. W., Yassa, M. A., Stark, S. M., Muftuler, L. T., & Stark, C. E. L. (2011). Distinct pattern separation related transfer functions in human CA3/dentate and CA1 revealed using high-resolution fMRI and variable mnemonic similarity. *Learning & Memory*, 18(1), 15–18.
- Lahl, O., Wispel, C., Willigens, B., & Pietrowsky, R. (2008). An ultra short episode of sleep is sufficient to promote declarative memory performance. *Journal of Sleep Research*, 17(1), 3–10.
- Laine, M., Niemi, J., Koivuselkä-Sallinen, P., & Hyönä, J. (1995). Morphological processing of polymorphemic nouns in a highly inflecting language. *Journal of Cognitive Neuropsychology*, 12(5), 457–502.
- Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review*, 107(2), 227–60.
- Lambon Ralph, M. A., Pobric, G., & Jefferies, E. (2009). Conceptual knowledge is underpinned by the temporal pole bilaterally: convergent evidence from rTMS. *Cerebral Cortex*, 19(4), 832–838.
- LaRocque, K. F., Smith, M. E., Carr, V. A., Witthoft, N., Grill-Spector, K., & Wagner, A. D. (2013). Global Similarity and Pattern Separation in the Human Medial Temporal Lobe Predict Subsequent Memory. *Journal of Neuroscience*, 33(13), 5466–5474.
- Lau, H., Alger, S. E., & Fishbein, W. (2011). Relational memory: A daytime nap facilitates the abstraction of general concepts. *PLoS ONE*, 6(11), 1–6.
- Lavric, A., Pizzagalli, D., Forstmeier, S., & Rippon, G. (2001). A double-dissociation of English past-tense production revealed by event-related potentials and low-resolution electromagnetic tomography (LORETA). *Clinical Neurophysiology*, 112(10), 1833–1849.
- Leach, L., & Samuel, A. G. (2007). Lexical Configuration and Lexical Engagement: When Adults Learn New Words. *Cognitive Psychology*, 55(4), 306–353.
- Lehtonen, M., & Laine, M. (2003). How word frequency affects morphological processing in monolinguals and bilinguals. *Bilingualism: Language and Cognition*, 6(3), 213–225.

- Leibold, C., & Kempter, R. (2008). Sparseness constrains the prolongation of memory lifetime via synaptic metaplasticity. *Cerebral Cortex*, *18*(1), 67–77.
- Leminen, A., Kimppa, L., Leminen, M. M., Lehtonen, M., Mäkelä, J. P., & Shtyrov, Y. (2016). Acquisition and consolidation of novel morphology in human neocortex: A neuromagnetic study. *Cortex*, *83*, 1–16.
- Lepage, M., Habib, R., & Tulving, E. (1998). Hippocampal PET activations of memory encoding and retrieval: the HIPER model. *Hippocampus*, *8*(4), 313–322.
- Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science*, *305*(5688), 1295–1298.
- Lewis, P. A., & Durrant, S. J. (2011). Overlapping memory replay during sleep builds cognitive schemata. *Trends in Cognitive Sciences*, *15*(8), 343–351.
- Li, P. (2002). Log linear models for the analysis of language acquisition data. *Journal of Cognitive Science*.
- Ling, X., & Marinov, M. (1993). Answering the connectionist challenge: a symbolic model of learning the past tense of English verbs. *Cognition*, *49*(3), 235–290.
- Liu, Y., Lin, W., Liu, C., Luo, Y., Bayley, P. J., Wu, J., & Qin, S. (2016). Memory consolidation reconfigures neural pathways involved in the suppression of emotional memories. *Nature Communications*, *7*, 1–12.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A Network Model of Category Learning. *111*(2), 309–332.
- MacWhinney, B., & Leinbach, J. (1991). Implementations are not conceptualizations: Revising the verb learning model. *Cognition*, *40*(1), 121–157.
- Mahon, B. Z., & Caramazza, A. (2009). Concepts and categories: a cognitive neuropsychological perspective. *Annual Review of Psychology*, *60*, 27–51.
- Maitra, R., Roys, S. R., & Gullapalli, R. P. (2002). Test-retest reliability estimation of functional MRI data. *Magnetic Resonance in Medicine*, *48*(1), 62–70.
- Manns, J. R., Hopkins, R. O., Reed, J. M., Kitchener, E. G., & Squire, L. R. (2003). Recognition memory and the human hippocampus. *Neuron*, *37*(1), 171–180.
- Marchman, V., Plunkett, K., & Goodman, J. (1997). Overregularization in English plural and past tense inflectional morphology: a response to Marcus (1995). *Journal of Child Language*, *24*(3), 767–779.
- Marcus, G. (1995). The Acquisition of the English Past Tense in Children and Multilayered Connectionist Networks. *Cognition*, *56*(3), 271–279.
- Marcus, G. (2001). *The algebraic mind: integrating connectionism and cognitive science*. Cambridge, MA: MIT Press.

- Marcus, G., Brinkman, U., Clahsen, H., Wiese, R., & Pinker, S. (1995). German inflection: the exception that proves the rule. *Cognitive Psychology*, 29(3), 189–256.
- Marr, D. (1970). A Theory for Cerebral Neocortex. *Proceedings of the Royal Society. B, Biological Sciences*, 176(1043), 161–234.
- Marr, D. (1971). Simple Memory: A Theory for Archicortex. *Philosophical Transactions. Biological Sciences*, 262(841), 23–81.
- Marschner, I. (2011). glm2: Fitting generalized linear models with convergence problems. *The R Journal*, 3(2), 12–15.
- Marshall, L., Helgadóttir, H., Mölle, M., & Born J. (2006). Boosting slow oscillations during sleep potentiates memory. *Nature*, 444, 610–613.
- Marslen-Wilson, W. D. (1987). Functional parallelism in spoken word-recognition. *Cognition*, 25(1–2), 71–102.
- Marslen-Wilson, W. D., & Teuber, H. L. (1975). Memory for remote events in anterograde amnesia: Recognition of public figures from newsphotographs. *Neuropsychologia*, 13(3), 353–364.
- Marslen-Wilson, W. D., & Tyler, L. K. (1997). Dissociating types of mental computation. *Nature*, 387(6633), 593.
- Marslen-Wilson, W. D., & Tyler, L. K. (1998). Rules, representations, and the English past tense. *Trends in Cognitive Sciences*, 2(11), 428–35. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/21227274>
- Marslen-Wilson, W. D., & Tyler, L. K. (2003). Capturing underlying differentiation in the human language system. *Trends in Cognitive Sciences*, 7(2), 61–62.
- Marslen-Wilson, W. D., & Tyler, L. K. (2005). The lexicon , grammar , and the past tense : Dissociation revisited. In M. Tomasello & D. Slobin (Eds.), *Beyond Nature-Nurture: Essays in Honor of Elizabeth Bates* (pp. 263–280). Mahwah, NJ: Lawrence Erlbaum.
- Marslen-Wilson, W. D., & Tyler, L. K. (2007a). Morphology, language and the brain: the decompositional substrate for language comprehension. *Philosophical Transactions of the Royal Society of London*, 362(1481), 823–836.
- Marslen-Wilson, W. D., & Tyler, L. K. (2007b). Morphology , language and the brain : the decompositional substrate for language comprehension, (March), 823–836.
- Martin, A. (2007). The Representation of Object Concepts in the Brain. *Annual Review of Psychology*, 58(1), 25–45.
- Martin, K. C., Casadio, A., Zhu, H., Yaping, E., Rose, J. C., Chen, M., ... Kandel, E. R. (1997). Synapse-specific, long-term facilitation of aplysia sensory to motor synapses: A function for local protein synthesis in memory storage. *Cell*, 91(7),

- 927–938.
- Mayes, A., Montaldi, D., & Migo, E. (2007). Associative memory and the medial temporal lobes. *Trends in Cognitive Sciences*, *11*(3), 126–135.
- McCarthy, J., & Prince, A. (1990). Prosodic morphology and templatic morphology. In *Presepectives on Arabic Linguistics II: Papers from the Second Annual Symposium on Arabic Linguistics* (pp. 1–55).
- McClelland, J. L. (2013). Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, *142*(4), 1190–210.
- McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, *14*(8), 348–356.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–57. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/7624455>
- McClelland, J. L., & Patterson, K. (2002a). “ Words or Rules ” cannot exploit the regularity in exceptions Reply to Pinker and Ullman. *Trends in Cognitive Sciences*, *6*(11), 464–465.
- McClelland, J. L., & Patterson, K. (2002b). Rules or connections in past-tense inflections: what does the evidence rule out? *Trends in Cognitive Sciences*, *6*(11), 465–472. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12457897>
- McClelland, J. L., & Patterson, K. (2002c). Rules or connections in past-tense inflections: What does the evidence rule out? *Trends in Cognitive Sciences*, *6*(11), 465–472.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *Psychology of Learning and Motivation*, *24*, 109–165.
- Mcgonigle, D. J., Howseman, A. M., Athwal, B. S., Friston, K. J., Frackowiak, R. S. J., & Holmes, A. P. (2000). Variability in fMRI : An Examination of Intersession Differences. *NeuroImage*, *11*, 708–734.
- McLaren, D. G., Ries, M. L., Xu, G., & Johnson, S. C. (2012). A generalized form of context-dependent psychophysiological interactions (gPPI): a comparison to standard approaches. *NeuroImage*, *61*(4), 1277–1286.
- McNaughton, B. L. (2010). Cortical hierarchies, sleep, and the extraction of

- knowledge from memory. *Artificial Intelligence*, 174(2), 205–214.
- Mcnaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, 10(10), 408–415.
- Medin, D. L., & Schaffer, M. M. (1978). Context Theory of Classification Learning, 85(3), 207–238.
- Mednick, S., Nakayama, K., & Stickgold, R. (2003). Sleep-dependent learning: a nap is as good as a night. *Nature Neuroscience*, 6(7), 697–698.
- Mensink, G., & Raaijmakers, J. G. (1988). A model for interference and forgetting. *Psychological Review*, 95(4), 434–455.
- Merks, M., Rastle, K., & Davis, M. H. (2011). The acquisition of morphological knowledge investigated through artificial language learning. *Quarterly Journal of Experimental Psychology* (2006), 64(6), 1200–20.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24, 167–202.
- Miller, G. A., & Selfridge, J. A. (1950). Verbal context and the recall of meaningful material. *American Journal of Psychology*, 63(2), 176–185.
- Minda, J. P., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 27(3), 775–99.
- Minda, J. P., & Smith, J. D. (2002). Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology*, 28(2), 275–92.
- Mirkovic, J., & Gaskell, M. G. (2016). Does sleep improve your grammar? Preferential consolidation of arbitrary components of new linguistic knowledge. *PLoS ONE*, 11(4), 1–26.
- Mirković, J., Seidenberg, M. S., & Joanisse, M. F. (2011). Rules Versus Statistics: Insights From a Highly Inflected Language. *Cognitive Science*, 35(4), 638–681.
- Misaki, M., Kim, Y., Bandettini, P. A., & Kriegeskorte, N. (2010). Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *NeuroImage*, 53(1), 103–118.
- Mishkin, M., Malamut, B., & Bachevalier, J. (1984). Memories and habits: Two neural systems. In G. Lynch, J. L. McGaugh, & N. M. Weinberger (Eds.), *Neurobiology of learning and memory* (pp. 65–77). New York, NY: Guildford Press.
- Mitchell, T. M., Hutchinson, R., Niculescu, R. S., Pereira, F., & Wang, X. (2004). Learning to decode cognitive states from brain images. *Machine Learning*, 57(1–2), 145–175.

- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76(2), 165–178.
- Moscovitch, M. (1992). Memory and Working-with-Memory: A Component Process Model Based on Modules and Central Systems. *Journal of Cognitive Neuroscience*, 4(3), 257–267.
- Moscovitch, M. (1995). Recovered consciousness: a hypothesis concerning modularity and episodic memory. *Journal of Clinical and Experimental Neuropsychology*, 17(2), 276–291.
- Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., ... Nadel, L. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *Journal of Anatomy*, 207(1), 35–66.
- Müller, G. E. & Pilzecker, A. (1900). Experimentelle Beiträge zur Lehre vom Gedächtnis. *Z. Psychol. Ergänzungsband*, 1, 1–300.
- Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Murphy, K., Birn, R. M., Handwerker, D. A., Jones, T. B., & Bandettini, P. A. (2009). The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *NeuroImage*, 44(3), 893–905.
- Nadel, L., & Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, 7(2), 217–227.
- Nakashiba, T., Buhl, D. L., McHugh, T. J., & Tonegawa, S. (2009). Hippocampal CA3 output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron*, 62(6), 781–787.
- Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI. *NeuroImage*, 56(2), 400–410.
- Newell, A. (1980). Physical Symbol Systems. *Cognitive Science*, 4(2), 135–183.
- Newell, A., & Simon, H. A. (1981). Computer science as empirical inquiry: Symbol and search. In J. Haugeland (Ed.), *Mind design: Philosophy psychology artificial intelligence* (pp. 35–66). Cambridge, MA: MIT Press.
- Newman, S. D., & Twieg, D. (2001). Differences in auditory processing of words and pseudowords: an fMRI study. *Human Brain Mapping*, 14(1), 39–47.
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W. D., & Kriegeskorte, N. (2014). A Toolbox for Representational Similarity Analysis. *PLoS Computational Biology*, 10(4).
- Nishida, M., Pearsall, J., Buckner, R. L., & Walker, M. P. (2009). REM sleep, prefrontal theta, and the consolidation of human emotional memory. *Cerebral Cortex*, 19(5), 1158–1166.

- Nokia, M. S., Mikkonen, J. E., Penttonen, M., & Wikgren, J. (2012). Disrupting neural activity related to awake-state sharp wave-ripple complexes prevents hippocampal learning. *Frontiers in Behavioral Neuroscience*, 6(84), 1–10.
- Noll, D. C., Genovese, C. R., Nystrom, L. E., Vazquez, A. L., Forman, S. D., Eddy, W. F., & Cohen, J. D. (1997). Estimating test-retest reliability in functional MR imaging. II: Application to motor and cognitive activation studies. *Magnetic Resonance in Medicine*, 38(3), 508–517.
- Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychological Review*, 110(4), 611–646.
- Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading : multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10(9), 424–430.
- Nosofsky, R. M. (1986). Attention, Similarity, and the Identification-Categorization Relationship. *Journal of Experimental Psychology*, 115(1), 39–57.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14(4), 700–708.
- Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. *Cognitive Psychology*, 23(1), 94–140.
- Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of “ multiple-system ” phenomena in perceptual categorization, 7(3), 375–402.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*.
- Nosofsky, R. M., & Zaki, S. R. (1998). Dissociations between Categorization and Recognition in Amnesic and Normal Individuals: An Exemplar-Based Interpretation. *Psychological Science*, 9(4), 247–55.
- O'Brien, F., & Cousineau, D. (2014). Representing Error bars in within-subject designs in typical software packages. *The Quantitative Methods for Psychology*, 10(1), 56–67.
- O'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Jozsef, C. (2010). Play it again: reactivation of waking experience and memory. *Trends in Neurosciences*, 33(5), 220–229.
- O'Reilly, J. X., Woolrich, M. W., Behrens, T. E. J., Smith, S. M., & Johansen-Berg, H. (2012). Tools of the trade: Psychophysiological interactions and functional connectivity. *Social Cognitive and Affective Neuroscience*, 7(5), 604–609.
- O'Reilly, R. C., Bhattacharyya, R., Howard, M. D., & Ketz, N. (2014). Complementary

- learning systems. *Cognitive Science*, 38(6), 1229–1248.
- O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, 4(6), 661–682.
- O'Reilly, R. C., & Norman, K. a. (2002). Hippocampal and neocortical contributions to memory: advances in the complementary learning systems framework. *Trends in Cognitive Sciences*, 6(12), 505–510. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12475710>
- O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive Representations in Learning and Memory: Principles of Cortical and Hippocampal Function. *Psychological Review*, 108(1), 311–345.
- O'Toole, A. J., Jiang, F., Abdi, H., & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, 17(4), 580–90.
- Onnis, L., Christiansen, M., Chater, N., & Gomez, R. (2003). Reduction of uncertainty in human sequential learning: Evidence from artificial language learning. In *Proceedings of The 25th Annual Conference of the Cognitive Science Society* (pp. 810–815).
- Onnis, L., Monaghan, P., Christiansen, M. H., & Chater, N. (2004). Variability is the spice of learning, and a crucial ingredient for detecting and generalizing in nonadjacent dependencies. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*.
- Orsolini, M., Fanari, R., & Bowles, H. (1998). Acquiring Regular and Irregular Inflection in a Language with Verb Classes. *Language & Cognitive Processes*, 13(4), 425–464.
- Orsolini, M., & Marslen-Wilson, W. D. (1997). Universals in morphological representation: Evidence from Italian. *Language & Cognitive Processes*, 12(1–47).
- Oudiette, D., & Paller, K. A. (2013). Upgrading the sleeping brain with targeted memory reactivation. *Trends in Cognitive Sciences*, 17(3), 142–149.
- Paller, K. A., & Voss, J. L. (2004). Memory reactivation and consolidation during sleep. *Learning & Memory*, 11(6), 664–670.
- Palmeri, T. J., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(3), 548–568.
- Patterson, K., Lambon Ralph, M. A., Hodges, J. R., & McClelland, J. L. (2001). Deficits in irregular past-tense verb morphology associated with degraded semantic knowledge. *Neuropsychologia*, 39(7), 709–724.
- Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you

- know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12), 976–987.
- Paz, R., Bauer, E. P., & Pare, D. (2007). Learning-Related Facilitation of Rhinal Interactions by Medial Prefrontal Inputs. *Journal of Neuroscience*, 27(24), 6542–6551.
- Peigneux, P., Laureys, S., Fuchs, S., Collette, F., Perrin, F., Reggers, J., ... Maquet, P. (2004). Are Spatial Memories Strengthened in the Human Hippocampus during Slow Wave Sleep? *Neuron*, 44(3), 535–545.
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: a tutorial overview. *NeuroImage*, 45(1 Suppl), S199–S209.
- Petrides, M., & Pandya, D. N. (1988). Association fiber pathways to the frontal cortex from the superior temporal region in the rhesus monkey. *The Journal of Comparative Neurology*, 273(1), 52–66.
- Pinker, S. (1991). Rules of language. *Science*, 253(5019), 530–535.
- Pinker, S. (1999). *Words and Rules: the ingredients of language*. New York, NY: HarperCollins.
- Pinker, S. (2001). Four Decades of Rules and Associations, or Whatever Happened to the Past Tense. In E. Dupoux (Ed.), *Language, Brain, and Cognitive Development* (pp. 157–179). Cambridge, MA: MIT Press.
- Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28(1–2), 73–193.
- Pinker, S., & Prince, A. (1994). Regular and Irregular Morphology and the Psychological Status of Rules of Grammar. In S. D. Lima, R. L. Corrigan, & G. K. Iverson (Eds.), *The Reality of Linguistic Rules* (pp. 353–388). Amsterdam: John Benjamins.
- Pinker, S., & Ullman, M. (2002a). Combination and structure, not gradedness, is the issue. Reply to McClelland and Patterson. *Trends in Cognitive Sciences*, 6(11), 472–474. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12457898>
- Pinker, S., & Ullman, M. T. (2002b). The past and future of the past tense. *Trends in Cognitive Sciences*, 6(11), 456–463.
- Plaut, D. C. (1995). Double Dissociation Without Modularity: Evidence from Connectionist Neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, 17(2), 291–321.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: computational principles in quasi-regular domains. *Psychological Review*, 103(1), 56–115.

- Plihal, W., & Born, J. (1997). Effects of early and late nocturnal sleep on declarative and procedural memory. *Journal of Cognitive Neuroscience*, 9(4), 534–547.
- Plihal, W., & Born, J. (1999). Effects of early and late nocturnal sleep on priming and spatial memory. *Psychophysiology*, 36(5), 571–582.
- Plunkett, K., & Juola, P. (1999). A Connectionist Model of English Past Tense and Plural Morphology. *Cognitive Science*, 23(4), 463–490.
- Plunkett, K., & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition*, 38(1), 43–102.
- Plunkett, K., & Marchman, V. (1993). From rote learning to system building: acquiring verb morphology in children and connectionist nets. *Cognition*, 48(1), 21–69.
- Plunkett, K., & Marchman, V. A. (1996). Learning from a connectionist model of the acquisition of the English past tense. *Cognition*.
- Plunkett, K., & Nakisa, R. C. (1997). A Connectionist Model of the Arabic Plural System A Connectionist Model of the Arabic Plural System. *Language and Cognitive Processes*, 12(5/6), 807–836.
- Poldrack, R. A. (2011). Inferring mental states from neuroimaging data: From reverse inference to large-scale decoding. *Neuron*, 72(5), 692–697.
- Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, 310(5756), 1963–1966.
- Poppenk, J., Evensmoen, H. R., Moscovitch, M., & Nadel, L. (2013). Long-axis specialization of the human hippocampus. *Trends in Cognitive Sciences*, 17(5), 230–240.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3), 353–363.
- Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, 83(2), 304–308.
- Power, J. D., Schlaggar, D. L., & Petersen, S. E. (2014). Studying brain organization via spontaneous fMRI signal. *Neuron*, 84(4), 681–696.
- Prasada, S., & Pinker, S. (1993). Generalisation of regular and irregular morphological patterns. *Language and Cognitive Processes*, 8(1), 1–56.
- Prasada, S., Pinker, S., & Snyder, W. (1990). Some evidence that irregular forms are retrieved from memory but regular forms are rule generated. In *31st annual meeting of the Psychonomic Society*.

- Preston, A. R., & Eichenbaum, H. (2013). Interplay of Hippocampus and Prefrontal Cortex in Memory. *Current Biology*, 23(17), 764–773.
- Preston, A. R., Shrager, Y., Dudukovic, N. M., & Gabrieli, J. D. (2004). Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus*, 14(2), 148–152.
- Pylyshyn, Z. (1984). *Computation and cognition: toward a foundation for cognitive science* (MIT Press). Cambridge, MA.
- Ragnarsdóttir, H., Simonsen, H. G., & Plunkett, K. (1999). Acquisition of past tense inflection in Icelandic and Norwegian children: an experimental study. *Journal of Child Language*, 26(3), 577–618.
- Rasch, B., Büchel, C., Gais, S., & Born, J. (2007). Odor cues during slow-wave sleep prompt declarative memory consolidation. *Science*, 315(5817), 1426–1429.
- Rastle, K., Harrington, J., & Coltheart, M. (2002). 358,534 nonwords: the ARC Nonword Database. *Quarterly Journal of Experimental Psychology*, 55(4), 1339–1362.
- Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, 97(2), 285–308.
- Rauchs, G., Bertran, F., Guillery-Girard, B., Desgranges, B., Kerrouche, N., Denise, P., ... Eustache, F. (2004). Consolidation of strictly episodic memories mainly requires rapid eye movement sleep. *Sleep*, 27(3), 395–401.
- Rauchs, G., Desgranges, B., Foret, J., & Eustache, F. (2005). The relationships between memory systems and sleep stages. *Journal of Sleep Research*, 14(2), 123–140.
- Rauchs, G., Feyers, D., Landeau, B., Bastin, C., Luxen, A., Maquet, P., & Collette, F. (2011). Sleep contributes to the strengthening of some memories over others, depending on hippocampal activity at learning. *Journal of Neuroscience*, 31(7), 2563–2568.
- Redondo, R. L., & Morris, R. G. M. (2011). Making memories last: the synaptic tagging and capture hypothesis. *Nature Reviews. Neuroscience*, 12(1), 17–30.
- Reed, K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3(3), 382–407.
- Rempel-Clower, N. L., Zola, S. M., Squire, L. R., & Amaral, D. G. (1996). Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. *The Journal of Neuroscience*, 16(16), 5233–5255.
Retrieved from
<http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=8756452&retmode=ref&cmd=prlinks%5Cnpapers3://publication/uuid/A1CBCD08->

E0C6-4325-9467-CB65A039B7E2

- Reymann, K. G., & Frey, J. U. (2007). The late maintenance of hippocampal LTP: requirements, phases, “synaptic tagging”, “late-associativity” and implications. *Neuropharmacology*, *52*(1), 24–40.
- Ribot, T. (1882). *Diseases of the Memory: An Essay in the Positive Psychology*. New York, NY: D. Appleton and Company.
- Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, *306*(5695), 443–447.
- Rihm, J. S., Diekelmann, S., Born, J., & Rasch, B. (2014). Reactivating memories during sleep by odors: odor specificity and associated changes in sleep oscillations. *Journal of Cognitive Neuroscience*, *26*(8), 1806–1818.
- Ritchey, M., Wing, E. A., Labar, K. S., & Cabeza, R. (2012). Neural Similarity Between Encoding and Retrieval is Related to Memory Via Hippocampal Interactions. *Cerebral Cortex*, *23*(12), 2818–2828.
- Rogalsky, C., & Hickok, G. (2008). Selective attention to semantic and syntactic features modulates sentence processing networks in anterior temporal cortex. *Cerebral Cortex*, *19*(4), 786–796.
- Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and Deterioration of Semantic Memory : A Neuropsychological and Computational Investigation. *Psychological Review*, *111*(1), 205–235.
- Rogers, T. T., & McClelland, J. L. (2003). *Semantic Cognition : A Parallel Distributed Processing Approach*. Cambridge, MA: MIT Press.
- Rolls, E. T., Treves, A., & Tovee, M. J. (1997). The representational capacity of the distributed encoding of information provided by populations of neurons in primate temporal visual cortex. *Experimental Brain Research*, *114*(1), 149–162.
- Rombouts, S. A. R. B., Barkhof, F., Hoogenraad, F. G. C., Sprenger, M., Valk, J., & Scheltens, P. (1997). Test-Retest Analysis with Functional MR of the Activated Area in the Human Visual Cortex. *American Journal of Neuroradiology*, *18*(7), 1317–1322.
- Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 111–144). New York, NY: Academic press.
- Rosch, E. (1977). Human categorization. In N. Warren (Ed.), *Studies in cross-cultural psychology* (pp. 3–49). London: Academic press.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*(4), 573–605.

- Rosenbaum, R. S., Winocur, G., & Moscovitch, M. (2001). New views on old memories: Re-evaluating the role of the hippocampal complex. *Behavioural Brain Research*, *127*(1–2), 183–197.
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, *65*(6), 386–408.
- Rosenblatt, F. (1962). *Principles of neurodynamics; perceptrons and the theory of brain mechanisms*. Washington: Spartan Books.
- Rotello, C. M., Macmillan, N. A., & Reeder, J. A. (2004). Sum-difference theory of remembering and knowing: A two-dimensional signal detection model. *Psychological Review*, *111*(3), 588–616.
- Roth, D. A., Kishon-Rabin, L., Hildesheimer, M., & Karni, A. (2005). A latent consolidation phase in auditory identification learning: time in the awake state is sufficient. *Learning & Memory*, *12*(2), 159–164.
- Rudoy, J. D., Voss, J. L., Westerberg, C. E., & Paller, K. A. (2009). Strengthening individual memories by reactivating them during sleep. *Science*, *326*(5956), 1079.
- Rumelhart, D. E. (1990). Brain style computation: Learning and generalization. In S. F. Zornetzer, L. Davis, & C. Lau (Eds.), *An introduction to neural and electronic networks*. San Diego, CA: Academic press.
- Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A General Framework for Parallel Distributed Processing. In *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I* (pp. 45–76). Cambridge, MA: MIT Press.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*(6088), 533–536.
- Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 216–271). Cambridge, MA: MIT Press.
- Sach, M., Seitz, R. J., & Indefrey, P. (2004). Unified inflectional processing of regular and irregular verbs: A PET study. *NeuroReport*, *15*(3), 533–537.
- Sajikumar, S., Navakkode, S., Sacktor, T. C., & Frey, J. U. (2005). Synaptic tagging and cross-tagging: the role of protein kinase Mzeta in maintaining long-term potentiation but not long-term depression. *Journal of Neuroscience*, *25*(24), 5750–5756.
- Saletin, J. M., Goldstein, A. N., & Walker, M. P. (2011). The role of sleep in directed forgetting and remembering of human memories. *Cerebral Cortex*, *21*(11), 2534–

- 2541.
- Sandrini, M., Cappa, S. F., Rossi, S., Rossini, P. M., & Miniussi, C. (2003). The role of prefrontal cortex in verbal episodic memory: rTMS evidence. *Journal of Cognitive Neuroscience*, *15*(6), 855–861.
- Schacter, D. L., & Tulving, E. (1994). *Memory systems*. (D. L. Schacter & E. Tulving, Eds.). Cambridge, MA: MIT Press.
- Schacter, D. L., & Wagner, A. D. (1999). Medial temporal lobe activations in fMRI and PET studies of episodic encoding and retrieval. *Hippocampus*, *9*(1), 7–24.
- Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*(17), 1622–1627.
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711), 20160049.
- Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., & Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, *26*(1), 3–8.
- Schiller, D., Eichenbaum, H., Buffalo, E. A., Davachi, L., Foster, D. J., Leutgeb, S., & Ranganath, C. (2015). Memory and Space: Towards an Understanding of the Cognitive Map. *Journal of Neuroscience*, *35*(41), 13904–13911.
- Schlichting, M. L., Mumford, J. A., & Preston, A. R. (2015). Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nature Communications*, *6*, 8151.
- Schreiner, T., & Rasch, B. (2015). Boosting vocabulary learning by verbal cueing during sleep. *Cerebral Cortex*, *25*(11), 4169–4179.
- Schreiner, T., & Rasch, B. (2017). The beneficial role of memory reactivation for language learning during sleep: A review. *Brain and Language*, *167*, 94–105.
- Schurger, A., Pereira, F., Treisman, A., & Cohen, J. D. (2010). Reproducibility distinguishes conscious from nonconscious neural representations. *Science*, *327*(5961), 97–99.
- Scott, S. K. (2005). Auditory processing — speech , space and auditory objects. *Current Opinion in Neurobiology*, *15*(2), 197–201.
- Scott, S. K., & Johnsrude, I. S. (2003). The neuroanatomical and functional organization of speech perception. *Trends in Neurosciences*, *26*(2), 100–107.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral

- hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, 20(1), 11–21.
- Seger, C. A. (2008). How do the basal ganglia contribute to categorization? Their role in generalization, response selection, and learning via feedback. *Neuroscience and Biobehavioral Reviews*, 32(2), 265–278.
- Seger, C. A., & Miller, E. K. (2010). Category Learning in the Brain. *Annual Review of Neuroscience*, 33, 203–219.
- Seger, C. A., & Peterson, E. J. (2013). Categorization = decision making + generalization. *Neuroscience and Biobehavioral Reviews*, 37(7), 1187–1200.
- Seidenberg, M. S. (1992). Connectionism without tears. In S. Davis (Ed.), *Connectionism: Theory and practice* (pp. 84–137). New York, NY: Oxford University Press.
- Seidenberg, M. S., & Arnoldussen, A. (2003). The brain makes a distinction between hard and easy stimuli: Comments on Beretta et al. *Brain and Language*, 85(3), 527–530.
- Seidenberg, M. S., & Bruck, M. (1990). Consistency effects in the generation of past the morphology. In *31st annual meeting of the Psychonomic Society*. New Orleans, LA.
- Seidenberg, M. S., & Hoeffner, J. H. (1998). Evaluating Behavioral and Neuroimaging Data on past Tense Processing. *Language*, 74(1), 104–122.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96(4), 523–56.
- Sejnowski, T. J., & Rosenberg, C. R. (1987). Parallel networks that learn to pronounce English text. *Complex Systems*, 1, 145–168.
- Sereno, J. A., & Jongman, A. (1997). Processing of English inflectional morphology. *Memory & Cognition*, 25(4), 425–437.
- Shepard, R. N. (1980). Multidimensional scaling, tree-fitting, and clustering. *Science*, 210(4468), 390–398.
- Shepard, R. N., & Chipman, S. (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology*, 1(1), 1–17.
- Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: hippocampal-midbrain encoding of overlapping events. *Neuron*, 60(2), 378–389.
- Simons, J. S., & Spiers, H. J. (2003). Prefrontal and medial temporal lobe interactions in long-term memory. *Nature Reviews. Neuroscience*, 4, 637–648.
- Simonsen, H. G. (2000). Past tense acquisition and processing in Norwegian: Experimental evidence. *Language and Language Behavior*, 3(2), 86–101.

- Sirota, A., Csicvari, J., Buhl, D., & Buzsaki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences*, *100*(4), 2065–2069.
- Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, *271*(5257), 1870–1873.
- Smith, C. (2001). Sleep states and memory processes in humans: procedural versus declarative memory systems. *Sleep Medicine Reviews*, *5*(6), 491–506.
- Smith, E. E., & Medin, D. L. (1981). *Categories and Concepts*. Harvard: University Press.
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the Mist : The Early Epochs of Category Learning, *24*(6), 1411–36.
- Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *26*(1), 3–27.
- Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., ... Consortium, W.-M. H. (2013). Resting-state fMRI in the Human Connectome Project. *NeuroImage*, *15*(80), 144–168.
- Solomon, R. L., & Postman, L. (1952). Frequency of usage as a determinant of recognition thresholds for words. *Journal of Experimental Psychology*, *43*, 195–201.
- Soon, C. S., Brass, M., Heinze, H. J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11*(5), 543–545.
- Soveri, A., Lehtonen, M., & Laine, M. J. (2007). Word frequency and morphological processing in Finnish revisited. *The Mental Lexicon*, *2*(3), 359 –385.
- Spaniol, J., Davidson, P. S., Kim, A. S., Han, H., Moscovitch, M., & Grady, C. L. (2009). Event-related fMRI studies of episodic encoding and retrieval: meta-analyses using activation likelihood estimation. *Neuropsychologia*, *47*(8–9), 1765–1779.
- Spiridon, M., & Kanwisher, N. (2002). How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron*, *35*(6), 1157–1165.
- Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*(2), 195–231.
- Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current Opinion in Neurobiology*, *5*(2), 169–177.
- Squire, L. R., & Bayley, P. J. (2007). The neuroscience of remote memory. *Current Opinion in Neurobiology*, *17*(2), 185–196.

- Squire, L. R., Clark, R. E., & Knowlton, B. J. (2001). Retrograde amnesia. *Hippocampus*, *11*, 50–55.
- Squire, L. R., & Knowlton, B. (2000). The medial temporal lobe, the hippocampus, and the memory systems of the brain. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 765–780). Cambridge, MA: MIT Press.
- Squire, L. R., Stark, C. E. L., & Clark, R. E. (2004). The medial temporal lobe. *Annual Review of Neuroscience*, *27*, 279–306.
- Staeren, N., Renvall, H., De Martino, F., Goebel, R., & Formisano, E. (2009). Sound categories are represented as distributed patterns in the human auditory cortex. *Current Biology*, *19*(6), 498–502.
- Stamatakis, E. A., Marslen-Wilson, W. D., Tyler, L. K., & Fletcher, P. C. (2005). Cingulate control of fronto-temporal integration reflects linguistic demands: A three-way interaction in functional connectivity. *NeuroImage*, *28*(1), 115–121.
- Staresina, B. P., Henson, R. N., Kriegeskorte, N., & Alink, A. (2012). Episodic reinstatement in the medial temporal lobe. *The Journal of Neuroscience*, *32*(50), 18150–18156.
- Stemberger, J. P., & MacWhinney, B. (1986). Form-oriented inflectional errors in language processing. *Cognitive Psychology*, *18*(3), 329–354.
- Sterpenich, V., Albouy, G., Darsaud, A., Schmidt, C., Vandewalle, G., Dang-Vu, T. T., ... Maquet, P. (2009). Sleep Promotes the Neural Reorganization of Remote Emotional Memory. *The Journal of Neuroscience*, *29*(16), 5143–5152.
- Stickgold, R., James, L., & Hobson, J. a. (2000). Visual discrimination learning requires sleep after training. *Nature Neuroscience*, *3*(12), 1237–1238.
- Stickgold, R., & Walker, M. P. (2013). Sleep-dependent memory triage: evolving generalization through selective processing. *Nature Neuroscience*, *16*(2), 139–45.
- Stickgold, R., Whidbee, D., Schirmer, B., Patel, V., & Hobson, J. A. (2000). Visual discrimination task improvement: A multi-step process occurring during sleep. *Journal of Cognitive Neuroscience*, *12*(2), 246–254.
- Stowe, L. A., Broere, C. A. J., Paans, A. M. J., Wijers, A. A., Mulder, G., Vaalburg, W., & Zwarts, F. (1998). Localizing components of a complex task : Sentence processing and working memory. *NeuroReport*, *9*(13), 2995–2999.
- Stowe, L. A., Haverkort, M., & Zwarts, F. (2005). Rethinking the neurological basis of language. *Lingua*, *115*(7), 997–1042.
- Strange, B. A., Duggins, A., Penny, W., Dolan, R. J., & Friston, K. J. (2005). Information theory, novelty and hippocampal responses: Unpredicted or unpredictable? *Neural Networks*, *18*(3), 225–230.
- Takashima, A., Bakker, I., van Hell, J. G., Janzen, G., & McQueen, J. M. (2014).

- Richness of information about novel words influences how episodic and semantic memory networks interact during lexicalization. *NeuroImage*, 84, 265–278.
- Takashima, A., Bakker, I., van Hell, J. G., Janzen, G., & McQueen, J. M. (2016). Interaction between episodic and semantic memory networks in the acquisition and consolidation of novel spoken words. *Brain and Language*, 167, 44–60.
- Takashima, A., Nieuwenhuis, I. L. C., Jensen, O., Talamini, L. M., Rijpkema, M., & Fernández, G. (2009). Shift from hippocampal to neocortical centered retrieval network with consolidation. *The Journal of Neuroscience*, 29(32), 10087–93.
- Tambini, A., Ketz, N., & Davachi, L. (2010). Enhanced brain correlations during rest are related to memory for recent experiences. *Neuron*, 65(2), 280–290.
- Tamminen, J., Davis, M. H., Merckx, M., & Rastle, K. (2012). The role of memory consolidation in generalisation of new linguistic information. *Cognition*, 125(1), 107–112.
- Tamminen, J., Davis, M. H., & Rastle, K. (2015). From specific examples to general knowledge in language learning. *Cognitive Psychology*, 79, 1–39.
- Tamminen, J., & Gaskell, M. G. (2008). Newly learned spoken words show long-term lexical competition effects. *The Quarterly Journal of Experimental Psychology*, 61(3), 361–371.
- Tamminen, J., Payne, J. D., Stickgold, R., Wamsley, E. J., & Gaskell, M. G. (2010). Sleep spindle activity is associated with the integration of new memories and existing knowledge. *Journal of Neuroscience*, 30(43), 14356–14360.
- Thompson-Schill, S. L., Bedny, M., & Goldberg, R. F. (2005). The frontal lobes and the regulation of mental activity. *Current Opinion in Neurobiology*, 15(2), 219–224.
- Tilley, A. J., & Empson, J. A. (1978). REM sleep and memory consolidation. *Biological Psychology*, 6(4), 293–300.
- Tong, F., & Pratte, M. S. (2012). Decoding patterns of human brain activity. *Annual Review*, 63, 483–509.
- Tononi, G., & Cirelli, C. (2003). Sleep and synaptic homeostasis: A hypothesis. *Brain Research Bulletin*, 62(2), 143–150.
- Tononi, G., & Cirelli, C. (2006). Sleep function and synaptic homeostasis. *Sleep Medicine Reviews*, 10(1), 49–62.
- Torgerson, W. S. (1958). *Theory and methods of scaling*. New York, NY: Wiley.
- Touretzky, D. S. (1987). Representing Conceptual Structures in a Neural Network. In *Proceedings of the IEEE First International Conference on Neural Networks*.
- Touretzky, D. S., & Geva, S. (1987). A distributed connectionist representation for

- concept structures. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*.
- Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, 2(2), 189–199.
- Treves, A., Tashiro, A., Witter, M. P., & Moser, E. I. (2008). What is the mammalian dentate gyrus good for? *Neuroscience*, 154(4), 1155–72.
- Tsao, D. Y., Freiwald, W. A., Knutsen, T. A., Mandeville, J. B., & Tootell, R. B. H. (2003). Faces and objects in macaque cerebral cortex. *Nature Neuroscience*, 6(9), 989–995.
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., ... Morris, R. G. M. (2007). Schemas and Memory Consolidation. *Science*, 316(7489), 76–82.
- Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., ... Morris, R. G. M. (2011). Schema-Dependent Gene Activation. *Science*, 333(6044), 891–895.
- Tucker, M. A., Hirota, Y., Wamsley, E. J., Lau, H., Chaklader, A., & Fishbein, W. (2006). A daytime nap containing solely non-REM sleep enhances declarative but not procedural memory. *Neurobiology of Learning and Memory*, 86(2), 241–247.
- Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.). New York, NY: Academic press.
- Tulving, E. (1985). Memory and consciousness. *Canadian Psychology*, 26(1), 1–12.
- Tulving, E., Kapur, S., Craik, F. I., Moscovitch, M., & Houle, S. (1994). Hemispheric encoding/retrieval asymmetry in episodic memory: positron emission tomography findings. *Proceedings of the National Academy of Sciences*, 91(6), 2016–2020.
- Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, 21(10), 1934–1945.
- Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience*, 30(33), 1177–1187.
- Tyler, L. K., Cheung, T. P. L., Devereux, B. J., & Clarke, A. (2013). Syntactic computations in the language network: Characterizing dynamic network properties using representational similarity analysis. *Frontiers in Psychology*, 4(271), 1–19.
- Tyler, L. K., DeMornay-Davies, P., Anokhina, R., Longworth, C., Randall, B., & Marslen-Wilson, W. D. (2002). Dissociations in Processing Past Tense

- Morphology: Neuropathology and Behavioral Studies. *Journal of Cognitive Neuroscience*, 14(1), 79–94.
- Tyler, L. K., Marslen-Wilson, W. D., & Stamatakis, E. A. (2005). Differentiating lexical form, meaning, and structure in the neural language system. *Proceedings of the National Academy of Sciences*, 102(23), 8375–8380.
- Tyler, L. K., Randall, B., & Marslen-Wilson, W. D. (2002). Phonology and neuropsychology of the English past tense. *Neuropsychologia*, 40(8), 1154–1166.
- Tyler, L. K., Shafto, M. A., Randall, B., Wright, P., Marslen-Wilson, W. D., & Stamatakis, E. A. (2010). Preserving syntactic processing across the adult life span: The modulation of the frontotemporal language system in the context of age-related atrophy. *Cerebral Cortex*, 20(2), 352–364.
- Tyler, L. K., Stamatakis, E. A., Post, B., Randall, B., & Marslen-Wilson, W. D. (2005). Temporal and frontal systems in speech comprehension: An fMRI study of past tense processing. *Neuropsychologia*, 43(13), 1963–1974.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., ... Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1), 273–289.
- Ullman, M. T. (1998). A role for declarative and procedural memory in language. *Brain and Cognition*, 37(1), 142–143.
- Ullman, M. T. (1999a). Acceptability Ratings of Regular and Irregular Past-tense Forms: Evidence for a Dual-system Model of Language from Word Frequency and Phonological Neighbourhood Effects. *Language and Cognitive Processes*, 14(1), 47–67.
- Ullman, M. T. (1999b). The functional neuroanatomy of inflectional morphology. *Behavioral and Brain Sciences*, 22(6), 1041–1042.
- Ullman, M. T. (2001). A neurocognitive perspective on language: the declarative/procedural model. *Nature Reviews Neuroscience*, 2(10), 717–726.
- Ullman, M. T. (2004). Contributions of memory circuits to language: The declarative/procedural model. *Cognition*, 92(1–2), 231–270.
- Ullman, M. T. (2006). The declarative/procedural model and the shallow structure hypothesis. *Applied Psycholinguistics*, 27(1), 97–105.
- Ullman, M. T., Bergida, R., & O'Craven, K. (1997). Distinct fMRI activation patterns for regular and irregular past tense. *NeuroImage*, 5, 549.
- Ullman, M. T., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Koroshetz, W. J., & Pinker, S. (1997). A Neural Dissociation within Language: Evidence that the Mental Dictionary Is Part of Declarative Memory, and that Grammatical Rules

- Are Processed by the Procedural System. *Journal of Cognitive Neuroscience*, 9(2), 266–76.
- van Dongen, E. V., Thielen, J. W., Takashima, A., Barth, M., & Fernández, G. (2012). Sleep supports selective retention of associative memories based on relevance for future utilization. *PLoS ONE*, 7(8), 1–6.
- van Kesteren, M. T. R., Fernández, G., Norris, D. G., & Hermans, E. J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 107(16), 7550–7555.
- van Kesteren, M. T. R., Rijpkema, M., Ruitter, D. J., & Fernández, G. (2010). Retrieval of associative information congruent with prior knowledge is related to increased medial prefrontal activity and connectivity. *Proceedings of the National Academy of Sciences*, 107(16), 7550–7555.
- Vandenberghe, R., Nobre, A. C., & Price, C. J. (2002). The Response of Left Temporal Cortex to Sentences. *Journal of Cognitive Neuroscience*, 14(4), 550–560.
- Vandierendonck, A. (1995). A parallel rule activation and rule synthesis model for generalization in category learning. *Psychonomic Bulletin & Review*, 2(4), 442–459.
- Vogt, B. A., Rosene, D. L., & Pandya, D. N. (1979). Thalamic and cortical afferents differentiate anterior from posterior cingulate cortex in the monkey. *Science*, 204(4389), 205–207.
- Wagner, U., Gais, S., & Born, J. (2001). Emotional memory formation is enhanced across sleep intervals with high amounts of rapid eye movement sleep. *Learning & Memory*, 8(2), 112–119.
- Waldvogel, D., van Gelderen, P., Immisch, I., Pfeiffer, C., & Hallett, M. (2000). The variability of serial fMRI data: correlation between a visual and a motor task. *NeuroReport*, 11(17), 3843–3847.
- Walker, M. P., Brakefield, T., Hobson, J. A., & Stickgold, R. (2003). Dissociable stages of human memory consolidation and reconsolidation. *Nature*, 425(6958), 616–620.
- Walker, M. P., & Stickgold, R. (2010). Overnight alchemy: sleep-dependent memory evolution. *Nature Reviews Neuroscience*, 11(3), 218.
- Whitfield-Gabrieli, S., & Nieto-Castanon, A. (2012). Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain Connectivity*, 2(3), 125–141.
- Wiestler, T., Mcgonigle, D. J., & Diedrichsen, J. (2011). Integration of sensory and motor representations of single fingers in the human cerebellum Integration of

- sensory and motor representations of single fingers in the human cerebellum. *Journal of Neurophysiology*, 105(6), 3042–3053.
- Wikenheiser, A. M., & Redish, A. D. (2015). Hippocampal theta sequences reflect current goals. *Nature Neuroscience*, 18(2), 1–8.
- Wilhelm, I., Diekelmann, S., Molzow, I., Ayoub, A., Mölle, M., & Born, J. (2011). Sleep selectively enhances memory expected to be of future relevance. *The Journal of Neuroscience*, 31(5), 1563–1569.
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172), 676–679.
- Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, 17(5), 766–780.
- Wu, X., & Foster, D. J. (2014). Hippocampal replay captures the unique topological structure of a novel environment. *Journal of Neuroscience*, 34(19), 6459–6469.
- Xue, G., Dong, Q., Chen, C., Lu, Z., Mumford, J. A., & Poldrack, R. A. (2010). Greater neural pattern similarity across repetitions is associated with better memory. *Science*, 330(6000), 97–101.
- Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 111(23), 8619–24.
- Yonelinas, A. P. (2002). The Nature of Recollection and Familiarity: A Review of 30 Years of Research. *Journal of Memory and Language*, 46(3), 441–517.
- Zaki, S. R., Nosofsky, R. M., Stanton, R. D., & Cohen, A. L. (2003). Prototype and Exemplar Accounts of Category Learning and Attentional Allocation : A Reassessment, 29(6), 1160–1173.
- Zalesak, M., & Heckers, S. (2009). The role of the hippocampus in transitive inference. *Psychiatry Research*, 172(1), 24–30.
- Zeithamova, D., Schlichting, M. L., & Preston, A. R. (2012). The hippocampus and inferential reasoning: building memories to navigate future decisions. *Frontiers in Human Neuroscience*, 6(70), 1–14.