

BRIEF ARTICLE

The role of memory consolidation in generalisation of new linguistic information

Jakke Tamminen^a, Matthew H. Davis^b, Marjolein Merkx^a, & Kathleen Rastle^a

^aDepartment of Psychology, Royal Holloway, University of London, Egham, TW20 0EX,
United Kingdom

^bMedical Research Council Cognition and Brain Sciences Unit, 15 Chaucer Road,
Cambridge, CB2 7EF, United Kingdom

Author note

Jakke Tamminen, jakke.tamminen@rhul.ac.uk; Matthew H. Davis, matt.davis@mrc-cbu.cam.ac.uk; Marjolein Merkx, marjoleinmerkx@hotmail.com; Kathleen Rastle, kathy.rastle@rhul.ac.uk.

Address correspondence to: Jakke Tamminen, Department of Psychology, Royal Holloway, University of London, Egham, TW20 0EX, United Kingdom. E-mail: jakke.tamminen@gmail.com. Tel: +44 (0)1784 414390, Fax: +44 (0) 1784 437520.

Abstract

Accounts of memory that postulate complementary learning systems (CLS) have become increasingly influential in the field of language learning. These accounts predict that generalisation of newly learnt linguistic information to untrained contexts requires offline memory consolidation. Such generalisation should not be observed immediately after training, as these accounts claim unconsolidated representations are context and hippocampus-dependent and gain contextual and hippocampal independence only after consolidation. We trained participants on new affixes (e.g., *-nule*) attached to familiar word stems (e.g., *buildnule*), testing them immediately or two days later. Participants showed an immediate advantage for trained affixes in a speeded shadowing task as long as these affixes occurred in the stem contexts in which they were learnt (e.g., *buildnule*). This learning effect generalised to words with untrained stems (e.g., *sailnule*) only in the delayed test condition. By contrast, a non-speeded definition selection task showed immediate generalisation. We propose that generalisation can be supported by initial context-dependent memories given sufficient processing time, but that context-independent lexical representations emerge only following consolidation, as predicted by CLS accounts.

Keywords: Language learning, generalisation, morphology, memory consolidation

The role of memory consolidation in generalisation of new linguistic information

1. Introduction

Language is a combinatorial system whereby information is transmitted by combining a limited set of units to represent a limitless spectrum of ideas. One illustration of this productivity arises in morphological systems. In English we combine a limited set of stem morphemes (e.g., *kind*) with a small number of prefixes (e.g., *unkind*) and suffixes (e.g., *kindness*) to form the vast majority of our existing and new words (Algeo, 1991). The key property of this kind of combinatorial system is that individual units can be arranged in entirely novel combinations that nonetheless convey clear meanings. For example, the affix *-able* can be attached to nearly any verb to form a noun that can be instantly understood by any user of English, even if they have never encountered the affix attached to that particular stem before (e.g., *tweetable*). It is in this sense that knowledge of the affix *-able* is *context-independent*; understanding of the affix is not limited to the contexts in which it has occurred in the past (the stem provides the context here). This ability to generalise familiar linguistic units to new contexts lies at the heart of the communicative power of language, and as such it is important to understand how people learn these units in such a way that they support generalisation. The present work examines the acquisition of new affixes as a means to explore this broader problem.

One solution to the problem of generalisation is offered by accounts of memory that propose complementary learning systems (CLS; e.g. McClelland, McNaughton, & O'Reilly, 1995). While the computational and neural implementations of these accounts differ to varying degrees, they all postulate two dissociable mechanisms that permit information encountered in a limited range of contexts to be generalised to novel contexts (e.g., Marr, 1971; Alvarez & Squire, 1994; Meeter & Murre, 2005; Winocur & Moscovitch, 2011). The

most prominent of these models (McClelland et al., 1995; O'Reilly & Norman, 2002) suggests that neocortical learning specialises in discovering context-independent properties of incoming information, using overlapping distributed representations that encode shared structure across different experiences to support generalisation. However, acquiring new overlapping representations in such a system requires particular computational mechanisms if new knowledge is not to displace old representations (Catastrophic Interference, cf. French, 1999). Therefore a second system in the hippocampus learns context-dependent, non-overlapping representations of new items that are immune to catastrophic interference. These structures enter into dialogue during offline periods such as sleep, when hippocampally-driven replay of new memories allows the gradual strengthening of neocortical memory traces, such that new context-independent neocortical representations develop (Frankland & Bontempi, 2005). Though CLS accounts have recently been used to account for the process of integrating novel words into the mental lexicon (e.g., Davis & Gaskell, 2009), here we use these theories as a framework to explore the problem of generalisation. Critically, CLS theories propose that while *context-dependent* information about new linguistic representations may be learnt rapidly and stored immediately, the *context-independent* representations required to support generalisation to novel contexts will develop only following offline memory consolidation.

However, this prediction appears to be contradicted by recent research on adult language learning, in which participants trained to read aloud words printed in an artificial orthography can immediately generalise learnt spelling-to-sound correspondences to untrained words (Taylor, Plunkett, & Nation, 2011). Though these data appear problematic for CLS accounts, it is possible that context-dependent representations might under some circumstances support generalisation. One possible mechanism is proposed in episodic models of memory (e.g., Hintzman, 1986, 1988). When a novel stimulus is presented,

multiple, context-dependent memory traces are retrieved and averaged to generate an abstract representation sufficient for generalisation. One key difference between these accounts and the CLS proposal is that generalisation in these episodic models is achieved by processes enacted at the time of testing. Such processes may take time and thus may not be available during speeded online language processing. Therefore it might be critical that the generalisation task used by Taylor et al. allowed participants unlimited time to respond, and hence, time to retrieve and combine episodic memory traces. Thus a key test of the CLS predictions about generalisation and consolidation is to use tasks which require rapid, online language processing, rather than non-speeded tasks in which there is time for more elaborate or explicit reasoning processes to support generalisation. Our aim was to assess generalisation using tests that require both speeded and unspeeded processing of novel linguistic combinations.

We trained participants on a set of novel affixes (e.g., *-nule*; Merckx, Rastle, & Davis, 2011) embedded in novel words that participants were instructed to learn along with their meanings (e.g., a *climbnule* is a person who climbs mountains with dangerous peaks; a *buildnule* is someone able to build furniture with remarkable speed). Our goal was to discover whether (and when) participants show knowledge of the new affixes sufficient to generalise to the online recognition of untrained exemplars (e.g., *sailnule*). We used a speeded shadowing task to measure online processing of words with embedded novel affixes. This task does not require metalinguistic judgments, can be performed on trained and untrained items alike, and response latencies are sensitive to lexical variables (Bates & Liu, 1996), making it ideal to probe the nature of learnt lexical representations. The power of this paradigm derives from the fact that during testing novel affixes can be presented in either trained or untrained contexts by manipulating the stem to which the affix is attached: the trained affixes could occur in the context of their trained stems (e.g., *climbnule*) or in the

context of new stems (e.g., *sailnule*), and were tested immediately or two days after training. By comparing shadowing of these stimuli to matched control stimuli with untrained novel affixes (e.g., *floathoke*, *griphoke*), we derived measures of context-dependent and context-independent affix learning before and after consolidation. Context-dependent knowledge is shown by enhanced performance for trained vs. untrained affixes when presented with trained stems. Context-independent affix knowledge is shown by enhanced performance for trained vs. untrained affixes when presented with untrained stems (i.e. outside of the learnt context).

Based on CLS accounts, we expected to see evidence in shadowing of the emergence of context-dependent affix representations immediately, and of context-independent affix representations only after consolidation. We also tested if our paradigm allowed generalisation in non-speeded tasks as shown by Taylor et al. (2011). Merx et al. (2011) reported that three days after training, participants could make non-speeded forced choice decisions about the meanings of untrained words with trained affixes (e.g., *sailnule*). We used the same task to investigate whether this generalisation occurs immediately after training. Finally, we expected to see reliable explicit recognition of trained stem-affix pairings immediately after training, as recognition can be supported by episodic memory, but hypothesised that performance would decline over time as the episodic memory trace decays. To this end, we used a recognition memory task to measure participants' memory for trained affixes, trained stems, and the contexts in which each affix was trained. The latter was accomplished by contrasting recognition accuracy for trained words (e.g., *climbnule*, *birdhalk*) and 'recombinant' words (novel combinations of trained stems and affixes, e.g., *birdnule* or *climbhalk*).

2. Method

2.1 Participants

Thirty-five paid native English speakers participated in the immediate test condition, and 34 in the delayed test condition.

2.2 Materials and Design

Sixteen novel affixes (Table 1) were divided into two lists (Merkx et al., 2011). Each participant was trained on the eight affixes from one list, and those from the other list served as untrained controls. The allocation of lists was counterbalanced. Each affix was trained with eight different monosyllabic verb or noun stems, making a total of 64 novel words to be learnt. Word definitions were composed so that each affix had a consistent meaning, similar to one of four existing English affix meanings (Table 1).

2.3 Procedure

In training, a novel word was simultaneously presented visually on screen and auditorily over headphones. The definition was presented below the word. After studying the word and definition, participants typed the novel word. Each novel word was presented nine times, interleaved with three single blocks of an active recall task in which a definition was presented on screen, and participants typed the corresponding novel word (see Merkx et al., 2011, Experiment 2).

Table 1. Examples of trained affixes and stems, their associated meanings, and untrained affixes in one counterbalancing list.

Affix	Examples of trained novel words (S+A+) and associated meanings
-nule	Climbnule is someone who climbs mountains with dangerous peaks. Buildnule is someone who is able to build furniture at a remarkable speed.
-ane	Lockane is the bank section containing the mechanism used to lock the vault. Bringane is the waiting room used for people who bring the queen presents.
-lomb	Knitlomb is a tool used to knit crossover patterns into woollen cloth. Pourolomb is a bottle cap used to pour exact measures of a liquor.
-esh	Creepesh is the price of buying stealth equipment used to creep noiselessly. Wrapesh is the extra cost of getting a shop assistant to wrap presents.
-tege	Whiptege is a leatherworker who has designed a new type of horse whip. Graintege is the person who buys the grain needed to produce chicken feed.
-ose	Crewose is a device used to measure the rum ration for sailing crew. Bombose is a delicate tool used to help defuse different types of bomb.
-halk	Birdhalk is a populated area where a rare bird has built a nest. Meathalk is the place on an exploration ship where dried meat was stored.
-uck	Vanuck is the tax paid for importing a van from the United States. Gunuck is the fine for illegal possession of a gun in Canada.

Note: The second list of affixes (untrained in this example) consists of -nept, -tund, -ort, -aph, -labe, -hoke, -ude, -ete. The affix meanings always denote a person, a place, a tool, or a cost.

The test phase took place either immediately or two days after training and consisted of three tasks in fixed order (shadowing, recognition memory, definition selection) using novel words that were learnt in the training session or never encountered before (untrained). The untrained items were created by combining trained stems (S+) with untrained affixes (A-), and untrained stems (S-) with trained (A+) or untrained affixes (A-). The conditions in each task are presented in Table 2. The shadowing task required participants to repeat aloud a spoken stimulus as quickly as possible. In the recognition memory task a word was presented visually and participants indicated if it was a trained or an untrained novel word. In addition

to the conditions used in shadowing, this task also included the recombinant condition (S_x+A_y+ , S_y+A_x+ , following training on S_x+A_x+ and S_y+A_y+). Finally, the definition selection task required participants to choose between visually presented target and foil definitions for S+A+ and S-A+ novel words in a two-alternative forced choice task. The target definitions were the same as those presented during training (for S+A+ words), or novel definitions consistent with the trained meaning of the affix (for S-A+ words). The foil definitions combined the meaning of the stem with the meaning of a different trained affix.

3. Results

Shadowing data were analysed by fitting linear mixed-effects models to log-transformed RTs (Table 2 and Figure 1 show retransformed RTs) of accurate responses (Baayen, Davidson, & Bates, 2008) with items and subjects as random factors. Fixed effects were centered and we report the t-statistic associated with the coefficient of each effect. A significant three-way interaction involving time-of-testing (immediate vs. delayed), context (training context vs. new context), and learning (trained affix vs. untrained affix), $t=1.97$, $p=.046$, justified analysing the data separately for the two contexts. For affixes presented in the context in which they had been trained, we observed a main effect of learning, $t=7.26$, $p<.001$, and no interaction with time-of-testing ($p=.60$), showing that there was a processing advantage for trained affixes both initially after learning and two days later. For affixes presented in a new, untrained context, a significant effect of learning, $t=2.17$, $p=.03$, was modulated by an interaction with time-of-testing, $t=1.92$, $p=.047$. This interaction reflected a significant effect of learning only in the delayed test, $t=3.25$, $p<.001$ (immediate test, $t=0.14$, $p=.90$). Critically, the three-way interaction is precisely as predicted by the CLS models: context-independent affix knowledge is not only absent immediately after training, but also

significantly enhanced by a delay between training and testing despite context-dependent affix knowledge being indistinguishable in both immediate and delayed testing¹.

Table 2. Conditions, number of items in each condition, and summary statistics (mean RT in shadowing, mean percent correct in recognition memory and definition selection) in the three test tasks. Standard error in parentheses.

Task and item type	Time of test	
	Immediate	Delayed
Shadowing		
Trained stem, trained affix (S _x +A _x +, N = 32)	1095 (±31) ms	1112 (±23) ms
Trained stem, untrained affix (S+A-, N = 32)	1129 (±33) ms	1143 (±26) ms
Untrained stem, trained affix (S-A+, N = 32)	1143 (±32) ms	1147 (±23) ms
Untrained stem, untrained affix (S-A-, N = 32)	1143 (±31) ms	1163 (±24) ms
Recognition memory		
Trained stem, trained affix (S _x +A _x +, N = 64)	86 (±1.6)%	88 (±1.4)%
Trained stem, untrained affix (S+A-, N = 32)	96 (±1.5)%	96 (±1.4)%
Untrained stem, trained affix (S-A+, N = 32)	98 (±0.9)%	95 (±1.5)%
Recombinant words (S _x +A _v +, N = 64)	81 (±2.2)%	67 (±3.5)%
Definition selection		
Trained stem, trained affix (S _x +A _x +, N = 64)	96 (±1.5)%	96 (±1.1)%
Untrained stem, trained affix (S-A+, N = 64)	86 (±2.9)%	81 (±3.4)%

Note: S+ = trained stem, S- = untrained stem, A+ = trained affix, A- = untrained affix. Subscript markings express the idea that in recombinant words stems and affixes are trained but derived from different trained novel words.

We analysed the recognition memory task using signal detection measures (d' , Snodgrass & Corwin, 1988) to control for response biases (Figure 1). Stem recognition was assessed by comparing z-transformed rates of correct “yes” responses to trained items (hits)

¹ An anonymous reviewer pointed out that CLS models might also predict that participants become faster over time in the S-A+ condition. Our between-subjects design provides only limited opportunity to observe this effect since the group of participants in the delayed condition are approximately 20 ms slower than those in the immediate group (based on performance in the untrained S-A- or trained S+A+ conditions). An exception to this general slowing is seen in the critical S-A+ condition in which RTs are virtually unchanged for the immediate and delayed groups. These comparisons are implicit in factorial analyses using linear mixed effects: the presence of a significant three-way interaction confirms that the generalised learning effect is absent in the group tested immediately and present in the group tested after a delay. A within-subject replication might provide a further opportunity to show speeding of S-A+ responses, though with the caveat that repeated testing permits other processes (e.g. repetition priming) to modify response times.

and incorrect “yes” responses to S-A+ items (false alarms). Affix recognition was computed similarly using false alarms to S+A- items, and whole-word recognition reflected the difference between hits and false alarms to recombinant words. These d' scores were analysed using analysis of variance (ANOVA). In all conditions, d' was significantly ($p < .05$) higher than zero, indicating above chance recognition. A 3x2 ANOVA with knowledge type (stem, affix, whole-word) and time-of-testing as factors showed a main effect of knowledge type, $F(2, 134) = 298.86$, $p < .001$, and an interaction between knowledge type and time-of-testing, $F(2, 134) = 5.47$, $p = .005$. Figure 1 suggests that this interaction is due to a larger time-of-testing effect in whole-word recognition than in the other two knowledge types, which showed little difference. This was confirmed by a significant interaction in a 2x2 ANOVA collapsed over stem/affix recognition, $F(1, 67) = 8.60$, $p = .005$. Whereas recognition memory for stems and affixes was the same for learners tested immediately or two days after learning, (both $p > .5$), there was a marginally significant difference in knowledge of the trained stem and affix combinations (whole-word) between the immediate and delayed test conditions, $t(67) = 1.75$, $p = .08$. Thus, episodic representations of trained affixes and the (stem) contexts in which they were trained appear to decay over time, although we note that the effect is statistically marginal and requires further support.

The definition selection task was analysed using mixed-effects logistic regression (Table 2). While performance was highly accurate in both contexts, a significant effect of context was revealed, $z = 8.20$, $p < .001$, reflecting better performance with items from the training context. There was no significant effect of time-of-testing ($z = 0.75$, $p = .45$) nor any interaction between these factors ($z = 1.54$, $p = .13$).

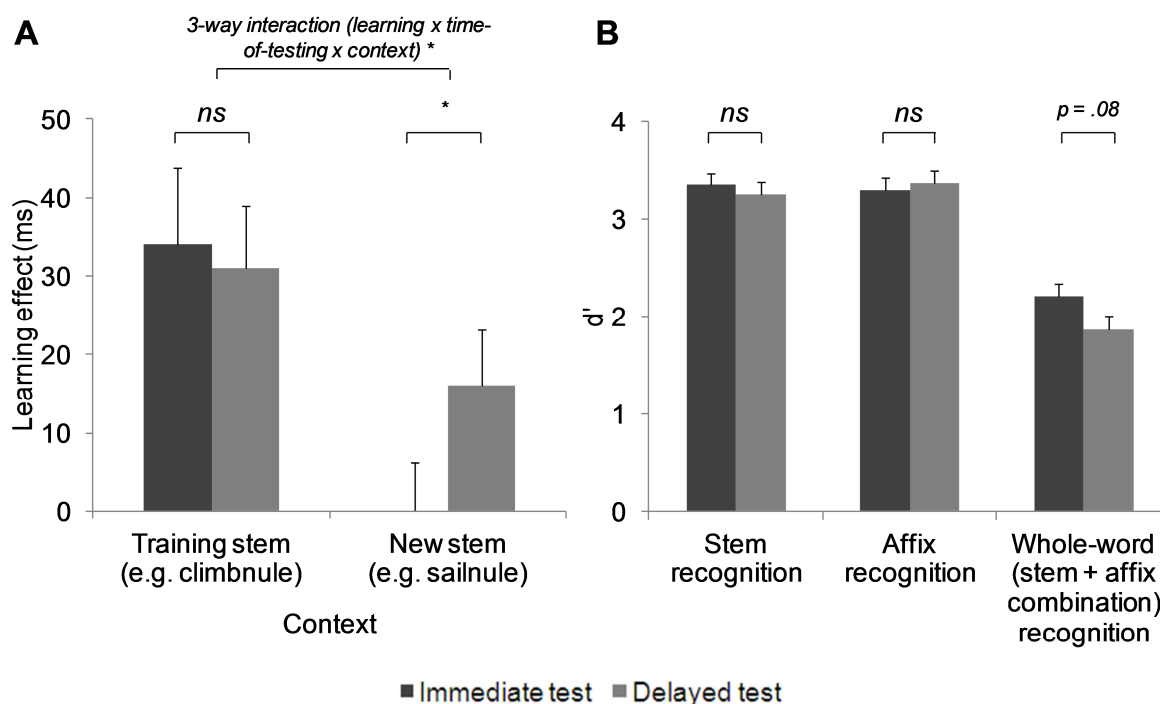


Figure 1. Magnitude of the learning effect (i.e. difference between RTs to words with untrained affixes and trained affixes) in the shadowing task, both when the affix was attached to the trained stem (showing learning of context-dependent affix representations) and with a new, untrained stem (showing learning of context-independent affix representations). Data are shown separately for the immediate and delayed time-of-testing conditions (A). D' scores in the recognition memory task (B). Error bars represent standard error. * $p < .05$, ns = not significant.

4. Discussion

Generalisation of linguistic knowledge from the context in which it was initially encountered to new contexts is a vital property of combinatorial systems at all levels of language processing (phonetic, lexical, semantic, etc). The power of these systems lies in the

ability of language users to combine context-independent units to form new wholes rapidly when speaking and to comprehend these novel combinations efficiently when listening. Our shadowing data show that in learning novel affixes evidence for the operation of context-independent linguistic representations is absent immediately after learning but is significantly enhanced following consolidation. Trained affixes enjoyed a processing advantage over untrained affixes when presented in a familiar context (i.e. with trained stems) in immediate and delayed testing. However, while context-dependent learning was apparent immediately, we saw a significant three-way interaction between affix learning, stem context, and time-of-testing, demonstrating that consolidation is required for affix knowledge to be apparent in speeded lexical processing outside of the training context (i.e. with untrained stems). This result is predicted by CLS accounts in which the emergence of generalisation in online tasks depends on consolidation processes that generate overlapping cortical representations.

We therefore suggest that previous reports of immediate generalisation in language learning reflect generalisation using context-dependent representations that may not be sufficient for online linguistic processing. Our non-speeded definition selection task also showed generalisation at both test times. Involvement of episodic memory in this task is supported by a significant correlation across all participants between whole-word recognition memory and accuracy in the definition selection task for the generalisation items (S-A+) ($r=.60$, $p<.001$, $r=.59$, $p<.001$ in immediate group, $r=.59$, $p<.001$ in delayed group), and by the very long definition selection RTs we observed to these ($M=4.4s$) and the trained ($M=3.1s$) items. No such correlation was observed with generalisation in shadowing in either group. The emergence of generalisation in shadowing occurred in parallel with changes in recognition memory performance. One intriguing possibility is that the emergence of context-independent affix representations is associated with a loss of episodic knowledge of the training context. However, despite a significant interaction demonstrating that episodic

memory performance declines more for learnt novel words than for learnt stems or learnt affixes, the decay in memory performance for learnt novel words was only marginally significant. Further evidence is needed to confirm whether there is any specific association between a decline in episodic knowledge of the training context and the development of context-independent lexical representations.

Our findings shed new light on the interpretation of related studies in this domain. For example, Gomez, Bootzin and Nadel (2006) conducted a study in which 15-month-old infants were exposed to an artificial language with an underlying syntactic rule, and reported that only infants who napped immediately after exposure applied the rule to new words at a later test measuring looking preference, suggesting that sleep-dependent consolidation was necessary for generalisation. The lack of immediate generalisation may have arisen because infants, unlike adults, are not equipped to use episodic knowledge for generalisation or because the preferential looking task, like shadowing in our case, was primarily sensitive to online linguistic processes. In either case, these data provide some precedent for our conclusions regarding consolidation in generalisation. Our findings add to those of Gomez and colleagues, though, by demonstrating in adult learning a clear dissociation between the time-course of generalisation in online and offline tasks. Similarly, Fenn, Nusbaum, and Margoliash (2003) showed immediate generalisation in phonetic learning which declined over the course of a day and rebounded after a night of sleep. Our data suggest that the immediate generalisation may have reflected operation of context-dependent phonetic representations followed by the emergence of context-independent representations after sleep.

Although our study was not designed to isolate the specific effect of sleep on learning and memory, there is an abundance of evidence suggesting that sleep may be the optimal brain state for consolidation (Diekelmann & Born, 2010). Neural instantiations of CLS models postulate a process of offline hippocampal replay that allows strengthening of

neocortical memory (e.g., Frankland & Bontempi, 2005) – a process that has been observed during sleep (e.g., Wilson & McNaughton, 1994; Rasch, Buchel, Gais, & Born, 2007).

Tamminen, Payne, Stickgold, Wamsley, and Gaskell (2010) showed that slow-wave sleep is associated with improving recognition of newly learnt words, while sleep spindles were associated with their integration in the mental lexicon. Future studies are needed to show whether these neurophysiological features of sleep also play a role in generalisation of newly learnt affixes.

In sum, we observed generalisation of new affix representations only two days after initial learning in a speeded lexical processing task. Episodic representations of new linguistic information were apparent immediately, and supported generalisation in a non-speeded task. This dissociation is problematic for theories of lexical representation that rely solely on episodic representations (e.g., Goldinger, 1998), and for theories which propose a single, abstractionist learning process that may be prone to catastrophic interference (e.g., Baayen et al., 2011; Plaut & Gonnerman, 2000). Our data provide unique support for CLS accounts which propose that the generation of new abstract linguistic representations requires consolidation. The critical feature of these accounts is that two distinct memory systems are involved in first encoding episodic representations of new linguistic knowledge, and second in generating abstract representations that support generalisation. We posit that an understanding of these general principles is invaluable to theories of language acquisition (Davis & Gaskell, 2009). For example, second language instruction depends on teaching pupils specific linguistic units and rules in a small set of trained contexts with the expectation that they will generalise this knowledge during natural comprehension and production. We provide the first evidence that the episodic representations that support initial learning are insufficient to support speeded generalisation of linguistic units outside of the trained context. Further research is required to establish the cognitive and neural processes responsible.

Acknowledgements

This research was funded by an Economic and Social Research Council grant to KR and MHD (RES-062-23-2268). We thank Lara Hemsworth for assistance in running the experiment.

References

- Algeo, J. (1991). *Fifty years among the new words: A dictionary of neologisms*. Cambridge, UK: Cambridge University Press.
- 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*, 7041-7045.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*, 390-412.
- Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, *118*, 438-481.
- Bates, E., & Liu H. (1996). Cued shadowing. *Language and Cognitive Processes*, *11*, 577-582.
- Davis, M. H., & Gaskell, M. G. (2009). A complementary systems account of word learning: Neural and behavioural evidence. *Philosophical Transactions of the Royal Society*, *364*, 3773-3800.
- Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, *11*, 114-126.
- Fenn, K. M., Nusbaum, H. C., & Margoliash, D. (2003). Consolidation during sleep of perceptual learning of spoken language. *Nature*, *425*, 614-616.

- Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature Reviews Neuroscience*, 6, 119-130.
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3, 128-135.
- Goldinger, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review*, 105, 251-279.
- Gomez, R. L., Bootzin, R. R., & Nadel, L. (2006). Naps promote abstraction in language-learning infants. *Psychological Science*, 17, 670-674.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428.
- Hintzman, D. L. (1988). Judgements of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95, 528-551.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London: Series B, Biological Sciences*, 262, 23-81.
- 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, 419-457.
- Meeter, M., & Murre, J. M. J. (2005). TraceLink: A model of consolidation and amnesia. *Cognitive Neuropsychology*, 22, 559-587.

- Merkx, M., Rastle, K., & Davis, M. H. (2011). The acquisition of morphological knowledge investigated through artificial language learning. *The Quarterly Journal of Experimental Psychology*, *64*, 1200-1220.
- 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*, 505-510.
- Plaut, D. C., & Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes*, *15*, 445-485.
- Rasch, B., Buchel, C, Gais, S., & Born, J. (2007). Odor cues during slow-wave sleep prompt declarative memory consolidation. *Science*, *315*, 1426-1429.
- Snodgrass, J. C., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34-50.
- 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*, 14356-14360.
- Taylor, J. S. H., Plunkett, K., & Nation, K. (2011). The influence of consistency, frequency, and semantics on learning to read: An artificial orthography paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 60-76.
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*, 676-679.

Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation.

Journal of the International Neuropsychological Society, 17, 766-780.