

Running head: CONSOLIDATION AND COGNITIVE MODELS

Models of cognition and constraints from neuroscience: A case study involving
consolidation

Stephan Lewandowsky

University of Western Australia

Ullrich K. H. Ecker

University of Western Australia

Simon Farrell

University of Bristol and University of Western Australia

Gordon D. A. Brown

University of Warwick and University of Western Australia

Stephan Lewandowsky

School of Psychology

University of Western Australia

Crawley, W.A. 6009, AUSTRALIA

email: stephan.lewandowsky@uwa.edu.au Web: <http://www.cogsciwa.com>

Abstract

Consolidation refers to the putative process by which existing memories are strengthened over time. There is widespread consensus within the neuroscientific community that consolidation is an important component of human memory. By contrast, the notion is rarely employed by cognitive modelers. We focus on behavioral data that have frequently been cited in support of consolidation—e.g., the Ribot gradient in amnesia and the temporal effects of retroactive interference—and show that (a) those data are in fact problematic for classic consolidation theory and (b) can be explained readily within a cognitive model based on temporal distinctiveness. We suggest that this changes the evidentiary landscape for consolidation and narrows the field of supporting evidence. Some of that evidence, gathered primarily in sleep studies, appears strong but is perhaps worthy of re-examination.

**Models of cognition and constraints from neuroscience: A
case study involving consolidation**

People forget. Information, once encoded, gradually becomes less accessible over time, being obscured or eroded by more recently encoded events. Notwithstanding this obvious truism, there is widespread agreement among neuroscientists that memories undergo a period of gradual *strengthening* upon encoding. This process, known as consolidation, is thought to render memories increasingly resistant to interference (Wixted, 2004, 2005). If people forget, how can they simultaneously also consolidate memories?

In this article we focus on behavioral evidence for the presumed consolidation of memories gathered in declarative-memory paradigms. To foreshadow briefly, we will draw two principal conclusions: First, some of the behavioral evidence presented in favor of consolidation, while potentially compatible with the notion, does not uniquely implicate consolidation. Second, the same behavioral evidence is naturally accommodated within an alternative framework based on temporal distinctiveness alone, and without recourse to consolidation. This article therefore challenges the evidentiary landscape for consolidation, and we conclude by identifying those areas of research that present the strongest evidence for consolidation and that could be productively re-evaluated to see if the notion withstands further scrutiny.

Consolidation presents an intriguing platform for the examination of theorizing in neuroscience and its relationship to purely cognitive approaches. First, consolidation is a curiously “invisible” construct; being a compensatory process that stands in opposition to forgetting. In most situations its presence can only be assumed but not directly demonstrated. That is, whatever the observed rate of forgetting, it can always be presumed that more forgetting would have occurred without consolidation. Consolidation is thus only detectable behaviorally in fairly special circumstances. One such circumstance is the case of retrograde amnesia, and in particular the temporal gradient of retroactive

memory loss that is apparent in many cases (for a review, see A. S. Brown, 2002). This temporal gradient, also known as the “Ribot” gradient after Theodule Ribot, who first suggested in 1881 that brain damage might impair recent memories more than temporally distant memories, refers to the fact that following onset of amnesia, patients are more likely to remember events from the distant past than more recent pre-morbid memories. This gradient is quite smooth and systematic when aggregated across individual studies, and it can extend over a remarkable length of time, being measured in decades in some instances (e.g., A. S. Brown, 2002).

In general, consolidation can be differentiated from other mechanisms by which memory might improve over time (e.g., rehearsal), by showing that there is a temporally limited window during which an amnestic agent can disrupt memory following the cessation of training (e.g., Dudai, 2004; Meeter & Murre, 2004). The temporal limit is crucial because if memory were impaired no matter when the agent is administered, then this might simply reflect an impairment of retrieval or destruction of a fully-formed trace. The reduction in effectiveness of the amnestic agent over time is therefore essential to all research on consolidation.

Another particularly elegant way to infer consolidation involves the build-up of interference observed on other, existing memorial content as a function of acquiring new information. For example, Gaskell and Dumay (2003) showed that exposure to novel pseudowords (“cathedruke”) slowed lexical decision of actual words (“cathedral”) on a delayed test but *not* on an immediate test, suggesting that the novel pseudowords had to be consolidated before taking on their interfering role in the lexicon.

Consolidation has been implicated at a number of different time scales. In addition to the Ribot gradient which spans years if not decades, consolidation has also been identified at the synaptic level (across time spans measured in minutes) and at the hippocampal level via lesioning studies in sub-human species (across a number of weeks).

In reviewing this distinction, Dudai (2004) characterized synaptic consolidation as a “. . .relatively fast type of process [that] takes place in local nodes in the neuronal circuit(s) that encode(s) the experience-dependent internal representation, i.e., the memory” (pp. 54-55). The evidence for synaptic consolidation is manifold; for example, injection of a protein synthesis inhibitor immediately before or upon learning will disrupt memory, but it will no longer do so if the injection is delayed one hour (Agranoff, Davis, & Brink, 1966; Meiri & Rosenblum, 1998). In humans, such consolidation over the short term can be identified using a dual-task methodology, which has yielded estimates of around .2 seconds per item for the time taken to “consolidate” an item into short term memory (Jolicoeur & Dell’Acqua, 1998).

The longer-term process, often referred to as “system consolidation” (Dudai, 2004), has been linked to the re-organization of memories over time, for example from initial temporary storage in the hippocampus to a more robust and long-lived cortical engram (McClelland, McNaughton, & O’Reilly, 1995). The evidence for system consolidation derives from lesioning studies, for example when the hippocampus of rats is lesioned some time after learning that a tone signals a subsequent electric shock. Similar to the pattern observed over shorter time scales with synaptic consolidation, the longer the hippocampal lesion is delayed, the less impairment of learning is observed relative to unlesioned (or sham-lesioned) control animals (e.g., Kim & Fanselow, 1992).

When joined together with the Ribot gradient, the various streams of evidence form the pattern shown in the three panels in Figure 1. The left-most panel shows synaptic consolidation, the center panel system consolidation, and the right-most panel shows the decade-long Ribot gradient summarized by A. S. Brown (2002). Although the data are presented in separate panels, they are connected conceptually by a common underlying (logarithmic) timeline. In all panels, performance is expressed as a relative percentage; that is, relative to a control group which was not exposed to the amnesic agent.

The figure clarifies the appeal of the consolidation concept as well as its more troubling aspects: On the one hand, consolidation provides an elegant account of why memory can seemingly improve over time, and that it can do so at various time scales (e.g., see also McGaugh, 2000). On the other hand, the zig-zag pattern across the three panels presents a conundrum to consolidation theorists: In particular, it is difficult to reconcile the seemingly complete consolidation after 25-30 days (center panel) with the further consolidation over several decades that originates with a “reset” to a lower point (right panel) under the single umbrella of “system” consolidation.

Admittedly the three panels involve very different data—and indeed different species—so the numbers are unlikely to be commensurate despite being expressed as relative scores. We nonetheless suggest that it is challenging to describe the three manifestations of consolidation across the entire range of time spans shown in the figure within existing theoretical distinctions. Specifically, given that the seeming discontinuity between the left and center panels has been adduced in support of the distinction between synaptic and systems consolidation (e.g., Dudai, 2004), the same inferential logic would imply the necessity for a further theoretical distinction based on the apparent discontinuity between the center and right panels.

The multiple temporal windows over which consolidation has been observed, with performance “reset” to a low point anew within each window, set the stage for our re-examination of the notion of consolidation. We suggest that further exploration of the pattern in Figure 1 might proceed along one of two broad avenues: First, one might seek further differentiations between different manifestations of consolidation, each with its own time scale and perhaps associated with distinct brain regions or neural substrates. This approach is not without merit and falls within the “taxonomic” tradition in the neurosciences that has been applied elsewhere with considerable success. For example, in recognition memory, neuroscientists have discovered distinct neural clusters for the

generation of recency, episodic familiarity, long-term familiarity, and novelty signals (Daselaar, Fleck, & Cabeza, 2006; Hölscher, Rolls, & Xiang, 2003; Xiang & Brown, 1998). Up to four distinct types of associative novelty/familiarity signals have been identified (Düzel, Habib, Guderian, & Heinze, 2004). However, a potential drawback of this approach is that it may yield an unwieldy number of distinctions, processes, and neural substrates that may not map onto an equal number of distinct psychological phenomena. This risk appears particularly pronounced in the case of consolidation, in which distinctions between processes and presumed substrates are based on temporal parameters; it is thus possible that increasingly fine-grained temporally-based distinctions will run the risk of diluting the theoretical value of the consolidation notion.

The second approach, by contrast, would seek an integrative explanation of the temporal pattern in Figure 1, even if that explanation is at least initially formulated at a purely cognitive level. We proceed as follows: We first show that some of the presumed temporal signature effects of consolidation—such as the Ribot gradient—can be explained by a cognitive process model based on temporal distinctiveness without any involvement of consolidation. Crucially, the zig-zag pattern across the panels in Figure 1 falls out of a deep property of the model, reflecting the fact that memory is cued on the basis of relative, not absolute, time.

Reconsidering Temporal Gradients

Despite the long-standing association between consolidation and the neuroscience of memory, some of the earliest evidence for consolidation was entirely behavioral: Müller and Pilzecker (1900) inferred the presence of consolidation from the pattern of interference observed with memory lists involving verbal material. As reviewed in great detail by Dewar, Cowan, and Della Salla (2007), it was Experiment 34 of Müller and Pilzecker (1900)'s study that provided the first behavioral evidence for consolidation. In that study,

the temporal interval between a first (critical) list and a subsequent (interfering) list was manipulated. First-list recall was found to increase considerably when the interfering list was withheld for 6 minutes rather than just 17 seconds. This finding is consonant with consolidation: Because memories are initially weak, they are overwhelmed by the second list if it is presented after a short time. Once memories have had a chance to consolidate, the interference caused by the second list is much reduced—not because the second list is encoded less well, but because its encoding causes less damage to the existing memories.

A first step towards re-examining the consolidation notion must therefore focus on these early behavioral data, and other related results reported subsequently. Wixted (2004) provided an elegant review of the behavioral literature relevant to consolidation. In particular, Wixted focused on 5 principal phenomena, all but one of which are isomorphs or derivatives of the temporal gradient of retroactive interference (the remaining one relates to the shape of the forgetting function which we briefly take up at the very end.)

Temporal gradient of retroactive interference

The basic paradigm that gives rise to the temporal gradient of retroactive interference (Müller & Pilzecker, 1900) has already been sketched out: The top panel of Figure 2 illustrates the methodology by placing the relevant events along a timeline. In all cases, the original material is presented first (labeled L1; usually paired associates A-B), which is followed by a retention interval that is constant across conditions (a) through (c). The time of recall at the right refers to retrieval of the first list (L1). The interfering second list (L2; usually pairs A-C) is presented at varying intervals during the retention interval.

Müller and Pilzecker (1900)'s data showed that performance improved as the interfering material was delayed (i.e., condition (a) vs. (b)), and Wixted (2004) reviewed a number of additional studies that show the same general pattern. The bottom panel of

Figure 2 shows representative data, taken from Newton and Wickens (1956, Experiment 1); the arrows connect the conditions to their observed means.

In many cases, including that shown in the figure, performance drops again when the interfering material is presented close to the test (condition (c)), but this decline is also frequently absent. Because this decline is often absent, and because it requires a second process to be explained even within the consolidation framework (Wixted, 2004), we do not focus on it here. Instead, we consider a temporal-distinctiveness model known as SIMPLE (G. D. A. Brown, Neath, & Chater, 2007), which describes memory performance as arising from the confusability between memory traces, and apply it to explaining the improved recall that arises when interfering material is delayed (i.e., condition (a) vs. (b) in Figure 2).

In SIMPLE, the confusability of any two traces depends on the *ratio* of the times that have elapsed between their encodings and the time of recall. The lower that ratio, the less the confusability among items and hence the more likely recall is to be correct. This simple assumption renders recent items less confusable and hence more memorable than more distant events. For example, items that were encoded 1 s and 2 s ago are less confusable (ratio of .5) than are items from 5 and 6 seconds ago (.83). The ratio mechanism also automatically favors items that were separated in time over others that occurred in close temporal proximity, all other things being equal. For example, items that occurred 5 s and 10 s ago (ratio .5) are less confusable than items that occurred 7 s and 8 s ago (.88), even though the average retention interval is equal for both pairs of items. It follows that items from further in the past, and items that occurred in greater temporal proximity, will be more difficult to recall.

From this description, it may already be apparent how the model might handle the classic temporal gradient of retroactive interference shown in Figure 2: In SIMPLE, all forgetting arises from interference, not trace decay. If interfering material is encoded in

close temporal proximity to to-be-remembered material (condition a), memory items will be made less temporally distinct. Temporal distinctiveness will increase with the temporal gap between to-be-remembered and interfering material (condition b). Thus the temporal gradient could emerge not because an increasing temporal gap allows for more consolidation, but because of the greater temporal distinctiveness conferred by the interfering material onto the to-be-remembered material.¹

More formally, in SIMPLE memory traces are positioned in a potentially multi-dimensional space in which one dimension—and often the only one—represents time. Events are placed along the temporal dimension as they occur, with their spacing determined by the time that elapses between events. The confusability or similarity of any two memory traces is defined by an exponential function of the distance between them in psychological space, the similarity-distance metric given by:

$$\eta_{ij} = e^{-c d_{ij}}, \quad (1)$$

where η_{ij} is the similarity between items i and j , d_{ij} is the distance between them, and c is a parameter that denotes the rate at which similarity declines with distance. If time is the only dimension of the presumed psychological space, then d_{ij} reduces to the time that has elapsed between items i and j . More specifically, the temporal dimension represents the logarithmically transformed time that has elapsed since encoding; that is, all times are expressed relative to a recall attempt in the presence. The logarithmic transformation allows for the conceptualization of similarity between two items in terms of the ratio of their times since encoding.

An item's distinctiveness or discriminability is given by:

$$D_i = \frac{1}{\sum_{k=1}^n (\eta_{i,k})}, \quad (2)$$

where D_i is the summed discriminability of the probed item i in relation to all n potentially recallable items (e.g., all other list items).

Recall probability is proportional to an item's discriminability as given by Equation 2 but is additionally affected by a recall threshold. Items that exceed a certain discriminability threshold are likely to be recalled, items that fall below the threshold are likely to be omitted. Hence final recall probability is given by:

$$P_i = \frac{1}{1 + e^{-s(D_i - t)}}, \quad (3)$$

where t is the omission threshold (i.e., the discriminability value below which items are not recalled), and s is the slope of the transforming function, which can be conceptualized as the noisiness of the threshold (i.e., large values of s would approximate a very sharp, precise threshold, low values of s would yield a noisy and more probabilistic threshold). This version of SIMPLE is thus defined by the three free parameters c , t , and s .

We applied this model to the retroactive-interference paradigm at various different time scales. In particular, we generated predictions for three prototypical experiments at roughly the same time scales as those shown in the three panels in Figure 1; namely, total retention intervals of 80, 36,000, and 21,000,000 minutes. Those values cover the ranges observed in Figure 1. Using a single-point representation for each of the two lists for simplicity, SIMPLE predicted the pattern of first-list recall shown in Figure 3. Each point in the figure corresponds to a unique delay between the two lists that is proportional to the total retention interval (viz., .1, .2, .4, and .8, respectively, in each panel).

The pattern predicted by SIMPLE closely resembles the data in Figure 1. It is important to note that SIMPLE's predictions arose entirely from the temporal distinctiveness mechanism discussed above. In particular, the predictions fall out of an important feature of temporal distinctiveness, namely its scale-invariance, which arises from the consideration of relative time ratios rather than absolute time in the calculation of temporal item similarity. In consequence, items studied 20 seconds and 40 seconds ago are as similar to each other as items studied 20 and 40 hours ago. SIMPLE therefore

predicts the *same* pattern across *different* timescales, based on the *same* process. This prediction holds irrespective of the number of timescales; that is, for SIMPLE it does not matter whether Figure 1 contains 3, 5, or 500 panels at different timescales. In contrast, explanations in terms of consolidation theory require the assumption of two, if not three—or indeed 500—different neural processes operating on different timescales. It is important to note that scale invariance cannot arise during consolidation because it is a property of the relative comparison between two times which are not known until the time of retrieval: Consolidation, by contrast, must operate on a fixed timescale which by definition is only related to the time of encoding (i.e., when the consolidation process starts), not the time of retrieval which at that point is unknown.

Our demonstration therefore shows that a simple temporal-distinctiveness model can account for one of the behavioral patterns taken to require neuroscientific distinctions—namely, the temporal gradients of retroactive interference across different timescales—without any such distinction but within a single process.

Generality of the Account

In a preliminary application of the distinctiveness account, we showed that SIMPLE can also handle two other pieces of evidence commonly taken to implicate consolidation, namely the Ribot gradient in retrograde amnesia, and the form of the forgetting curve (G. D. A. Brown & Lewandowsky, 2010).

The Ribot gradient in retrograde amnesia—the temporally graded loss of access to recent memories—is one of the most commonly cited phenomena in support of consolidation theory (e.g., Meeter & Murre, 2004). According to the standard model of consolidation, the Ribot gradient results from damage to brain areas—primarily the hippocampus—that are involved in the storage recent memories (e.g., Axmacher, Draguhn, Elger, & Fell, 2009). Because memories are gradually transferred or re-encoded

in neocortical areas, the extent of memory loss diminishes for events that have receded further in time and hence have become less dependent on hippocampal storage. This model is, however, by no means the only possible neuroscientific explanation for the Ribot gradient (see, e.g., Moscovitch, 2008), and here we extend an alternative account based on SIMPLE that was sketched by G. D. A. Brown and Lewandowsky (2010).

Crucial to this account is the fact that memory representations in SIMPLE frequently involve not just a temporal dimension, but are augmented by other dimensions such as an ordinal representation of events (cf. Lewandowsky, Duncan, & Brown, 2004) or similarity among items (G. D. A. Brown et al., 2007). Indeed, “although temporal organization is assumed to be primary, we assume that psychological space will become organized along whatever dimensions are most accessible and useful for a particular task at hand. In particular, the inclusion of a nontemporal positional dimension may be required” (G. D. A. Brown et al., 2007, p. 568). Given that both temporal and positional dimensions are typically involved in representing memories, the Ribot gradient emerges from SIMPLE as follows: the logarithmic compression of time (see earlier discussion surrounding Equation 2) implies that it will be a good retrieval cue for recent items and events—because they are spaced far apart—whereas its utility as a cue declines with absolute time. It would thus make sense for the memory retrieval system to pay relatively greater attention to the temporal dimension when retrieving recent items, and to pay relatively less attention to the temporal dimension and correspondingly greater attention to other dimension(s) for items further in the past. SIMPLE can therefore model the temporally-graded loss of recent memories by assuming that access to the temporal dimension—but not other representational dimensions—is lost in retrograde amnesia. This assumption meshes well with neuroscientific results that have implicated the hippocampus specifically in the formation of *temporal* associations (e.g., Downes, Mayes, MacDonald, & Hunkin, 2002; Howard, Kahana, & Wingfield, 2006). Thus, hippocampal

damage can elicit the Ribot gradient by impairing the use of specifically temporal information, rather than by revealing the cortical transfer of memories assumed by the standard consolidation model. G. D. A. Brown and Lewandowsky (2010) reported a simulation of the Ribot gradient based on this assumption within SIMPLE. The results are shown in Figure 4.

Another piece of evidence that has been taken to support consolidation theory is the form of the forgetting curve. In particular, according to “Jost’s Second Law” (Wixted, 2004), when two memory traces have equal strength, the older trace will be less prone to forgetting than the newer one. This is consistent with the idea that the older item has had more time to consolidate and is thus more resistant to forgetting (cf. Meeter, Murre, & Janssen, 2005). However, SIMPLE’s simulations of forgetting curves can produce several forms of the forgetting curve—in line with various forms described in the literature—all of which adhere to Josts’ Second Law without involvement of a consolidation process (G. D. A. Brown & Lewandowsky, 2010). This arises from a natural property of ratio-based temporal distinctiveness because the confusability of two items will necessarily increase as time passes, and this increase will be rapid at first but will gradually slow down. For example, two items that were encoded 10 and 20 seconds ago (confusability ratio $10/20 = 0.5$) will be less confusable than two items that were encoded 110 and 120 seconds ago (ratio $110/120 = 0.917$). Letting an additional 10 seconds pass, the confusability of the first pair of items will have increased markedly ($20/30 = .67$), whereas the confusability of the second item pair will have only increased by a marginal amount ($120/130 = 0.923$).

In summary, we note that using a single cognitive principle—temporal distinctiveness—has enabled us to explain behavioral data that have often been used to support consolidation theory. In particular, we could successfully bridge ostensibly different manifestations of consolidation across different time scales with a single process.

There have been other demonstrations of unifying principles being able to accommodate memory phenomena across vastly different timescales (Moreton & Ward, 2010; Watkins & Peynircioglu, 1983); and indeed, this effort towards greater parsimony and scale-invariance can also be found in other domains, such as linguistic behavior, psychophysics, and motor control (Kello et al., 2010).

Conclusions

Before we draw conclusions from our analysis, several qualifying thoughts are in order. Perhaps most important is to clarify what we did *not* say: We did not argue against the existence of consolidation, and we are not putting forth our temporal-distinctiveness model as an exclusive alternative to the long-standing research tradition on consolidation. Instead, we focused on temporal-behavioral data that have frequently been offered in support of consolidation theory and showed (a) that those data are in fact at best compatible with consolidation but not uniquely supportive of it, and (b) that the very feature that makes the data problematic for consolidation theory—viz. the existence of multiple competing time scales—falls out naturally from our distinctiveness model.

However, we are very aware that the literature on consolidation is vast and cannot be done justice in a single rather brief article. Beyond the data that we have explained within temporal distinctiveness, we can point to three lines of evidence that appear to offer particularly compelling support for consolidation theory: First, there is the large body of research surrounding sleep, which has provided strong evidence for some type of consolidation process (e.g., Born, Rasch, & Gais, 2006; Born, 2010; Dumay & Gaskell, 2007). Second, there is evidence that memory can be improved during sleep through the application of transcranial magnetic stimulation (TMS) by inducing forms of brain activity that have been implicated in consolidation by other sleep research (e.g. ?, ?; Marshall, Helgadottir, Mölle, & Born, 2006; ?, ?). Finally, there are reports of memory being

qualitatively restructured—rather than being merely improved—during consolidation (e.g. Wagner, Gais, Haider, Verleger, & Born, 2004; Walker & Stickgold, 2010).

Where does this leave us and where do we go from here? We suggest that further progress can be achieved along two prongs: First, the sources of strong evidence just cited might possibly also yield to further re-examination, similar to the way that the behavioral-temporal evidence has yielded to re-examination in this article.

Second, and more generally, progress in this area as any other will be tied to systematic model comparison. Our analysis could succeed only by showing that SIMPLE produced the principal effects at a quantitative level, an outcome that verbal analysis alone could not have ascertained. What remains to be seen is how far this model can be pushed: In the same way that single-process signal-detection models can form a useful baseline against which other, more complicated recognition models can be evaluated (e.g., DeCarlo, 2008; Wixted, 2007), the temporal-distinctiveness model developed here can serve as a baseline against which other candidate proposals can be evaluated by formal model-comparison means (i.e., taking into account differences in complexity; see Lewandowsky & Farrell, 2011, Chapter 5).

Several obvious candidates exist that remain to be formally evaluated. For example, the “TraceLink” model of Meeter and Murre (e.g., Meeter & Murre, 2004, 2005) represents an impressive effort to account for existing data on amnesia and consolidation—including some of those modeled above. Meeter and Murre (2004) presented 9 simulations of TraceLink, ranging from an application to the “permastore” through the Ribot gradient to implicit learning in amnesia. Their work provides an impressive sufficiency proof for the power arising from a model that contains a consolidation process; however, what remains to be seen is whether the inclusion of a consolidation process is *necessary* to retain the model’s power and whether its additional complexity (some 8-10 parameters compared to

3 in SIMPLE) is statistically justified relative to the increase in power. We consider this to be an exciting avenue for future modeling.

Furthermore, computational modeling can inform research from cognitive neuroscience, where model parameters can be linked to brain activity (as recorded by EEG or fMRI), and can be used to make inferences about changes in cognitive mechanisms that would not otherwise be possible (e.g., Ho, Brown, & Serences, 2009). We suggest that there is a natural synergistic relationship between modeling and neuroscience; neuroscience provides constraining data for models that can go beyond behavioral measures, and models give us a better handle on our neuroscience data. Some of the other papers in this special issue exemplify this synergistic approach perfectly.

References

- Agranoff, B. W., Davis, R. E., & Brink, J. J. (1966). Chemical studies on memory fixation in goldfish. *Brain Research*, *1*, 303–309.
- Axmacher, N., Draguhn, A., Elger, C. E., & Fell, J. (2009). Memory processes during sleep: beyond the standard consolidation theory. *Cellular and Molecular Life Sciences*, *66*, 2285–2297.
- Born, J. (2010, June). Slow-wave sleep and the consolidation of long-term memory. *World Journal of Biological Psychiatry*, *11*, 16–21.
- Born, J., Rasch, B., & Gais, S. (2006). Sleep to remember. *The Neuroscientist*, *12*, 410–424.
- Brown, A. S. (2002). Consolidation theory and retrograde amnesia in humans. *Psychonomic Bulletin and Review*, *9*, 403–425.
- Brown, G. D. A., & Lewandowsky, S. (2010). Forgetting in memory models: Arguments against trace decay and consolidation failure. In S. Della Sala (Ed.), *Forgetting* (pp. 49–75). Hove, UK: Psychology Press.
- Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*, 539–576.
- Daselaar, S. M., Fleck, M. S., & Cabeza, R. (2006). Triple dissociation in the medial temporal lobes: Recollection, familiarity, and novelty. *Journal of Neurophysiology*, *96*, 1902–1911. (cited By (since 1996) 103)
- DeCarlo, L. T. (2008). Process dissociation and mixture signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 1565–1572.
- Dewar, M. T., Cowan, N., & Della Salla, S. (2007). Forgetting due to retroactive interference: A fusion of müller and pilzeckers (1900) early insights into everyday forgetting and recent research on anterograde amnesia. *Cortex*, *43*, 616–634.

- Downes, J. J., Mayes, A. R., MacDonald, C., & Hunkin, N. M. (2002). Temporal order memory in patients with korsakoff's syndrome and medial temporal amnesia. *Neuropsychologia*, *40*, 853-861. (Article 538QQ English Times Cited:8 Cited References Count:36)
- Dudai, Y. (2004). The neurobiology of consolidation, or, how stable is the engram? *Annual Review of Psychology*, *55*, 51–86.
- Dumay, N., & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words. *Psychological Science*, *18*(1), 35–39.
- Düzel, E., Habib, R., Guderian, S., & Heinze, H. (2004). Four types of novelty-familiarity responses in associative recognition memory of humans. *European Journal of Neuroscience*, *19*(5), 1408-1416. (cited By (since 1996) 17)
- Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition*, *89*(2), 105–132.
- Ho, T. C., Brown, S., & Serences, J. T. (2009). Domain general mechanisms of perceptual decision making in human cortex. *Journal of Neuroscience*, *29*, 8675–8687.
- Hölscher, C., Rolls, E. T., & Xiang, J. (2003). Perirhinal cortex neuronal activity related to long-term familiarity memory in the macaque. *European Journal of Neuroscience*, *18*(7), 2037-2046. (cited By (since 1996) 26)
- Howard, M. W., Kahana, M. J., & Wingfield, A. (2006). Aging and contextual binding: Modeling recency and lag-recency effects with the temporal context model. *Psychonomic Bulletin & Review*, *13*, 439–445.
- Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology*, *36*, 138-202.
- Kello, C. T., Brown, G. D. A., Cancho, R. Ferrer-i, Holden, J. G., Linkenkaer-Hansen, K., Rhodes, T., et al. (2010). Scaling laws in cognitive sciences. *Trends In Cognitive Sciences*, *14*(5), 223–232.

- Kim, J. J., & Fanselow, M. S. (1992). Modality-specific retrograde amnesia of fear. *Science*, *256*, 675–677.
- Lewandowsky, S., Duncan, M., & Brown, G. D. A. (2004). Time does not cause forgetting in short-term serial recall. *Psychonomic Bulletin & Review*, *11*, 771–790.
- Lewandowsky, S., & Farrell, S. (2011). *Computational modeling in cognition: Principles and practice*. Thousand Oaks, CA: Sage.
- Marshall, L., Helgadottir, H., Mölle, M., & Born, J. (2006). Boosting slow oscillations during sleep potentiates memory. *Nature*, *444*, 610–613.
- 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.
- McGaugh, J. L. (2000). Memory—a century of consolidation. *Science*, *287*, 248–251.
- Meeter, M., & Murre, J. M. J. (2004). Consolidation of long-term memory: Evidence and alternatives. *Psychological Bulletin*, *130*, 843–857.
- Meeter, M., & Murre, J. M. J. (2005). Tracelink: A model of consolidation and amnesia. *Cognitive Neuropsychology*, *22*, 559–587.
- Meeter, M., Murre, J. M. J., & Janssen, S. M. J. (2005). Remembering the news: Modeling retention data from a study with 14,000 participants. *Memory & Cognition*, *33*, 793–810.
- Meiri, N., & Rosenblum, K. (1998). Lateral ventricle injection of the protein synthesis inhibitor anisomycin impairs long-term memory in a spatial memory task. *Brain Research*, *789*(1), 48–55. (cited By (since 1996) 54)
- Moreton, B. J., & Ward, G. (2010, August). Time scale similarity and long-term memory for autobiographical events. *Psychonomic Bulletin & Review*, *17*(4), 510–515.

- Moscovitch, M. (2008). The hippocampus as a stupid, domain-specific module: Implications for theories of recent and remote memory, and of imagination. *Canadian Journal of Experimental Psychology*, *62*, 62–79.
- Müller, G. E., & Pilzecker, A. (1900). Experimentelle Beiträge zur Lehre vom Gedächtnis [experimental contributions to the science of memory]. *Zeitschrift für Psychologie*, *1*, 1–300.
- Newton, J. M., & Wickens, D. D. (1956). Retroactive inhibition as a function of the temporal position of the interpolated learning. *Journal of Experimental Psychology*, *51*, 149–154.
- Wagner, U., Gais, S., Haider, H., Verleger, R., & Born, J. (2004). Sleep inspires insight. *Nature*, *427*, 352–355.
- Walker, M. P., & Stickgold, R. (2010). Overnight alchemy: sleep-dependent memory evolution. *Nature Reviews: Neuroscience*, *11*.
- Watkins, M. J., & Peynircioglu, Z. F. (1983). Three recency effects at the same time. *Journal of Verbal Learning and Verbal Behavior*, *22*, 375–384.
- Wixted, J. T. (2004). The psychology and neuroscience of forgetting. *Annual Review of Psychology*, *55*, 235–269.
- Wixted, J. T. (2005). A theory about why we forget what we once knew. *Current Directions in Psychological Science*, *14*, 6–9.
- Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, *114*, 152–176.
- Xiang, J.-Z., & Brown, M. W. (1998). Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology*, *37*, 657–676. (cited By (since 1996) 137)

Author Note

Preparation of this paper was facilitated by a Discovery Grant from the Australian Research Council and an Australian Professorial Fellowship to the first author. Address correspondence to the first author at the School of Psychology, University of Western Australia, Crawley, W.A. 6009, Australia. Electronic mail may be sent to stephan.lewandowsky@uwa.edu.au. Web: <http://www.cogsciwa.com>.

Footnotes

¹ In its current unmodified form, SIMPLE cannot explain the decline between conditions (b) and (c) in Figure 2. With an auxiliary assumption involving relative memorability—i.e., involving List 2 in the computation of retrieval probabilities—SIMPLE can explain the full inverted U-shaped pattern. For simplicity, this modification is omitted here and thus SIMPLE's predictions are limited to the ascending limb of the inverted-U in Figure 2.

Figure Captions


Figure 1. Consolidation at three time scales, ranging from minutes (left-hand panel) to days (center panel) and years (right-hand panel). The left panel shows data from Agranoff et al. (1966), replotted as a percentage of consolidated memory following Dudai (2004). In the experiment, goldfish learned to escape an electric shock in a shuttle-box tank. A protein synthesis inhibitor was administered to separate groups of fish at the times after training shown on the abscissa. Consolidation appeared complete after one hour; that is, the drug did not impair performance relative to a no-treatment control. The center panel shows data from a fear-condition experiment by Kim & Fanselow (1992), also re-expressed as a percentage of consolidated memory. Rats were fear conditioned via electric shock before their hippocampus was lesioned at the time shown on the abscissa after training. The longer the lesion was delayed the less it impaired performance, suggesting that memory for the fear-invoking stimulus had become increasingly consolidated. The right-hand panel shows data from A. S. Brown (2002, Figure 1) that reflect performance of amnesic patients as a percentage of the performance of matched control participants. See text for details.

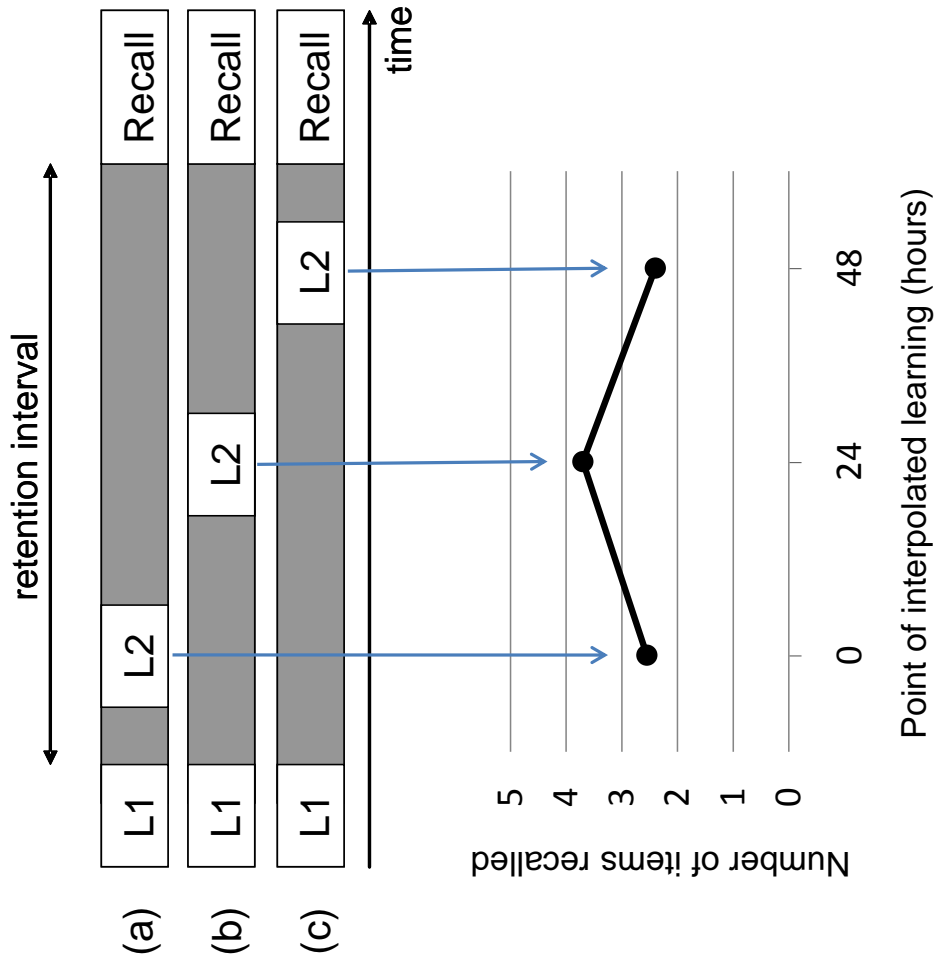
Figure 2. Top panel shows the temporal structure of three conditions (a through c) to observe the temporal gradient of retroactive interference. See text for explanation. The bottom panel shows data from a representative experiment (Newton & Wickens, 1956), and the arrows connect conditions to their corresponding outcome.

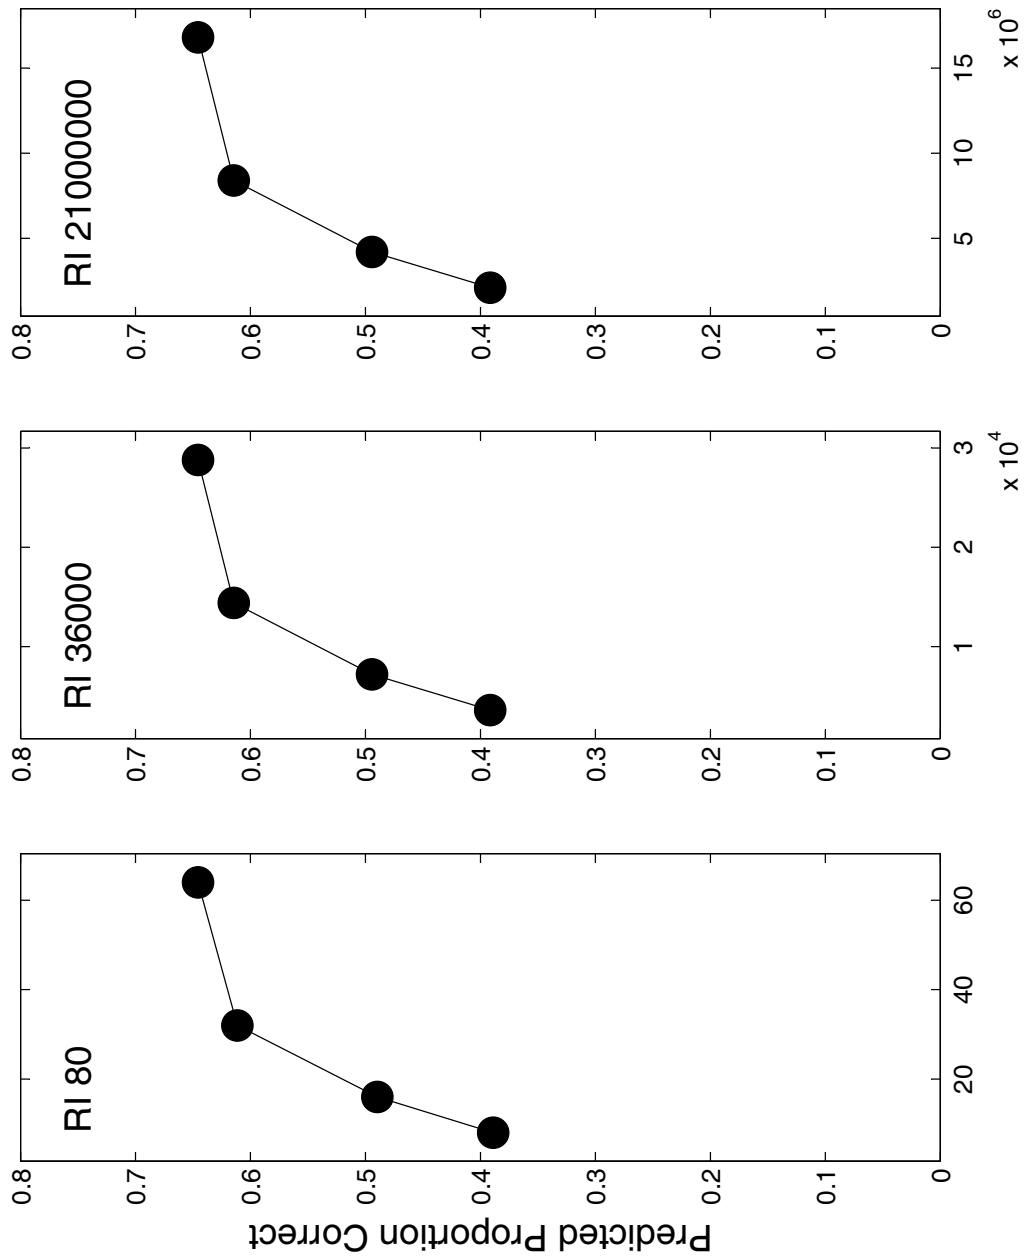
Figure 3. SIMPLE's predicted pattern of first-list recall (L1 in Figure 2) for three retroactive interference experiments at three different time scales. The three panels correspond, roughly, to the same time scales as those observed in Figure 1. Across all panels, parameters were held constant at $c = 6$, $t = .8$, and $s = 3$. See text for further details.

Figure 4. Illustration of how a typical recency gradient (top line) may be transformed into temporally graded amnesia (bottom line) within SIMPLE if access to a temporal dimension is lost due to hippocampal damage. The simulation assumes that the attentional weight given to the temporal dimension during recall reduces as a linear function of the temporal distance of the to-be-retrieved memory, reflecting the notion that temporal discrimination is most useful for recent memories whereas more distant memories are best retrieved on the basis of other information. See text for further details.

Consolidation and Cognitive Models , Figure 1







Delay between List 1 (Target) and List 2 (in minutes)

