The Interaction of Learning Speed and Memory Interference:

When Fast is Bad

Øystein Sandvik

Master of Philosophy in Psychology

Cognitive Neuroscience

Department of Psychology

University of Oslo

May 2013

Acknowledgments

I would like to thank my supervisor Dr. Guido Biele for constructive guidance throughout the work on this thesis. Dr. Biele's advice was critical when developing the idea for my project. He also provided invaluable advice during the data analysis as well as feedback on the writing of my thesis. The author, however, will have to be held accountable for the end result. All stages of the project, from idea to final thesis, has been planned and carried out by me.
I would also like to thank my wife Siri. If it had not been for her moral support, this thesis would simply not be. Let's just hope it was worth it.

Abstract

Research on individual differences in speed of learning has suggested that forgetting rates could be different for fast and slow learners. Studies have shown either no difference or slower forgetting over time for fast learners. The present study extends this area of research by investigating the possibility that fast and slow learning are differentially vulnerable to interference. Based on neural network models and the encoding variability hypothesis, two novel hypotheses were built and tested in two experiments by a paired-associates task. The hypotheses suggested that fast learning will be more prone to interference when similarity of the learning material is high. Hence, an interaction of learning speed and interference (i.e., similarity) was predicted. Experiment 1 ($N = 22$) compared retention of Chinese characters for fast and slow learning (both subject and item-specific speed) by manipulating similarity (high vs. low) of the characters learned. Results of Experiment 1 were inconclusive. Experiment 2 ($N = 21$) had the same basic design as Experiment 1, but included a number of procedural improvements. Interactions in the predicted direction were found both when comparing learning speed between subjects as well as for item-specific speed. However, only the interaction of between-subjects learning speed and similarity was significant. A joint analysis, including data from both experiments, yielded significant interactions for both subject speed and item-specific speed, indicating that the lack of a significant interaction of item-specific speed and similarity in Experiment 2 was probably due to the low sample size. The findings are discussed in relation to previous research on individual differences in learning speed and forgetting.

Contents

Introduction

Some things are learned fast, while other requires more effort. However, once something is learned, does it matter how much time it took to learn it? Will the things that were difficult to learn be more easily forgotten? What about individual differences? Some people naturally acquire new knowledge faster than others, but if the slow learner is granted enough time, would the knowledge of the fast and the slow learner be comparable? Is it possible to equate associational strength between individuals learning at different speeds? If the slow learner was given a head start and both the fast and slow learner discontinued learning once a common level of associational strength had been reached, would there be a difference in associational strength 24 hours later?

These are questions that have troubled researchers for a long time, but any firm conclusions are yet to be produced. Research on individual differences in learning and forgetting has provided inconsistent results with respect to these questions. The results of prior studies have either indicated slower forgetting for fast learners or no individual differences in forgetting. The present investigation attempts to inform this area of research by addressing two questions that have not been properly considered in prior studies: (a) Do fast and slow learning affect established representations differentially? In other words, do fast and slow learning produce different levels of interference on existing memory representations? (b) Are memory representations established by fast and slow learning comparable even if response strength has been equated? The interpretation of the results of prior research depends strongly on these questions. Because forgetting has usually been treated as a function of time in prior studies, the possibility that memory representations established by fast and slow learning have different vulnerability to interference has been neglected. Furthermore, if fast and slow learning produces different levels of interference on prior knowledge, does it even make sense to compare forgetting rates over time for fast and slow learners? If fast and slow learning are differentially affected by subsequent learning, any differences in forgetting rates could be more related to the specific learning material used as stimuli rather than any general differences in forgetting over time.

The investigation reported on in the following considers several areas of research that collectively suggest a common role of interference in the relationship between learning speed and forgetting. More specifically, it is suggested that interference – operationalized by stimuli similarity – affects memory representations established by fast and slow learning differentially. By considering research results from simulations on computational models of cognition in relation to question (a) it is suggested that fast learning will generally cause higher levels of interference when similarity of the learning material is high. Theoretical considerations of the encoding variability hypothesis informs (b) by suggesting that representations established by fast and slow learning will be different and therefore have differential vulnerability to interference, even when response strength of learned associations are equated. Interestingly, they both suggest the same behavioral level effect. Based on this, two specific hypotheses are constructed and tested experimentally. The general suggestion of these hypotheses is that fast learning is more prone to interference and retention will therefore be more strongly impaired compared to slow learning when interference (i.e., similarity of subsequent learning) increases. In order to establish proper theoretical justifications for these hypotheses, a somewhat wide, and at times discursive, introduction of prior research will be required. However, before embarking, a few terminological remarks are in order.

The terms *learning rate*, *rate of acquisition*, and *speed of learning* are used somewhat interchangeably in the literature. In the current report, consistency will be approached by using learning rate when referring to model parameters, while learning speed and rate of acquisition will refer to the speed or number of trials required for an agent (human or model) to reach a state that enables recall or recognition of the learning material.

The term *interference* can be somewhat confusing in the context of the present investigation. In one way, interference can be a label referring to the role of an independent variable that is manipulated in order to produce differential forgetting. Another way to use this term is to refer to interference as the very effect that is produced by some factor. For instance, different levels of similarity of some learning material can cause different levels of interference. The level of interference produced could then be measured by retention, but interference would still be the effect that is produced rather than the manipulation that is causing the effect. In other words, interference can be – and indeed has been – used as a label for an independent variable as well as a dependent variable. In the present investigation this distinction is critical because manipulations usually referred to as different levels of interference in the literature is suggested to interact with a third variable (learning speed) in its affect on retention. Suggesting that one level of interference causes different levels of interference depending on the level of the third variable would not be very informative. However, because referring to manipulations (i.e., independent variables) as interference is so common in the literature that adhering to a strict consistency was found very challenging in the present report. This is especially true when referring to prior studies. However, when presenting the experimental design and results of the present investigation, some degree of consistency has been approached by only referring to interference as the effect that is produced and not as an independent variable causing an effect on memory.

Computational Models of Cognition

The relationship between learning speed and forgetting has a long history in psychological research, but has rarely been directly connected to results on model simulations. In general, the focus of psychological research using model simulations has mainly been on building models and running simulations that can account for and possibly explain empirical results from behavioral studies. It is less common to see literature and studies addressing the inverse relationship. In other words, results from behavioral studies generally precede model creation rather than conducting behavioral studies based on results from model simulations. This order is of course the most natural and rational way of conducting psychological research. However, if models of psychological phenomena are to be valuable they should not only describe behavioral results but also be able to provide explanations as well as novel predictions for human behavior. Models attempting to account for the relationship between a low number of variables by a mathematical function can be of great value by providing predictions that can be empirically tested. However, since such models are purely mathematical in nature, they are rarely capable of producing predictions for variables outside its intended scope. More specifically, the output of the model is limited to the variables that goes into it. In contrast, models attempting to simulate the mechanisms behind cognition and behavior, rather than the functional relationship between variables, can have a much greater area of impact. One such modeling framework is (artificial) neural networks. Neural networks are based on a small set of principles relating to the biology of cognitive functioning. The most influential neural network approach within psychology is parallel distributed processing (PDP), also known as connectionism, as suggested and formalized by Rumelhart, McClelland and the PDP Research Group (1986). Because neural network simulations are based on a framework inspired by the low level construction and mechanisms of the central nervous system, such simulations can provide results that were not part of the initial research question but still have the potential to generate novel hypotheses for biological neural networks. So even though neural network simulations are usually conducted to address a specific mechanism in relation to a clearly defined hypothesis, such results can have a much wider impact if one is willing to accept the possibility that this framework has at least some bearing on the mechanisms of biological networks. The most common approach in this area of research has been to assess the validity of this framework by evaluating how well the models can simulate behavioral results. If the direction is turned around, two important reasonings can be made. First, if findings initially considered as simulation artifacts are found to be general across studies, they could be used to build novel hypotheses about human cognition and behavior. Second, by investigating such hypotheses experimentally, the models behind the hypotheses can be evaluated based on the empirical results of the experiments. If the results of such experiments are found to support the hypotheses, it should be considered as evidence in support of the notion that such models do indeed serve an important role in research on human cognition and behavior. In fact, such evidence should have a stronger impact than studies attempting to account for behavioral phenomena by modeling the functional relationship between behavioral level variables. This is because such hypotheses would be based on intrinsic properties related to the very nature of the models rather than being dependent on manipulable model parameters. The present study is inspired by and attempts to build on one such general finding from neural network simulations. When simulations are run on neural network models using distributed and overlapping representations, it turns out that fast, sequential learning – caused by high values on the learning rate parameter and massed repetitions of input presentations – is associated with more interference than slow, interleaved learning. This finding is highly related to the notion of "catastrophic interference" in neural networks and will therefore be further introduced within this context.

Interference in Neural Network Models

Shortly after neural networks became popular and widespread among cognitive scientist, an important and almost devastating finding was made. When networks were trained to learn associations of input–output representations using the same units and weights to represent different associations (i.e., overlapping representations), sequential presentation of new associations almost completely destroyed previously learned associations (McCloskey & Cohen, 1989; Ratcliff, 1990). Since the same units and weights are used for different representations, changing these weights when learning new associations will therefore interfere with the network's existing knowledge. Hence, new learning would cause strong levels of retroactive interference on previously learned associations. The level of interference was in fact so strong that comparison with interference effects in humans seemed meaningless. This motivated some researchers to deem the neural network approach inappropriate as a tool for investigating human learning and memory. Instead of abandoning the concept, several researchers carried out a great deal of research on catastrophic interference attempting to show that neural networks are in fact able to account for human performance after all. Such research has provided a lot of technical knowledge of how and why interference effects are so “catastrophic” in neural networks and what can be done to avoid such high levels of interference (French, 1999; Lewandowsky & Li, 1995). This research has provided more biologically plausible learning algorithms and simulation results comparable to human performance, but more importantly it has also provided knowledge of a few common causes of interference that are more or less independent of the different learning algorithms. O’Reilly and Munakata (2000) have emphasized four important factors affecting interference in neural networks: Sequential learning; overlapping representations; weight scaling for context input; and learning rate. If the assumption is made that human memory interference is caused by the same mechanisms that cause interference in neural networks, novel hypotheses can be built for human performance. So while catastrophic interference has been considered by some researchers to be the Achilles’ heel of connectionism (Lewandowsky & Li, 1995), such results can also have the ability to motivate and enable scientific progress. In the present study, novel hypotheses are inspired and partially built on results from neural network simulations. More specifically, can the factors identified as important for interference in neural networks be connected to factors involved in human learning and memory mechanisms? In order to explore this possibility and formalize specific hypotheses, several areas of research on human learning and memory must be considered. However, before introducing theoretical aspects and empirical results from behavioral studies, a few more details about the factors causing interference in neural networks will be introduced.

Sequential learning. The order of input presentation is a fundamental factor when it comes to interference in neural networks. Since the same units and weights are used to represent different associations, the update of weights when learning a new association will necessarily interfere with previously learned associations. In general, this is true for both sequential and interleaved presentation of input-output associations as long as the network uses overlapping representations. The reason why the order of input presentation matters is related to the internal representations of the network. When inputs are presented sequentially, the internal representations of the network are shaped by learning to discriminate between items of the current list. Interleaved learning of all associations will, in contrast, shape internal representations in a way that makes the network discriminate between all items to be learned. As a result, interference resolution between items is more strongly in effect during interleaved learning than sequential learning. In humans, interference effects have usually been shown in sequential learning procedures as well. For instance, the classical AB-AC paired associates list learning task (Barnes & Underwood, 1959) use such a setup. First, participants are asked to learn a list of word associations, AB, where A represents a set of words to be associated with another set of words, B. After the AB associations are learned, a new list of associations, AC, is learned. In this second list, A are the same words as in AB, but C is a set of novel words. After the AC list has been learned, participants are tested on both lists. Experiments using the AB-AC paradigm have shown that increasing the number of learning trials on the AC list causes increasing levels of interference on the previously learned AB list (Barnes & Underwood, 1959). But even though the neural network simulations show the same overall pattern, the problem is that neural networks do not show a smooth and gradual effect of interference, as humans do, but the effect is rather catastrophic and immediate (Figure 1). Nevertheless, one essential parallel of behavioral studies and neural networks is the order of stimuli/input presentation. In the present investigation, sequential learning will be treated as a precondition for producing interference rather than a factor of interest.

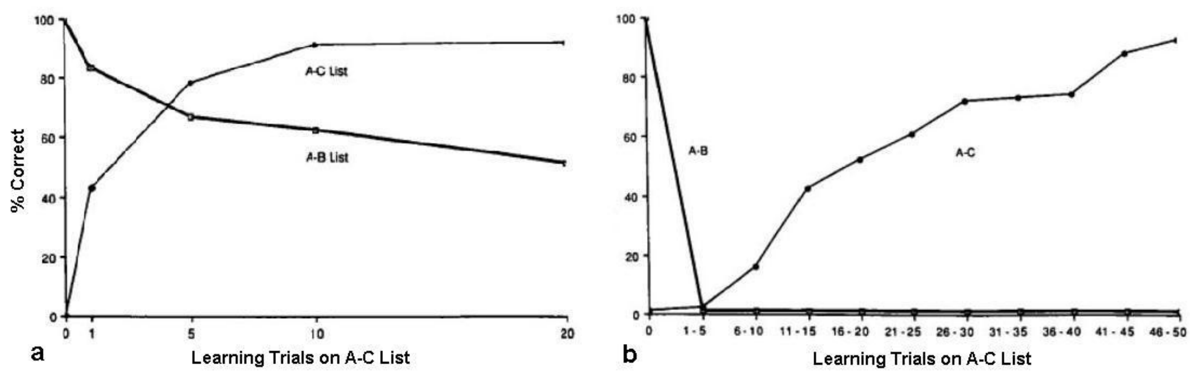

Figure 1. (a) Human performance on the AB-AC task in Barnes and Underwood (1959). Percentage correct as a function of number of trials on the AC list. (b) Models shows catastrophic level of interference when simulating the AB-AC task (McCloskey and Cohen, 1989). Copyright (a) 1959 by the American Psychological Association, Inc.; (b) 1989 by Academic Press, Inc. Reprinted with permission.

Distributed and overlapping representations. Much of the research that followed the original catastrophic interference finding consistently pointed to overlapping representations as the main cause for interference (French, 1992; McRae & Hetherington, 1993; for a review, see French, 1999). The basic idea is that interference is a result of re-using the same units and weights for different associations. While reducing overlap between representations can be accomplished by a number of different techniques, the problem with this is that using sparser, separated and less overlapping representations would contradict some of the most appealing qualities of neural networks as models of human cognition. Overlapping representations is what enables neural networks to effectively generalize from prior learning to novel associations (Hinton, McClelland & Rumelhart, 1986; McClelland, McNaughton, & O'Reilly, 1995). In addition, overlapping representations also provides a more storage efficient way of representing knowledge since several associations can be represented by the same units and weights (Hinton et al., 1986). What follows from this, is that the degree of representational overlap constitutes a tradeoff between generalization and storage capabilities on one side and vulnerability to interference on the other. When moving away from overlapping representations towards more separated representations with less overlap, an important question becomes apparent: Does this mean that neural networks are no good as models of human cognition or does it simply reflect a tradeoff that biological networks also must deal with? Several proposals of how the human brain has evolved to deal with this tradeoff have been suggested. One influential proposal suggests that the brain deals with this tradeoff by having different representational organizations in different brain structures (McClelland et al., 1995). By this account, memories are initially stored by sparse representations with low overlap in the hippocampal system, while remote memory accumulates by building dense and overlapping representations through small synaptic changes in posterior neocortex. This way, the hippocampal system supports fast and relatively interference free learning of recent experiences, while the posterior neocortex slowly discover general structures over a number of experiences. The initial storage of memories involves widely distributed representations in the neocortical system, but it is the conjunctive representations of the hippocampal system that supports the rapid and relatively interference free learning. McClelland and his colleagues primarily relate their proposal to theories of consolidation. However, their work also addresses the role of interference in learning and memory. They argue that the hippocampal system exist to allow retention of acquired knowledge about specific experiences without interfering with existing knowledge in the neocortex. The appeal of such an approach is that interference is strongly reduced for recent learning without completely forsaking the long term benefits of overlapping representations.

In addition, this account is not only consistent with what is known about the hippocampal system and its role in learning and memory (e.g., Squire & Wixted, 2011), it also provides an account of *why* there are separate but complementary learning systems in the hippocampus and neocortex. The slow and gradual change in dependence on the hippocampal system over time, known as consolidation, could be a direct consequence of the tradeoff associated with overlapping representations rather than being some arbitrary property of the biological network of the nervous system. Either way, a neural network account of cognition does anyhow seem to implicate overlapping representations when it comes to questions of interference. While the level of interference can be reduced by different techniques – including reduction of representational overlap – it seems that some degree of overlap is required if interference is to be produced. In other words, if there is no overlap between representations, there is no reason for interference to be present.

Weight scaling for context input. The third factor that is important when it comes to interference in neural networks is context. O'Reilly and Munakata (2000) have provided a simulation example of the AB-AC task of Barnes and Underwood (1959). In this paradigm, successful retrieval will depend on whether the participants are able to connect the learned associations to the correct list. And since the lists are presented sequentially, larger differences in learning context between the lists should improve the participants ability to retrieve the correct association and successfully disambiguate between associations learned in the different contexts. Increasing the relative impact of the context will therefore support such disambiguation. In O'Reilly and Munakata's simulation, the learning context was represented by a separate input layer. But when the network formed input-output representations of the learning material, the internal representations were shaped by integrating both input items and input context. The input context patterns were similar within lists but still slightly different for each item on the same list. This was intended to reflect the notion that while the external environment might be constant, the internal perception changes as time is passing. Increasing the weight scaling parameter for context inputs strongly reduces the onset of the catastrophic between list interference. This makes sense because increasing the impact of context inputs that have between lists differences should encourage the network to form internal representations that depends more on properties that differ between lists. However, what is really happening at the level of the internal representations is that the network uses more non-overlapping representations. Which technically only means that changing the scaling parameter for the context input is effectively the same as using sparser and less overlapping representations, as already pointed out in the previous section.

Learning rate. The final and, for the current study, most important factor, is the learning speed of the network. By changing how strongly the weights of the network are updated on each learning trial, usually controlled by the learning rate parameter, the speed at which a network acquires input-output associations can be manipulated. Increasing the learning rate generally speeds up the learning process. However, when comparing faster learning rates with slower rates, it turns out that the faster learning rates produce more interference (O'Reilly & Munakata, 2000). While this may seem counter-intuitive when thinking in terms of human performance, in the neural network framework this makes perfect sense. In order to speed up learning, larger weight changes are performed on each trial, so when the AC items (i.e., the second list of associations) on the AB-AC task are learned, these large updates will also undo the AB learning faster. Does this mean that fast learning should always be associated with fast forgetting? Results from studies addressing the relationship between learning speed and amount retained in humans are in fact indicating the complete opposite. Fast learners seem to retain more than slow learners over time. It is important to note that results have been ambiguous, and because of this the results of the human studies will be introduced more properly in a separate section below. From a neural network point of view, such a simple relationship between learning rate and forgetting of previously learned associations is too simplistic. The role of distributed and overlapping representations, introduced above, must also be considered when addressing this relationship. Catastrophic interference is only problematic when representations are overlapping. Hence, there is an interaction of learning rate and degree of representational overlap. Stated differently, the effect of learning rate on forgetting depends on the degree of overlap between representations. If the different associations learned by a neural network are represented by different units and weights (i.e. no overlap), there will be no need for updating the weights that represent previously learned associations when novel associations are learned.

Taken together, the four factors emphasized by O'Reilly and Munakata (2000) boils down to two important factors: Overlapping representations and learning rate. Sequential versus interleaved presentation could also be considered a potential moderator of the relationship of interest, but for the sake of simplicity and because both prior modeling and human studies have usually treated this factor as a precondition rather than a moderator, it will not be further investigated in the present study. The importance of context should not be neglected either, but since it can be shown that its role in interference in neural networks is technically the same as varying the degree of overlap at the level of representations, it is not a critical prerequisite for drawing a parallel between models and humans. If it is assumed that neural network models do share some fundamental properties with biological networks, one could use the identified relationship to make predictions about human cognition. In order to do this, the human variables corresponding to the variables of interest in neural networks must be identified.

When it comes to degree of overlap between representations, the theoretical transfer is straightforward, but operationalized measurements necessary for experimental procedures are more problematic. However, by including the assumption that the degree of representational overlap is correlated with the degree of feature similarity between stimuli, it should be possible to construct a variable suitable for experimental use. This assumption is not necessarily unproblematic and will therefore be more properly treated later on.

The other factor, learning speed, is more easily operationalized, but is more problematic theoretically. This is because different learning speeds, both between and within individuals, could be influenced by a number of different factors. The learning rate parameter in neural network simulations corresponds most naturally to any low level biological differences affecting the strength of synaptic changes. However, while it would be reasonable to assume such differences between individuals (and perhaps also within), attempting to connect such low level differences to an operationalized measurement at the behavioral level could be problematic. This is because behavioral level measurements of learning speed will be influenced by a range of other factors, such as existing knowledge, motivation, attention, and learning strategies. The reason why people acquire new knowledge at different rates, as well as why each individual acquire new associations or units of knowledge at different rates, is most likely a complex combination of these factors. Motivation, attention, and learning strategy will necessarily be strongly guided by and interact with existing knowledge. Because of this, a hard differentiation between these potential sources of learning speed will probably be overly simplified. However, by looking at how each of them could be expected to affect learning speed, more nuanced predictions can be made. These factors will be considered within the context of prior research on the relationship between learning speed and forgetting.

Speed of learning and amount retained

Does rate of acquisition predict rate of forgetting? Research aimed at informing this question has been unable to provide a straightforward answer. Underwood's seminal analysis (1954) of prior research addressed this question by evaluating different methods attempting to adjust for levels of initial learning. Prior to Underwood's analysis the general conception was that fast learners retains more than slow learners over time and therefore will perform better when memory is tested at a later point in time. By arguing that the methods used to equate the degree of learning for fast and slow participants were inadequate Underwood deemed the prior studies inappropriate as evidence in this matter. Underwood developed his own approach in order to more appropriately equate the degree of initial learning between participants. When testing this approach in a set of experiments on learning and retention of paired nonsense syllable lists, Underwood found no difference in forgetting over 24 hours between fast and slow learners. Underwood's approach has been well received and a number of subsequent studies using different methods to equate the initial learning between participants have come to the same conclusion: There is no difference in rate of forgetting for fast and slow learners (Gentile, Monaco, Iheozor-Ejiofor, Ndu, & Ogbonaya, 1982; Schoer, 1962; Shuell & Keppel, 1970; Stroud & Carter, 1961; Stroud & Schoer, 1959).

In more recent years, however, the general opinion seems to have shifted back towards the pre-Underwood conception. A few studies have provided compelling evidence in favor of the notion that fast learners do in fact retain more than slow learners over time, even though the level of initial learning is equated. Kyllonen and Tirre (1988) used an item dropout procedure to ensure equal learning for 685 participants on a paired associates task. A list of name-number pairs were learned in cycles. Once a pair was learned to a criterion it was dropped from the list to avoid overlearning. The dropout criterion varied between conditions from one correct response to two or three successive correct responses. When learning of the initial list was completed, a new list containing the same items - but randomly repaired associations - was learned. Since the items were the same in both lists the second list served as a source of interference to the first list. The strength of this interference was varied between conditions by different dropout criterions in the same way the strength of associations also varied by condition in the first list. Finally, retention and trials necessary to relearn the original list was measured. Kyllonen and Tirre found that item-specific learning speed was a significant predictor of retention and reacquisition speed across all forgetting conditions of their experiment. General learning speed, as measured by an independent test, also predicted retention and relearning on the paired associates task.

The same connection of individual differences in learning speed and subsequent recall has also been reported in a study by MacDonald, Stigsdotter-Neely, Derwinger, and Bäckman (2006). MacDonald et al. investigated several predictors of forgetting using a multilevel modeling technique. 136 participants memorized 4-digit numbers and retention was tested at different intervals ranging from immediately after memorization to 8 months later. In order to control for strategy confounds the participants received mnemonic training prior to memorization. In addition to age and cognitive abilities (episodic memory, perceptual speed, working memory), learning speed was found to predict retention, especially at 24 hours after initial learning. Participants who required more trials to reach the learning criterion forgot significantly more than fast learners. When attempting to explain the discrepancies between their own and past results MacDonald et al. identified a number of factors including the possibility of flawed research designs of preceding studies, statistical procedures used, and selection and psychometric properties of measures (e.g., MacDonald et al. used recall rather than recognition as outcome measure). Considering such factors in connection with the different trends of early and late results, it could be claimed that differentiated weights should be allocated to the different results. In general, later studies always have the benefits of learning from earlier studies when it comes to design issues. Moreover, the same is true for the availability of technical aids such as computers for stimuli presentation and measurement accuracy and more sophisticated statistical procedures. So even though some authors (e.g., Gentile, Voelkl, Pleasant, & Monaco, 1995, p. 185) have called results in line with Underwood's conclusion “a consistent finding”, this question should not be considered settled. If any conclusions are to be made, it should be that fast learners retain more or that there are no differences in amount retained between fast and slow learners. Any suggestions toward the alternative, that slow learners retain more, has little, if any, support in the existing body of research.

The early studies on the relationship between learning speed and amount retained focused mainly on equating the initial learning for fast and slow learners. As a consequence, the reason *why* some people learn faster than others did not receive much attention in these studies. In later years, however, this question has been more carefully addressed. Among the sources investigated are learning strategies, effectiveness of cue selection, and prior knowledge.

Learning strategies. In general, it can be difficult to categorize and make strong distinctions between different learning techniques. However, a number of studies have still shown that the use of specific techniques increase learning speed and subsequent recall when compared to other techniques. The keyword mnemonic, for instance, has been shown to be effective when it comes to learning and remembering second-language vocabulary (e.g., Atkinson, 1975; Desrochers, Gelinas, & Wieland, 1989; Pressley et al., 1980) as well as other types of learning (Bellezza, 1981). The keyword mnemonic is a technique that uses a familiar word or phrase that is acoustically similar to a novel word. By linking the two words with an interactive visual image that includes an object or concept that is easily associated with the meaning of the novel word, subsequent encounters with the novel word will elicit the keyword and provide access to the word meaning through the image. For instance, the Spanish word for food is “alimento”. This sounds a bit like “a lemon toe” in English and by imagining a person who is eating a lemon of his toe, one could make the meaning of the word more easily available. When this keyword method has been experimentally compared to rote memorization, it has been shown that this technique is superior when it comes to learning speed, immediate recall, and delayed recall (e.g., Rosenheck, Levin, & Levin, 1989). Because both immediate and delayed recall has been found to be superior for the keyword technique, it has been suggested that forgetting rates for associations acquired through these two techniques are indifferent. However, Wang, Thomas, and Ouellette (1992) have challenged these findings. Wang and his colleagues argued that the measures of long term retention in the previous studies may have been confounded because they used within-subjects designs to test for differences between immediate and delayed intervals. Because immediate tests of recall could have inflated memory measures at the delayed tests, Wang et al. argued that a between-subjects design would be more appropriate for testing forgetting rates for material learned with the different techniques. By using a 2 x 2 factorial design with learning condition (keyword vs. rote rehearsal) and retention interval (immediate vs. one week delay) as between-subjects factors, Wang et al. conducted a series of experiments to test their assumptions. At immediate tests of recall, the keyword technique was found to be superior to rote memorization in all experiments. Effectively, replicating the results of previous studies and providing further support for the notion that the keyword technique should be associated with faster learning. However, in contrast with previous findings, forgetting rates of the different techniques were different. Condition x Time interactions were significant in all experiments indicating stronger forgetting over time for the keyword condition (Figure 2). Interestingly, by interpreting the time factor as equivalent to low and high levels of retroactive interference on the learned material, these results are in accordance with the predictions of the modeling results presented above.

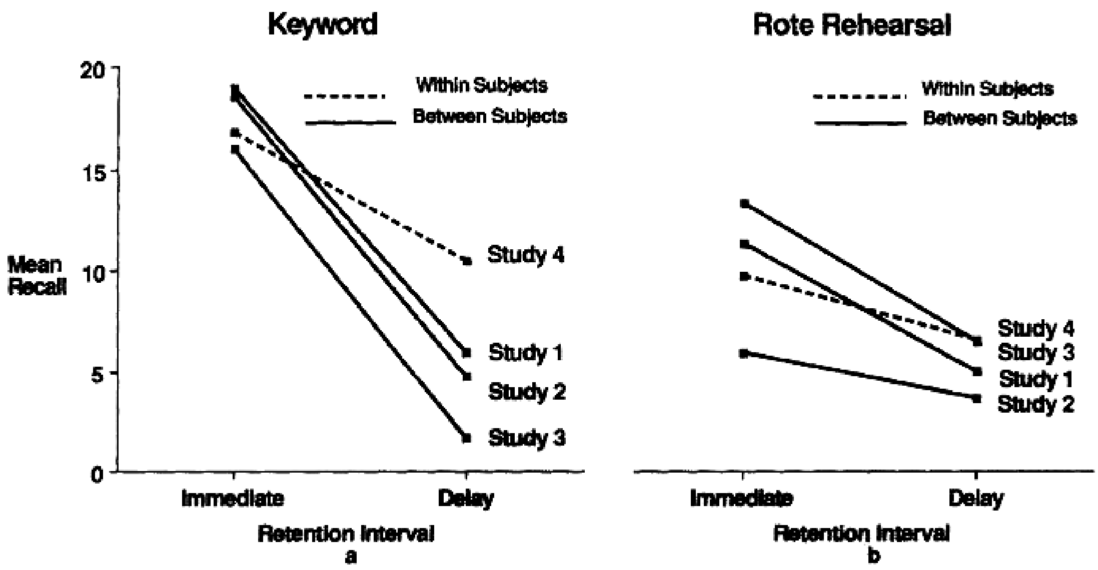

Figure 2. Mean recall as a function of learning strategy for 2nd-language vocabulary words of the four experiments of Wang et al. (1992). All experiments indicated stronger forgetting over time for (a) the keyword mnemonic compared to (b) rote rehearsal. Copyright 1992 by the American Psychological Association, Inc.

Differential forgetting over time for different learning techniques has been replicated by Wang and Thomas (1992). In a study investigating long-term retention of Chinese characters learned by different techniques, Wang and Thomas found a significant interaction of learning condition (mnemonic vs. rote learning) and time (immediate vs. two-day delay), with a stronger decline in retention over time for characters learned by the mnemonic method. Recall was superior at immediate testing for participants using the mnemonic technique, while rote memorization was found to be more robust over time, producing higher recall at the delayed test when compared to the mnemonic group.

Even though the studies on the effect of learning strategies on rate of acquisition and forgetting are somewhat inconsistent, the results of Wang et al. (1992) and Wang and Thomas (1992) suggests that the use of different strategies could lead to different rates of forgetting. Why this happens is less clear. Wang et al. (1992) suggested that

> preexperimental associations of the native-language keyword may interfere with retrieval of the desired interactive images over time. Although the keyword serves initially as an effective retrieval cue, preexperimental associations may regain their prominence over time and hinder retrieval of the mnemonic image. (pp. 526-527)

Stated more generally, the keyword mnemonic relies on a strong, yet single, cue for retrieval. Successful retrieval of the learned association would therefore be highly dependent on this one cue (i.e., the keyword), rendering the association more prone to interference. The slower learning of the rote memorization could, in contrast, lead to several weaker but perhaps more distributed paths to retrieval. Which would mean that retrieval could be harder in general, but the association would be more robust to interference due to its distributed nature. Such reasoning strongly parallels the arguments of the encoding variability hypothesis (e.g., Martin, 1968) and will be further introduced within this context later on. If, however, such an interaction of learning strategy and time is related to differential vulnerability of retrieval paths, it would probably be insufficient to only consider the interference of preexisting associations. Interference caused by subsequent learning could have the same effect on this relationship, and investigating this possibility by experimental manipulation of interference rather than time could therefore be an interesting extension of this path of research.

Cue selection effectiveness. Another potential source of individual differences in learning speed is effectiveness of cue selection. This area of research is a narrow field within experimental psychology. However, since such research has provided some counterintuitive – yet consistent – results related to speed of learning, a brief introduction will be included.

Efficiency of cue selection has been defined as the degree to which a response is learned to only one component of a compound stimulus (Richardson, 1971, 1973). In a series of experiments, Richardson (1973) studied the extent to which efficiency of cue selection contributes to increased speed of learning. It was hypothesized that cue selection efficiency, among other factors, would be related to fast learning. However, contrary to what was expected, slow learning was consistently found to be associated with higher levels of cue selection efficiency. The experimental stimuli consisted of consonant trigrams and was learned to be associated with a digit response. After the paired-associate list was learned to a certain criterion, the individual letters from the trigrams were presented individually and tested for recall of the corresponding digit. Efficiency scores were computed by dividing the number of different correct recall responses by the total number of correct recall responses.

> The efficiency score is the percentage of the total correct recall responses that would be necessary for S to give a single correct response to each trigram stimulus which is represented by a correct recall to one or more of the three stimulus letters. (Richardson, 1973, p. 398)

Results similar to those of Richardson (1973) has also been found by Parsons (1968) and Postman and Greenbloom (1967). A later review on component selection by Richardson (1976) concluded that the finding that slow learners select more efficiently than fast learners is a consistent finding. How should such a finding be interpreted? Richardson (1976) offers few clues toward any explanations. However, another study by Richardson (1972) has indicated that selection efficiency can be increased by emphasis and instructions, suggesting that fast learners are as able as slow learners to select efficiently but they usually do not.

The term “functional stimulus” refers to the notion that humans often select a part of a stimulus and make an abstraction of it (e.g., Richardson, 1971). Shepard (1963) suggested that some stimuli are perceived as unitary and unanalyzable wholes, while others are analyzed into components and dimensions. Because it seems reasonable to assume individual differences in how stimuli are perceived and what parts of stimuli are made functional, it could be suggested that such differences is part of the explanation for why slower learners are more efficient at cue selection. If fast learners are faster because they more easily make meaningful abstractions of stimuli wholes, then this could explain why the fast learners have been found to be less efficient in studies operationalizing cue selection efficiency by the degree of learning one concrete subpart of a compound stimuli. Consideration of such an assumption in relation to the role of interference as a moderator of the relationship between learning speed and forgetting calls for some interesting remarks. If abstractions of higher level attributes of stimuli are associated with faster learning, then the reasoning of the previous section could again be applied. A stimulus that is encoded by making an abstraction of its high level attributes could, at least in some cases, be more prone to interference when compared to encoding based on a specific part of the stimulus. This would of course depend on the nature of the potential interference. However, in order to show that this assumption would be reasonable for at least some types of interference, an example is provided. The consonant trigram *bcd* could be encoded both as a high level abstraction such as “the first three consonants of the alphabet” or by making a part of it functional, for instance “starts with the letter b”. If subsequent learning then interferes by associating a similar stimulus such as *dcb* with a different response, the high level abstraction would clearly become more ambiguous than the association emphasizing the single component of the stimulus. An example of the opposite case can easily be constructed as well, however, the point is not to make a case for general prediction but rather to show one possible explanation for the cue selection effectiveness results as well as informing the relationship of interest in the present investigation. When such an effect is likely to be present depends strongly on how stimuli are encoded as well as which features of the interfering stimuli are similar. Martin (1968) and Shepard (1963) have emphasized the importance of stimulus meaningfulness for encoding, which suggests that not only physical similarity of stimuli are of importance but also how stimuli are perceived and interpreted in a meaningful context. Because encoding will depend on how subjects actively perceive and organize the stimuli, the role of existing associations and prior knowledge must be considered.

Prior knowledge. Existing associations and knowledge is obviously important for how new knowledge is acquired. If prior knowledge is the main source of the differences in speed of learning, an interaction of speed and interference is not necessarily predicted. This is because the general effect of interference should be strongly diminished for participants with strong prior knowledge of the learning material. Prior knowledge will guide motivation, attention, strategy, and cue selection, and the more familiar the to-be-learned material is, the more likely it is that the learning will be faster (Booth, Koedinger, & Siegler, 2007; Martin, 1968; Richardson, 1976). Moreover, it could be argued that it is more likely that attention will be directed to features of the learning material capable of resolving potential interference with prior learning. Hence, if such effects are assumed there seems to be no reason to predict accelerated forgetting for fast learning when interference is high. In other words, if prior knowledge is the primary reason for differences in learning speed the interaction predicted by computational models will probably not hold. However, prior knowledge can also be considered in a different way. Martin (1968) assumed that encoding variability is inversely related to stimulus meaningfulness. If stimuli are meaningful, Martin argued that there is a higher probability that the same encoding will be used for both first and second list learning in the AB-AC paradigm, producing negative transfer and retroactive inhibition. In contrast, less meaningful stimuli will not have one specific encoding readily available and will therefore have a greater probability of being associated with different encodings of the first and second list. By taking into account that prior knowledge and perceived meaningfulness of stimuli are closely related, the role of prior knowledge does not seem to offer any straightforward predictions for the relationship of interest in the current investigation. However, even though Martin (1968) was mostly concerned with stimulus meaningfulness and paired-associate transfer, his encoding variability hypothesis has later been extended and still serves an important role today.

Encoding Variability

The spacing effect, first demonstrated by Ebbinghaus (1885/1913), is a thoroughly studied and replicated finding in psychology. It refers to the fact that memory is superior when learning repetitions are distributed and interleaved rather than when repetitions are massed and adjacent in time. The spacing effect has been shown for various learning materials and memory tasks (for reviews and meta-analyses, see Cepeda, Pashler, Vul, Wixted, & Rohrer, 2006; Delaney, Verkoeijen, & Spirgel, 2010; Dempster, 1989; Donovan & Radosevich, 1999; Greene, 1989; Janiszewski, Noel, & Sawyer, 2003; Raaijmakers, 2003) and many theories and models have attempted to account for the empirical results. Unfortunately, even though the spacing effect is a robust and well studied phenomenon, there is still disagreement among researchers about how this effect should be explained (Delaney et al., 2010). An extensive review of the different explanations that have been suggested is outside the scope of the present paper. However, one of the proposed explanations for the spacing effect seems to be highly relevant for the relationship under investigation in the present study and will therefore be introduced.

While it is interesting to note the striking parallel of the distributed and interleaved nature of the spacing effect and the importance of these same properties in neural network simulations, the connection between the spacing effect and learning speed is perhaps not very intuitive. However, by considering one of the proposed mechanisms underlying the spacing effect, an important connection can be made. The encoding variability hypothesis (e.g., Martin, 1968; Melton, 1970; Richardson, 1976) has emphasized the importance of encoding variability for subsequent memory tests. The hypothesis suggests that subsequent memory will benefit from higher variability of the features or context encoded across repeated exposures to a stimulus. Stated differently, if a stimulus is processed differently on each exposure, subsequent memory will be better when compared to stimuli with more similar processing across exposures. As mentioned above, this hypothesis is only one of a number of suggested explanations for the spacing effect. There are also different versions of the hypothesis emphasizing different aspects of importance (e.g., meaningfulness, context, components or features). However, with respect to the spacing effect, a recent review has emphasized encoding variability as one of the most appealing accounts in the literature (Delaney et al., 2010), even though it is not able to fully account for the empirical data (Delaney et al., 2010; Greene, 1989).

Studies on the spacing effect which includes learning speed as a factor are not common. However, the operational definition of learning speed makes it straightforward to connect it to the concept of encoding variability. Slower learning is usually defined by a higher number of trials-to-criterion or longer stimuli exposure times required to establish subsequent recall or recognition. A higher number of learning trials or longer exposure time will allow, perhaps even encourage, a higher level of processing variability. Since slower learning, by definition, will include a higher number of trials or longer temporal exposures, it could be suggested that memory representations established by slow learning will be more distributed when compared to fast learning. If the representations resulting from fast and slow learning are in fact different, the reasoning can be extended further in order to make predictions for how this could be related to the interaction predicted by the neural network models. The more distributed representations resulting from slow learning would depend on a higher number of encoded features or retrieval paths when compared to representations established by faster learning. Because slow learning in general seems to be associated with inferior memory when compared to fast learning, the strength of each of these retrieval paths should be relatively weaker than the low numbered paths of fast learning representations. However, if initial learning and the subsequent memory test is interrupted by learning which has similar features to the initial learning, the representations established by fast and slow learning could be differentially affected by the intervening interference. The more distributed representations of slow learning should be more robust to interference. This is because they depend on a higher number of retrieval paths which means that the probability that all available paths will be interfered with is lower when compared to the narrower but stronger representations of fast learning. Importantly, this would only be true when the level of interference is defined by an overall degree of stimulus similarity. If specific features of the stimulus are being interfered with, the effect of the interference on the initial learning would depend on which features were attended to during the encoding. If the features attended to enable discrimination between the initial learning and the interference learning, the superiority of fast learning would be further strengthened rather than reduced. So even though fast learning might depend on a low number of strong paths for retrieval which generally could be more prone to interference, what is actually learned and interfered with is essential for making predictions about the outcome.

What is interesting about the above reasoning is that the assumptions of the general case are in line with the predictions made by neural network models. If slow learning leads to more distributed memory representations and distributed representations are generally more resistant to interference, it follows that interference should be a moderator of the relationship between learning speed and forgetting. However, before formally stating any hypotheses there is one more factor that needs to be considered. The factors emphasized by neural network models included overlapping representations as important for interference effects. In fact, as stated earlier, overlapping representations is the very reason why interference happens in neural networks in the first place. Whether this also is true for interference effects in human memory is still an open question. However, since this factor has been found to be important in neural network models and the present paper investigates hypotheses built on such results, it will be important to consider this factor in humans as well.

Distributed and overlapping neural representations

The question of whether knowledge is coded in a distributed or localist manner in the brain is still debated among researchers (e.g., Bowers, 2009; Plaut & McClelland, 2010). The localist account argues that words, objects, and concepts are coded distinctly by non-overlapping representations (Bowers, 2009). It is related to the notion of "Grandmother cells" (e.g., Gross, 2002) and is largely based on empirical results from single-cell recording studies (Bowers, 2009, 2010). However, it is important to note that localist accounts are generally more concerned with emphasizing the presence of localist representations, rather than arguing for a total absence of distributed representations (Page, 2000). This point is essential for the current investigation because the predicted interaction does not rely on a fully distributed and overlapping account. It merely requires that representational overlap exists. The idea that representations are distributed and overlapping in the brain is a core principle of the PDP approach (Bowers, 2009; Seidenberg, 1993), but it is not an important principle for neural network theories in general (Bowers, 2009; Feldman & Ballard, 1982).

Imaging techniques currently available suffer from low spatial resolutions when it comes to informing the question of representational overlap. However, a recently developed statistical analysis technique of functional imaging data is promising in this matter. While traditional imaging studies have primarily focused on the differential contributions of distinct neural structures, novel analysis techniques targeting distributed activation patterns have gained increased interest among researchers during the past decade (Rissman & Wagner, 2012). More precisely, multivoxel pattern analysis (MVPA) is capable of decoding information represented within distributed activity patterns of functional imaging data. Such a technique has potential to inform both the question of representational overlap as well as unresolved questions related to the encoding variability hypothesis.

Activation pattern similarity. Haxby et al. (2001) used functional magnetic resonance imaging to measure patterns of response in ventral temporal cortex while participants were looking at images of different categories. The categories included faces, cats, different categories of man-made objects, and nonsense pictures. The distinct patterns of neural response for the different categories successively predicted which category was being viewed. More interestingly, the categories could also be identified even when the regions that responded maximally to a category were excluded from the analysis. Furthermore, response patterns enabling discrimination between all categories were also found within the regions that responded maximally to only one category. Haxby and his colleagues concluded that representations in ventral temporal cortex must be widely distributed and overlapping. Similar results have also been found for category processing in ventral and dorsal occipital cortex and in superior temporal sulcus (Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999; Ishai, Ungerleider, Martin, & Haxby, 2000) as well as for within category exemplar classification in lateral occipital complex (Cichy, Chen, & Haynes, 2011; Eger, Ashburner, Haynes, Dolan, & Rees, 2008). Even more convincing are the results of O'Toole, Jiang, Abdi, and Haxby (2005) who demonstrated that shared image-based attributes is a factor driving neural similarity. In other words, similar object categories shares more voxels than dissimilar categories because the similar objects share more attributes. These results, however, should be nuanced by results indicating that the regions of the human visual system are neither interchangeable nor equipotential and contains regions that are primarily involved in the analysis of single class stimulus (Spiridon & Kanwisher, 2002). Even though the debate is not completely settled (e.g., Tong & Pratte, 2012), recent studies (e.g., Ewbank, Schluppeck, & Andrews, 2005; Haxby et al., 2001; Shohamy & Wagner, 2008) and reviews (Martin, 2007; Rissman & Wagner, 2012) are converging on a conclusion in favor of representational overlap. However, because there seem to be some degree of inconsistency of the empirical results, perhaps a graded theoretical middle ground is a more appropriate approach (e.g., Plaut & Behrmann, 2011).

Pattern analysis and encoding variability. As mentioned above, one of the suggested explanations for the spacing effect is the encoding variability hypothesis. Because MVPA attempts to account for the fact that representations are distributed in the brain it has been employed as a tool for investigating the validity of the encoding variability hypothesis. Xue et al. (2010) applied representational similarity analysis to fMRI data in a series of experiments.

The experiments included memorization tasks for face recognition and word recall with varying interrepetition intervals (1-20 consecutive repetitions) and repetition lags (4–9 trials). Pattern similarity analyses revealed that better subsequent recognition and recall was associated with greater similarity between neural activity patterns across repetitions. This would indicate that successful episodic memory encoding occurs when the same neural representations are reactivated across study repetitions. Because the encoding variability hypothesis, in contrast, would predict that successful encoding and subsequent retrieval should rather be associated with greater dissimilarity between activity patterns across study presentations, such results appears to be in contradiction to the hypothesis. However, before making any conclusions, a few considerations are in order. As pointed out by the authors, “fMRI data are a relatively coarse aggregate measure of the responses of large populations of neurons and, thus, may not necessarily capture all of the aspects of encoding variability that might be at play” (Xue et al., 2010, p. 100). Furthermore, the fact that similar processing across study repetitions support memory retention in general should be no surprise (e.g., Kahn, Davachi, & Wagner, 2004; Nyberg et al., 2001), but the simple conclusion that such results are in disfavor of the encoding variability hypothesis does not follow from this. The results of Xue et al. (2010) does not favor a strong version of Martin's (1968) original notion of encoding variability which was primarily concerned with stimulus meaningfulness. However, if the concept of encoding variability is extended to include context (Melton, 1970), component selection (Richardson, 1976), or a combination of variability and retrieval (Delaney et al., 2010) the approach used by Xue et al. is only capable of supporting the encoding variability hypothesis, it can not disconfirm it. More generally, investigations of encoding variability by pattern analysis would be more informative if a discontinuous similarity function is assumed. Up to the point of successful reactivation of representations of previous study trials, increasing levels of pattern similarity should predict better retention. While better subsequent memory due to encoding variability would only predict greater dissimilarity between activity patterns across study presentations once a certain threshold has been met. Such a threshold would represent successful reactivation and be necessary for connecting subsequent encoding to representations established at preceding presentations. However, once such a threshold has been met, the encoding variability hypothesis would predict that greater dissimilarity between activity patterns on following presentations should lead to more distributed memory representations supporting retrieval at subsequent memory tests. The preceding reasoning is supported by research on study-phase retrieval and its role in explaining the spacing effect. Study-phase retrieval is simply a fancy term for "recognizing that something is repeated when you see it" (Delaney, 2010, p. 91) and was first proposed by Hintzman and Block (1973). It has survived a number of thorough investigations (e.g., Appleton-Knapp, Bjork, & Wickens, 2005; Hintzman, Summers, & Block, 1975; Johnston & Uhl, 1976; Paivio, 1974; Sahakyan & Goodmon, 2007; Storm, Bjork, & Bjork, 2008) and is still considered one of the major theories of the spacing effect (Delaney, 2010). In fact, even though there seems to be no single factor capable of explaining the spacing effect, the recent review by Delaney et al. (2010) emphasized a hybrid encoding variability and retrieval account as the most appealing.

Interference as a Moderator

Learning speed, forgetting, and the relationship between them are high level constructs that have usually been operationalized by somewhat crude measurements at the behavioral level. The complex composition of factors affecting learning speed has made it difficult to interpret experimental results and make any strong general conclusions. However, even though the relationship between how fast something is learned and how fast it is subsequently forgotten is complex and therefore makes it difficult to state any general predictions, there seem to be one mediating factor that has not been properly considered yet. Interference has long been known to be important for forgetting in general (McGeoch, 1932; Roediger, Weinstein, & Agarwal, 2010), but is it also possible that this factor is the very reason for the inconsistency of the learning speed results of the past? Different studies have produced different results and suggested different conclusions regarding the general relationship between learning speed and forgetting. If representations resulting from fast and slow learning are differentially affected by different levels or types of interference, then this could potentially explain the disagreeing conclusions in this area of research. Because between-studies differences in tasks and stimuli used have not been properly considered it is possible that different studies using different tasks and stimuli have produced different results simply because different levels or types of interference are produced by the different experimental tasks and stimuli. Such between-studies differences of prior research have not been specifically investigated in the present project and will therefore not be further considered.

It should be noted though, that a few studies have investigated the possibility of differential effects of interference for fast and slow learners (Schoer, 1962; Stroud & Carter, 1961). These studies did not produce significant interactions of learning speed and interference. Interference was manipulated by list length (Stroud & Carter, 1961) and by length of interpolated lists (Schoer, 1962). Both studies predicted that slow learners would be more prone to interference, but no such evidence was produced. The theoretical basis for this prediction was somewhat thin, as evident by Stroud and Carter's justification for their hypothesis:

> If ability differences are associated with resistance to inhibition as just suggested and if this should turn out to be a general phenomenon, then slow learners should be more susceptible to retroactive inhibition effects than fast learners, and should, in terms of widely accepted theory of forgetting, retain less well what they learn. Perhaps all this is a bit tenuous. (1961, p. 30)

Other studies, however, have produced results that provide empirical support for the opposite effect without stating specific predictions.

The main findings of the large-scale investigation of individual differences in learning and forgetting by Kyllonen and Tirre (1988) have already been introduced. However, the results of this study include an additional finding that was not specifically predicted beforehand. In this study, four different strength levels of interference between initial learning and test of retention were included. The different strength levels were introduced by varying the criterion for successful recall and list drop-out across conditions. The criterions included 0 (no interfering list learning), 1, 2, or 3 successive correct recall trials of the items of the interfering list. Importantly, the interference learning did not include different degrees of similarity to the initial learning, but rather varying degrees of associational strength of the items that were specifically designed to interfere with the associations established by the initial learning. At moderate amounts of interference (i.e. Criterion = 2 successful recall trials), the difference in retention for fast and slow learners were less than the difference in the other interference conditions. This interaction was present across all strength levels of the initial learning (Figure 3). The authors did not go to great lengths when attempting to explain the effect, but did provide the following suggestion:

> Assume that an intermediate degree of strength of an interfering trace produces maximal interference (i.e., ones knowledge of Columbus-1492, presumably a very strong trace, does not interfere with new name-number associations). Call this value of

strength *m*. For fast learners, strength = *m* for interfering traces that had been taken to Criterion 2. For slow learners, strength only approaches *m* for traces that had been taken to Criterion 3. Assuming that the criterion manipulation equates for strength initially, the differences in criterion needed to reach strength = *m* at retention time then simply reflect differences between fast and slow learners in forgetting rate. (Kyllonen & Tirre, 1988, pp. 406-407)

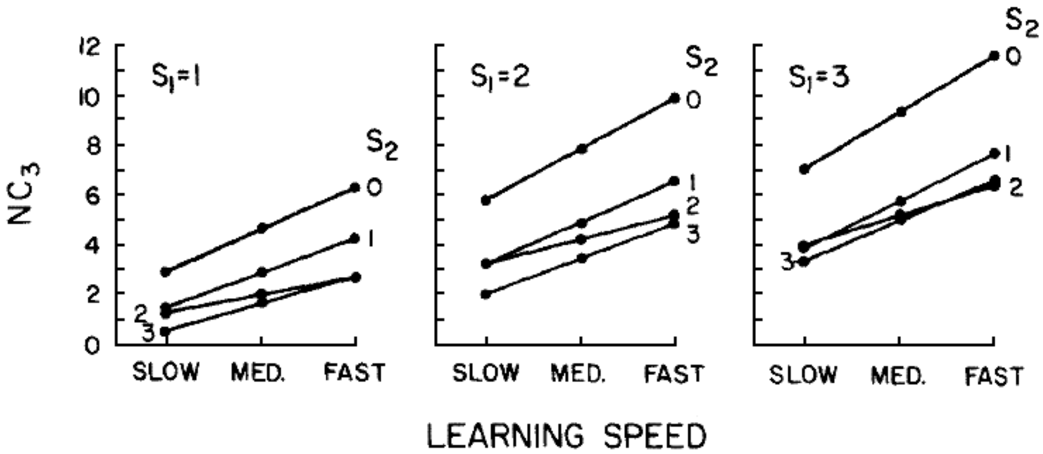

Figure 3. Retention (NC_3: number correct) as a function of learning speed in Kyllonen and Tirre (1988). S_1 (1-3 successive responses) show strength levels of the initial learning, while S_2 (different lines) represent different strength levels for the interfering list. Note the interaction of learning speed and interference at the intermediate level of interference strength. Copyright 1988 by Elsevier, Inc. Printed with permission.

Given the overall results of the study, that forgetting is slower for fast learners, the explanation suggested by the authors is indeed consistent with the data. However, this reasoning seem to build on an additional assumption not explicitly expressed by the authors, that forgetting over time and interference are two separable sources of memory loss. While such an account is consistent with both traditional (for a historical perspective, see Roediger, et al., 2010) as well as more recent (e.g., Hardt, Nader, & Nadel, 2013) theories of decay, it fails to address the possibility that interference can affect fast and slow learning differently. By making the same assumption as Kyllonen and Tirre, “that an intermediate degree of strength of an interfering trace produces maximal interference” (1988, p. 406), an account emphasizing a moderating role of interference between speed of learning and forgetting would also be consistent with these results.

A proper treatment of the classical decay versus interference debate in the forgetting literature will not fit within the limits of the current paper. However, it should be noted that completely neglecting the importance of this question will undermine the role of the empirical results in building the hypotheses of the current investigation. This is true for both the interpretation of the interaction found by Kyllonen and Tirre (1988) as well as for the learning strategy results presented earlier. Nevertheless, these empirical results together with the prediction of neural network simulations as well as the theoretical extensions of the encoding variability hypothesis are all pointing toward a question that has not been properly investigated yet. It is important to note that even though the prediction of neural network simulations and the encoding variability hypothesis are both suggesting an interaction of learning speed and interference (i.e., similarity), the mechanisms behind these predictions are not the same. In other words, the predictions they make are actually quite different but suggest the same behavioral level effect. In general, associations established by fast learning are more prone to interference from subsequent learning when the similarity of the initial and interfering learning is high. Because the encoding variability hypothesis suggest that this effect can be connected to the different numbers of learning trials or exposure times of fast and slow acquisition of associations, the effect is also predicted within as well as between individuals. The following hypotheses are therefore proposed:

1. Fast learners are more prone to interference than slow learners when similarity of the original and interfering learning is high; and
2. Associations learned fast are more prone to interference than associations learned slowly when similarity of the original and interfering learning is high.

Testing these hypotheses experimentally would involve treating learning speed (between subjects and item-specific) and similarity (of the initial and interfering learning material) as independent variables and post-interference retention as the dependent variable. The hypotheses would then predict that retention will depend on the interaction of learning speed and similarity. Stated differently, the relationship between learning speed and retention will be moderated by similarity of the intervening learning material. Testing these hypotheses does not need to make any assumptions with respect to the prior research on learning speed and retention. They simply suggest that interference will increase more for fast learning when similarity increases. Whether fast learning is generally superior to slow learning, as suggested by some prior studies, is not essential. In fact, the inconsistency of prior studies could potentially be explained by the specific stimuli used in the different studies if the between subject hypothesis of the present investigation is assumed. However, if superior retention for fast learners is assumed, it should be noted that the between-subjects hypothesis can still be supported even if retention is better for fast learners across all similarity conditions (F1 in Figure 4 represents such a case). The hypotheses simply suggest that the relative difference will be stronger for fast learning when similarity increases. Hypothetical graphs depicting the hypothesized effect are presented in Figure 4.

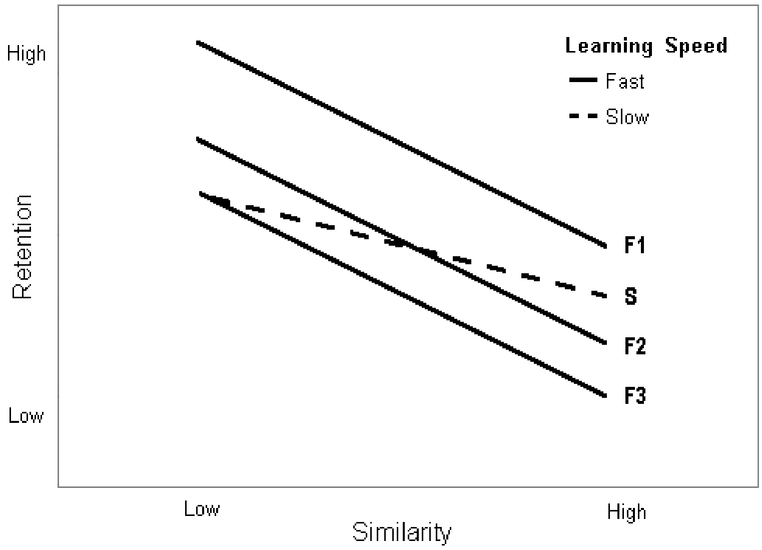

Figure 4. Hypothetical graphs. Downward orientation of slopes represents a general interference effect of similarity. S is slow learning and serves as a reference for the other lines. F1-F3 represents fast learning and shows different ways the hypothesis could be manifested. Whether F1, F2, or F3 is expected will depend on the similarity scale as well as which assumptions are made with respect to the general case (i.e., whether it is assumed that fast learning is generally superior to slow learning or not).

One problem, however, with testing these hypotheses experimentally is that learning speed cannot be randomly assigned and therefore breaks the assumption of independency required for a true experiment. The investigations reported in the following sections are nevertheless referred to as experiments, but are perhaps more appropriately categorized as quasi-experiments or correlational.

In order to test these hypotheses, two experiments were designed and carried out. The two experiments had the same basic design. However, because some of the parameters of the initial experiment were found to be ineffective in establishing a procedure that was sensitive to the hypothesized effects, the second experiment was simply an improved version of the first experiment. A joint analysis including the results of both experiments was also carried out. However, before the combined results are presented, the separate results of the experiments will be reported. The purpose of this is to emphasize the differences between these experiments and relate these differences to the overall results.

Experiment 1

The experiment consisted of learning a number of Chinese characters and their Norwegian translations. First, a set of character-translation associations was learned (L1). Then, a new set (L2) of character-translation associations with high visual similarity to half of the items of L1 was learned. Finally, memory and level of retroactive interference was assessed by a final test of retention of the initially learned associations of L1.

Method

Participants. 22 participants were recruited for the experiment (13 women, 9 men, M_{age} = 21.8, SD_{age} = 2.9, age range: 19 - 28 years). One participant was excluded from the analysis based on a multivariate outlier analysis. Mahalanobis distance (MD) was calculated based on a matrix with learning speed and the retention scores of the different conditions. (Subjects with a significantly deviating MD were excluded). All participants were fluent Norwegian speakers and none had any familiarity with Chinese characters.

Design. Learning speed was used both as a between-subjects and within-subjects factor. Each participant's general learning speed was assessed by mean trials-to-criterion (TTC) on the initial set of associations (L1). Participants were then categorized as either a fast or slow learner depending on whether the mean TTC was lower or higher than the median TTC of all participants. For within-subjects comparison, each association learned was categorized as fast or slow based on its TTC compared to the participant's median TTC. Similarity was manipulated within-subjects by randomly assigning half of the characters learned in L1 to either the high or low similarity condition. For items in the high similarity condition there was a similar Chinese characters – with a different meaning – included in L2 (Figure 5). The dependent variable was retention of the initially learned associations of L1 following completion of the interference learning. This resulted in a 2 x 2 x 2 mixed factorial design with one between-subjects factor (fast vs. slow) and two within factors, item-specific learning speed (fast vs. slow) and similarity (high vs. low).

Figure 5. Example of Chinese characters causing retroactive interference. (a) Was part of the initial learning (L1) and was learned to be associated with the word *alltid* (always). When this character was assigned to the high similarity condition, the interfering list (L2) included (b) which was then learned to be associated with *vann* (water). Learning to associate (a) with one word followed by learning to associate (b) with a different word was intended to produce retroactive interference due to the visual similarity of the characters.

Stimuli and procedure. The stimuli consisted of 40 pairs of Chinese characters and their Norwegian translations (Appendix A). The characters were selected based on their visual similarity to other Chinese characters with a different meaning. The overall stimuli pool consisted of 80 (40 x 2) Chinese characters for which each participant was to initially learn 40 (L1) characters and their meaning before half of these were interfered with by learning 20 (L2) characters with high similarity to the initially learned characters. The initial learning (L1) was split into two sublists in order to control for order effects. When learning a list of items for subsequent memory testing and categorizing the items by how fast they are learned there is a natural confound of time and interference from other items not part of the controlled interference. The experiment used an item-dropout procedure in order to control for overlearning, and because of this, time between learning and testing as well as interference from other items will be higher for items learned fast. The split of L1 into two separated lists allowed some control of this possible confound. After learning of L1 was completed, all participants learned the list containing the interfering items (L2). Half of these were similar to items of the first sublist of L1 while the other half targeted items of the second sublist. For each participant, random selection determined which items went into the first and second list of L1 as well as which items were interfered with. When the interference learning was completed, all participants were tested, in random order, on all items of L1. The flow of the experimental procedure is illustrated in Figure 6.

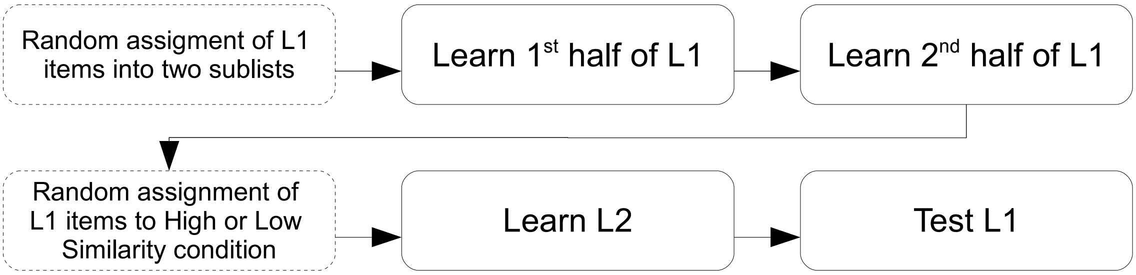

Figure 6. Flow chart of experimental procedure. L1 items were divided into two sublists. When learning of L1 was completed, each list item was categorized as fast or slow based on the number of trials required for learning (TTC). Half of the fast and half of the slow items were then randomly assigned to the High Similarity condition, which meant that a visually similar character would be included in L2. L2 was learned by the same procedure used for learning the sublists of L1 (see text and Figure 7 for details on the list learning procedure). Following learning of L2, retention of all L1 items was tested.

Each learning sequence was carried out by first presenting all items of the list in random order. Each character and its Norwegian translation was presented for 1500 ms and was immediately followed by presentation of the next list item. Between learning trials and test trials of each round, three math tasks were given in order to interfere with working memory strategies and reduce recency effects. The math tasks asked participants to indicate the validity of an equation (e.g., 34 + 43 = 76?). On the test trials participants were prompted, in random order, with the Chinese characters of the current list and were given four options to choose from. Each test trial was time limited and proceeded automatically after six seconds, even if the participant had not responded. The options included three randomly chosen translations of the current list in addition to the correct translation. The four options were displayed with different on-screen spatial orientations (high, low, left, right) and responses were given by pressing the keyboard arrow key corresponding to the on-screen orientation. The position of the correct option was randomly chosen on each trial. When an item was correctly responded to on two consecutive rounds, the item was dropped from the list. This procedure was carried on until all items of the list were learned (Figure 7). The same procedure was used on both sublists of the initial learning (L1) as well as the interference learning list (L2). For items in the high similarity condition, the final test options included the translation of the similar character in addition to the correct and two randomly chosen translations.

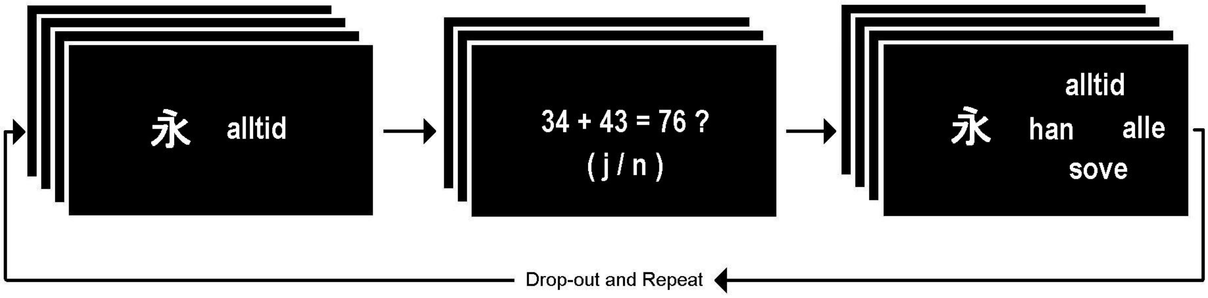

Figure 7. List learning procedure used in Experiment 1. Learning trials for character-translation pairs were presented for 1500 ms. Working memory interference (math tasks) was included between each round of learning and testing. Test trials included four options and were limited to six seconds. Responses were given by pressing the keyboard arrowkey corresponding to the on-screen orientation. Items were dropped from the list when a correct response had been given on two consecutive rounds. The procedure was repeated until all items of the list were learned.

A pilot study was conducted in order to select stimuli and experimental settings that would provide a sensitive measure on the dependent variable. The pilot study was found to be critical for this purpose since many of the initially chosen character pairs did not even show a general effect of interference. It was also important for avoiding ceiling and floor effects on the dependent variable by adjusting stimulus presentation and testing times as well as choosing appropriate list lengths.

The experiment was programmed using the Presentation software (version 16.3, www.neurobs.com). Stimuli were presented in white font on black background on a 24 inch monitor. Chinese characters were in font type SimSum and 72-point size.

Results

Because interactions were specifically predicted and investigated, any main effects of speed and similarity, or lack thereof, must consequently be interpreted with caution. However, if the interaction is ordinal and predicted a priori, interpretations of significant main effects are arguably still somewhat meaningful. In contrast, if the interaction is disordinal (cross-over) or main effects fail to reach significance, interpretations could be more problematic. Before any statistical hypothesis tests were carried out, ANOVA assumptions were assessed by Levene's test for homogeneity of variance (across groups), Mauchly's test for sphericity (no corrections were necessary), and by visual inspection of various graphical representations of the data. Because the dependent variable was a measure of accuracy – bound between 0 and 1 – and therefore inconsistent with the normal distribution, arcsine transformation was applied. Data analysis was carried out in R (version 2.15.2, www.R-project.org) and IBM SPSS Statistics (version 20, www.ibm.com/software/analytics/spss).

F-tests of a full factorial repeated measures ANOVA model including subject learning speed, within-subject item speed, and similarity was conducted. For the between-subjects comparison of learning speed on post-interference retention, there was a significant main effect, $F(1,19) = 20.80$, $p < 0.001$, $\eta_G^2 = 0.191$ (The effect size reported is generalized eta-squared, η_G^2, because it provides comparability across between-subjects and within-subjects designs, and is therefore recommended for repeated measures designs (Bakeman, 2005)). In addition to a negative correlation of -0.6 ($p = 0.005$) between subject speed and retention, this indicated that fast learners generally remembered more than slow learners. A main effect of similarity, $F(1,19) = 58.31$, $p < 0.001$, $\eta_G^2 = 0.480$, showed that the learning of L2 generally caused a significant level of retroactive interference on L1. In other words, when learning to associate a Chinese character and its translation in the initial list was followed by learning a similar Chinese character with a different meaning, the probability of successful disambiguation and recognition of the initially learned association dropped significantly. Whether this effect of interference was equally strong for participants requiring different numbers of trials to learn the paired associates would be indicated by an interaction of learning speed and similarity, and effectively be informing the between-subjects part of the hypothesis of interest. The interaction term of the repeated measures ANOVA model showed no such relationship $F(1,19) < 1$ (Figure 8a). For item specific learning speed (items acquired fast or slow within subjects), there was a tendency in the direction of the hypothesized interaction (Figure 8b), but the interaction term was not significant $F(1,19) = 1.20$, $p = 0.288$, $\eta_G^2 = 0.021$. However, there was a significant main effect of within-subjects item speed, $F(1,19) = 13.24$, $p = 0.002$, $\eta_G^2 = 0.091$, indicating the same relationship that was found in the between-subjects analysis; Retention of fast items were generally superior when compared to items learned slowly. Summary statistics for Experiment 1 are presented in Table 1.

Analysis including the potential confound of trials between learning and final test was also carried out, but did not change the main results of the experiment and will therefore not be reported separately. This was also the case for Experiment 2, so in order to avoid redundancy, such statistics will be reported in the Joint Analysis later on.

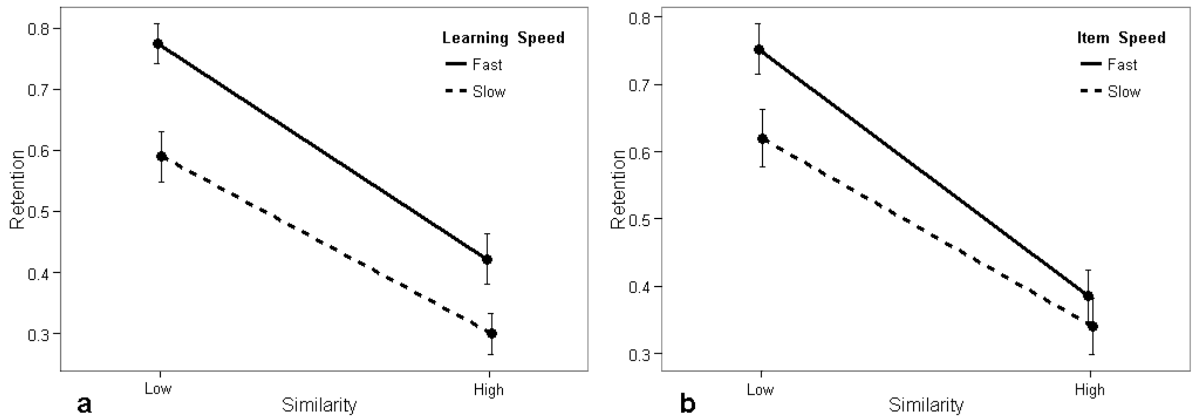

Figure 8. (a) Retention as a function of low and high similarity for fast and slow learners in experiment 1. Parallel lines indicates no interaction of subject learning speed and similarity. (b) Lines represent items learned fast or slow within subjects in Experiment 1. The steeper line of fast items indicates an interaction in the predicted direction. However, the interaction was not significant. Error bars show standard error of the mean (between subjects).

Table 1

Condition means, standard deviations, and standard errors of the mean for Experiment 1

	Low Similarity			High Similarity		
Learning Speed	*M*	*SD*	*SE*	*M*	*SD*	*SE*
Fast Subjects	0.77	0.15	0.03	0.42	0.19	0.04
Slow Subjects	0.59	0.19	0.04	0.30	0.15	0.03
Fast Items	0.75	0.17	0.04	0.39	0.17	0.04
Slow Items	0.62	0.19	0.04	0.34	0.19	0.04

Discussion

The overall results of Experiment 1 did not provide any evidence in support of the predicted interaction of learning speed and similarity. However, subsequent considerations of the experimental setup and results revealed a number of factors indicating some deficiencies of the chosen setup. More specifically, four aspects that had potential for improvements were identified.

First, differences in mean TTC of the first and second sublists of L1 were varying a lot between participants, suggesting there could be differences between participants in how well the preceding task instructions had been understood. This could have affected the overall results because estimates of participants general learning speed were based on their mean

TTC on L1. If such differences were in fact the cause of this variance, measures that were intended to reflect general learning speed could have been strongly influenced by between-participant differences in how well they had understood the task instructions. Improving the instructions and including a training session to familiarize participants with the procedure before any TTC measurements are made could help to stabilize and make the estimates of learning speed more valid.

Second, even though the variance between lists was high, the mean and variance of within-list TTCs was surprisingly low. The mean of mean TTC on L1 was 3.67 with a standard deviation of 1.05. Considering that the drop-out-criterion required that at least two trails were correct before an item was considered learned, a mean of only 1.6 trails over this criterion could be problematic. It could mean that the fastest participants had overlearned some of the items because they did not even need two presentation trials to learn the association. If this was the case, the experimental setup would not have been sensitive to the hypotheses under investigation because such overlearning could have led to representations that were more robust and less vulnerable to interference. Moreover, the low mean and variance across participants also indicates that there was a low number of trials separating fast and slow learners as well as items learned rapidly and slowly. Since such small differences between fast and slow learning does not have much room for differences in encoding variability for specific items, this experiment would probably not be very capable of detecting such effects of encoding variability if they in fact were present. One way to improve on this could be to shorten the presentation time of the learning trials. Because a higher number of learning trials would then be required for each item, this should help in avoiding overlearning, push the mean, and leave more room for differences between fast and slow learning.

Third, even though some effort was put in when selecting stimuli based on the results of the pilot study, an analysis of stimulus-specific effects in this experiment revealed that some items still had, on average, little or even inversed interference effects. In other words, some of the character-translation associations that were intended to produce retroactive interference showed the opposite effect and instead enhanced memory of the associations they were supposed to be interfering with. This suggested that there was room for improvement with respect to stimulus selection.

The final factor identified was the testing procedure. Testing memory by recognizing and selecting the target translation among four options means that one out of four test trials, on average, will be a random hit if participants guess when faced with an item for which the correct translation is not recognized. In fact, in order to reduce variance related to participants reluctance to respond at different degrees of uncertainty, the instructions encouraged guessing. Because the same testing procedure was used for assessing post-interference retention on the final test, it is possible that the dependent variable of the experiment included some error variance due to such random effects. When sample size is increased, such effects will average out and diminish. However, considering the rather small sample size of the current experiment, the presence of such effects can not be ruled out. In order to further investigate the hypotheses, a second experiment was designed with special attention to these problems.

Experiment 2

Method

Participants. 21 participants (14 women, 7 men, M_{age} = 25.1 years, SD_{age} = 4.8, age range: 19-35 years). Three participants were excluded from the analysis. One participant failed to produce any correct answers on the final test while the other two exclusions were based on multivariate outlier analysis (MD). All were fluent Norwegian speakers and none had familiarity with Chinese characters.

Stimuli and procedure. The design of Experiment 2 was the same as Experiment 1. However, the stimuli and procedure used had some important changes. The stimuli of this experiment was a subset of the stimuli used in the preceding experiment. 20 pairs (40 character-translation associations; Appendix 2) were selected based on the analysis of stimulus-specific effects in Experiment 1. In order to stabilize the TTC measurements, a training session was included before the participants started learning the Chinese characters and their meaning. The training session consisted of six Greek letters and used the same setup as the main learning procedure of the experiment. The training session was conducted in order to familiarize participants with the procedure and thereby reducing differences between the first and subsequent lists of the actual experiment due to procedural learning effects and differences in understanding the task instructions. The problems related to overlearning and small differences between fast and slow learning were attempted improved by shortening the presentation time of the learning trials. Presentation time for learning trials throughout Experiment 2 was set to 1000 ms. The final factor identified as problematic in Experiment 1 was improved by increasing the number of test options and having two tests for each item on the final test of post-interference retention. Test options throughout the experiment now had six options which included five randomly chosen translations of the current list in addition to the correct translation (Figure 9). The options were numbered and participants responded by pressing the number corresponding to their choice on a numeric keypad. The final test of retention also adopted the new testing procedure and in addition to testing each item twice and only counting items that were correctly responded to on both runs as successfully retained this reduced the chance of a random hit to 1/36.

Figure 9. Example of test trial in Experiment 2. Trials included six options and responses were given by pressing the corresponding number on a numeric keypad.

Results

The interaction of subject learning speed and similarity was significant, $F(1,16) = 6.70$, $p = 0.020$, $\eta_G^2 = 0.062$, and disordinal (Figure 10a). Fast learners retained more than slow learners when similarity was low. However, when similarity was high, fast learners were inferior to slow learners. Main effect of learning speed was no longer present, $F(1,16) < 1$, however, this was likely due to the disordinal nature of the interaction. As mentioned earlier, such interpretations can be meaningless in the presence of interactions, especially when there is a cross-over. However, in order for such an interaction to be present, there must necessarily be important effects of both factors, and interpretations of main effects alone when there is an interaction could be misleading. The main effect of similarity, however, was significant, $F(1,16) = 21.73$, $p < 0.001$, $\eta_G^2 = 0.176$. Both fast and slow learners retained less when similarity was high. While an interaction in line with the predicted within-subjects effect also was present in the data (Figure 10b), the F-test did not yield the interaction term significant, $F(1,16) = 2.43$, $p = 0.139$, $\eta_G^2 = 0.040$. Main effect of item-specific learning speed was not significant either, $F(1,16) = 1.60$, $p = 0.224$, $\eta_G^2 = 0.021$, but interpretation of such a statistic, as emphasized above, must be related to the possibility of an interaction with similarity. Summary statistics are presented in Table 2.

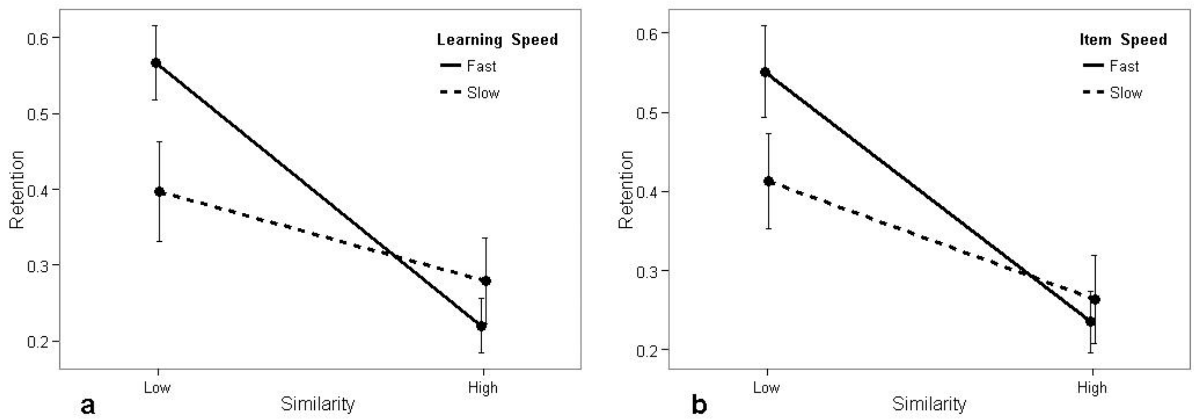

Figure 10. (a) Interaction of subject learning speed and similarity in Experiment 2. Fast learners retained more when similarity was low, but were inferior to slow learners when similarity was high. Interaction was significant (p = 0.020). (b) The same relationship was observed for fast and slow items within subjects, but the interaction was not significant (p = 0.139). Error bars show standard error of the mean (between subjects).

Table 2

Condition means, standard deviations, and standard errors of the mean for Experiment 2

	Low Similarity			High Similarity		
Learning Speed	*M*	*SD*	*SE*	*M*	*SD*	*SE*
Fast Subjects	0.57	0.21	0.05	0.22	0.15	0.04
Slow Subjects	0.40	0.28	0.07	0.28	0.24	0.06
Fast Items	0.55	0.24	0.06	0.24	0.16	0.04
Slow Items	0.41	0.26	0.06	0.26	0.23	0.06

Discussion

The results of Experiment 2 provided evidence in support of the between subject hypothesis, indicating that the measures taken to improve sensitivity of Experiment 1 were effective. Not only was the subject learning speed and similarity interaction significant, it was also disordinal. The current hypotheses did not necessarily predict such a cross-over; It merely suggested that fast learning would be associated with more interference when similarity is high. If it is assumed that fast learners generally retain more over time, such a prediction could still be supported even without a cross-over.

When it comes to the within-subjects part of the hypothesis, both Experiment 1 and 2 failed to produce evidence for the predicted interaction. While such results are in disfavor of the hypothesis, the fact that both these experiments had a relatively low number of participants leaves the possibility that the failure to produce evidence for the predicted effect was due to lack of statistical power. In order to investigate if the small sample sizes of these experiments were the primary source of the lack of statistically significant results, an analysis of the combined results of Experiment 1 and 2 was carried out.

Joint Analysis

Normalization and Data Preparation

Prior to merging the datasets, the dependent variables of both experiments were z-normalized. In order to control for differential effects of the two experiments, a variable indicating which experiment each data point originated from was included. The multivariate outliers that were excluded from the separate analyses were also kept out of the combined analysis. No further exclusions were made, resulting in a combined dataset of 39 subjects.

Results

Experiment as a factor. An initial analysis included experiment, subject speed, within-subject item speed, and similarity as factors in a 2 x 2 x 2 x 2 repeated measures ANOVA. Subject learning speed and similarity interaction was significant, $F(1,35) = 6.43$, $p = 0.016$, $\eta_G^2 = 0.035$. More interestingly, the item-specific speed and similarity interaction now also approached significance, $F(1,35) = 3.63$, $p = 0.065$, $\eta_G^2 = 0.030$. None of these interactions had significant higher order interactions with the experiment factor. This does not necessarily mean that the effects are comparable across experiments. However, if such interactions had been found to be present, justification of the combined analysis could have been more problematic. The above results including the experiment factor were obtained using Type II sum of squares (hierarchical) which is justified by the present hypotheses.

Main effects. Both main effects of subject speed, $F(1,35) = 7.84$, $p = 0.009$, $\eta_G^2 = 0.075$, and item speed, $F(1,35) = 8.24$, $p = 0.007$, $\eta_G^2 = 0.048$, were significant, indicating superior retention for fast learners as well as for items acquired fast within subjects. Similarity had a significant main effect and a large effect size, $F(1,35) = 74.45$, $p < 0.001$, $\eta_G^2 = 0.324$.

The confound. As mentioned in the method section of Experiment 1, the initial learning (L1) was separated into two sublists. The purpose of this was to allow some control of different numbers of intervening trails between learning and post-interference test for fast and slow learning. By calculating the number of trials between drop-out and final test for each item and including this as a covariate in the analysis, it should be possible to inform the role of this potential confound. This variable, trials-to-test (TTT), was scaled by the same procedure used for the dependent variable before it was included in the analysis. Surprisingly, the inclusion of this covariate in fact strengthened the role of the within-subject speed and similarity interaction, $F(1,34) = 4.93$, $p = 0.034$. The subject speed and similarity interaction did not change much, $F(1.34) = 6.26$, $p = 0.018$, but a loss of explanatory power for the between-subject interaction was not expected either because fast learners on average had lower TTT's than slow learners (difference: 0.69 standard deviations on the scaled variable). Hence, fast learners had the advantage of having fewer intervening trials and shorter time delay between learning and final test, suggesting that the experimental procedure used in the present investigation was actually quite conservative with respect to the between-subjects part of the hypothesis. For the within-subjects hypothesis, this potential confound is theoretically more problematic and the observed strengthening of the interaction term was not predicted. This matter will be further discussed in the General Discussion.

Discussion

Even though Experiment 1 alone did not provide any evidence for the hypotheses, a number of factors were identified as problematic regarding the experiment's sensitivity to the hypothesized effects. When proper adjustments were made and tested in Experiment 2, results were in accordance with both hypotheses. However, only the between-subjects part of the hypothesis was statistically significant, while the result of the within-subjects interaction was inconclusive. By merging the data of both experiments in a combined analysis, evidence for both hypotheses was produced. This suggests that the failure of Experiment 2 to produce a firm conclusion regarding the within-subjects hypothesis was most likely due to the small sample size.

The measures taken to improve the sensitivity of the experiment were mainly thought to be effective by reducing error variance, both on the learning speed variable as well as the dependent variable. However, there was one additional difference between the two experiments that could be of importance. The reduction of the number of associations to be learned was primarily done to insure that the effect of the interference learning (L2) was strong. Since learning 60 associations (40 L1 items + 20 L2 items) was found to be a demanding task, the reduction was also motivated by a need to shorten the time required to complete the experiment. However, the difference in the number of list items between the two experiments could also be directly related to the hypothesized effects. Because Experiment 1 had twice the number of associations to be learned, it is likely that more intra-list interference resolution was required to complete the initial learning (L1). This could have resulted in more robust representations of these associations and therefore made them less prone to interference from the items of L2. Because the presentation time was changed from Experiment 1 to Experiment 2, comparing the TTC's between experiments would be inappropriate. In fact, because of this difference, as well as the other between-experiment differences that were implemented, the data available does not have a statistic that would be appropriate for such a comparison. This matter must therefore be left as an open question for future inquiries.

The difference in list length could also have caused the differential results in a more indirect way. Being presented with 20 items in each list could have been experienced as a more overwhelming task and therefore led participants to undertake different strategies. If attention was strongly focused on a subset of the items on each round while the other items were more or less ignored, the TTC's of the items learned late would not be an appropriate measure of item-specific learning speed. This possibility is supported by informal post-experimental reports by participants of Experiment 1. However, the low average TTC of Experiment 1 is not consistent with such an explanation.

Experiment 1 did not produce any strong evidence for any of the hypotheses, but the trend was stronger within subjects than between. In Experiment 2, this trend was reversed; The between-subject learning speed and similarity interaction was significant while the within-subject interaction was not. This pattern brings attention to a more general aspect of the present investigation. When these effects are investigated in the same experiment, the hypothesized effects are actually working against each other. This point is perhaps easiest to explain through visual representation. Consider the graphs of fast and slow learning within subjects in Figure 10b. In order to have a strong interaction, the difference in retention of fast versus slow items in the low similarity condition must be large. However, when this difference increases, between-subject differences in retention in the low condition (which is an average of both fast and slow items) will converge toward some middle value. When the difference increases within subjects, the difference between subjects (in the low similarity condition) will be smaller, resulting in a diminishing between-subject interaction. Consequently, measures taken to increase experimental sensitivity to within-subjects differences can reduce the sensitivity of between-subjects differences. The hypotheses of interest are therefore perhaps more easily investigated in separate experiments. However, it should also be noted that the above reasoning builds on the assumption that fast learning is to be associated with superior retention when similarity is low. Given the inconsistent results of prior research, such an assumption might not be appropriate.

General Discussion

Based on research on computational models of cognition, the encoding variability hypothesis, and behavioral research on individual differences in associative learning and forgetting, two novel hypotheses were built and tested. The overall results of the present investigation suggest that prior research on learning speed and forgetting has not properly considered the role of the learning material used in specific studies. Previous studies on individual differences in learning speed have investigated whether learning speed is a reliable predictor of differences in forgetting. Such studies have mainly focused on differences in forgetting over time and therefore neglected the possibility that: (a) Memory representations established by fast and slow learning are different and therefore have differential vulnerability to interference; and (b) Fast and slow learning between initial learning and memory testing interacts with the level of similarity of the initial and intervening learning. The current experiments addressed these problems by investigating whether similarity of intervening learning moderates the relationship between learning speed and retention. An interaction of learning speed and stimuli similarity was hypothesized for both between-subjects speed and item-specific speed. The experimental results provided evidence for both hypotheses by showing that retention of Chinese characters and their meaning depends on the interaction of how fast they were learned and the similarity of subsequently learned characters. This was found both when comparing learning speed between subjects as well as for different speeds of acquisition of items within subjects.

Explaining the Results

The present investigation was primarily concerned with establishing that the interaction of interest is in fact present in associative learning at the behavioral level. Because of this, focus was not on designing an experiment that would be capable of differentiating between the potential causes behind this effect. Any attempts towards such conclusions based on the present findings would have poor empirical foundation. However, a few notes on this is nevertheless in order. Two factors were emphasized when building the hypotheses: (a) Learning rate of neural network models with overlapping representations; and (b) The encoding variability hypothesis. As noted in the introduction, it is important to distinguish between these. Even though they make the same prediction, the mechanisms behind are different. Neural network model simulations suggest that fast learning (due to high levels on the learning rate parameter) will cause more interference when prior knowledge and novel learning have common features (i.e., high similarity). In other words, stronger updates of connections that are part of an existing representation will lead to faster forgetting. In contrast, the encoding variability hypothesis can be extended to suggest that different learning speeds will result in memory representations with different structures, even though associative strength is comparable at the behavioral level, and will therefore have different levels of vulnerability to interference. The learning rate explanation would primarily predict between-subjects differences while the encoding variability hypothesis relates to the number of learning trials or exposure time and would therefore be consistent with both between and within subject learning speed interactions with similarity of subsequent learning. Even though the present results cannot rule out the learning rate explanation, the fact that interactions were observed both between and within subjects would be more consistent with the encoding variability account. Moreover, if an effect in line with the learning rate explanation was the primary source of the observed interaction, it seems probable that there should also be a higher order interaction of subject speed, item speed and similarity. The result of the present investigation indicated no such relationship.

Other Explanations

Even though the experiments produced evidence for the hypotheses under investigation, there are also some alternative explanations for these results that are not necessarily consistent with the hypotheses.

Willingness to respond. If the participants of these experiments were varying a lot in how willing they were to respond at different levels of certainty, this could have led to the pattern of the between-subjects data. If some participants were more reluctant to respond than others, this could have caused systematic overlearning among participants categorized as slow learners. Such overlearning could have led to more robust associations less vulnerable to interference. While this possibility can not be ruled out for the between-subjects effect, the interaction of item-specific speed and similarity is not consistent with such an explanation.

Floor effect. Given the main effects of learning speed and item-specific speed (i.e., fast learners and fast items were generally superior), there is a possibility that a floor effect could have imposed the interactions observed. Because retention was generally lower for slow participants and slow items, the lower bound of zero retention could have been more restrictive for slow participants and slow items in the high similarity condition. However, a number of observations suggest that this is not a likely explanation. First, the fact that disordinal interactions were observed is inconsistent with this possibility. Second, if slow participants and slow items were more restricted by the lower bound than fast participants and fast items, there should be a difference in the number of participants and items touching the lower limit of zero retention in the high similarity condition. The data indicated no such differences.

Incidental testing effect. The covariate (TTT) which was intended to control for order effects by controlling for the number of intervening trials between initial learning and final testing could have been reflecting something very different. The pilot study revealed a surprising effect that can be related to this potential problem. The initial setup of the pilot study did not have options in the testing trials. Rather, testing was done by having the participants type their response when prompted with a Chinese character. This procedure resulted in an inversed interference effect of similarity. In other words, participants actually performed better on the items that were in the high similarity condition. Informal analysis of these results suggested that the interference learning (L2) functioned as a testing opportunity for L1 items rather than interference. When participants were prompted with a character in test trials, responses were given for the corresponding L1 character. This could go on for a high number of trials before the difference between the L1 and L2 character was realized. However, once this was realized, a testing effect (e.g., Roediger & Karpicke, 2006) on the L1 associations had been in effect, resulting in improved memory for the L1 associations and therefore also improving the participants ability to distinguish between the interfering characters on the final test. Even though the final experiments had a strong interference effect of L2, the presence of such testing effects cannot be ruled out. However, there are no strong reasons to believe that this effect would be different for fast and slow participants and should therefore not have any implications for the results obtained for the between-subjects interaction. For the within-subjects effect, however, this is a bit more problematic. If there was in fact a testing effect of learning L2, this effect could have been different for fast and slow items due to the different numbers of intervening trials between learning a specific association in L1 and learning the interfering association in L2. Furthermore, if a testing effect was present, this could invalidate the TTT measure as a proper control of the order effects. If it is assumed that testing effects were indeed present and that slow items benefited more from such testing, the unpredicted strengthening of the interaction term when the covariate (TTT) was included – as reported in the Joint Analysis above – now makes sense. However, given the strong interference effect of L2 in both experiments, the assumption that a testing effect was present seems unlikely.

Limitations and Future Research

As should be clear by the alternative explanations discussed in the above section, the present investigation had a number of important limitations that should be addressed. The limitations of these experiments are not only important with respect to informing the specific hypotheses under investigation, but also when assessing the observed effects in a larger context. For instance, the critique of previous research for not properly distinguishing between time and interference is also valid for the experiments of the present investigation. Given the inconclusive results related to the covariate, order and time delay should be addressed more carefully in future studies. For instance, having a higher number of sublists would provide more levels of the order factor and therefore a more sensitive distinction between the number of intervening trials and the controlled interference. Another way to overcome this problem is to not have any sublists at all, but rather continuously update one list by introducing a new item whenever an item has reached the drop-out criterion. The benefit of such a procedure would be that order can be controlled by comparing items with different TTC's that were dropped at the same time.

Specifically comparing forgetting over longer time intervals as well as manipulating interference would also be an interesting approach. If the goal is to specifically address differential forgetting rates between fast and slow individuals, which has been the objective of many of the previous studies discussed in the present report, manipulating both time and similarity should reveal some interesting high-order interactions.

Another obvious shortcoming of the present experiments was the low number of levels on each factor. Both mechanisms (learning rate; encoding variability) that led to the current hypotheses would predict a threshold level of similarity for which the interaction would be expected above but not below. The neural network prediction was based on learning rate and overlapping representations. Because at least some degree of overlap will be required for an interaction, there should be a corresponding level of similarity for such a threshold. The encoding variability account would also predict a threshold if reactivation (study-phase retrieval; e.g., Delaney et al., 2010) is assumed. The problem with this is that such a prediction would not be able to distinguish these accounts since they both predict a threshold. The prediction of such a threshold could be addressed experimentally by including a higher number of levels on the similarity factor.

It should be noted that the learning speed and item speed variables in the present experiments were continuous and therefore included more information than what was reflected by the median splits and two-level factors used in the repeated measures ANOVA. It could be argued that binomial regression would be a more appropriate approach for data analyses in the present investigation. Such a proposition is not disputed, however, the choice of statistical tests mainly reflected three factors: Simplicity, correspondence with previous studies, and correspondence between statistical tests and the graphical representations of the results (Figures 8 and 10).

One major challenge of conducting experiments on learning speed is to identify different sources of learning speed. This makes it difficult to relate behavioral effects to low level mechanisms. As mentioned in the introduction, different sources of learning speed can lead to different predictions for forgetting. This was also a problem in the present investigation. Because the source of the differences in learning speed, manifested as different TTC's, could not be identified, it was not possible to connect the observed interactions to a specific underlying mechanism. Future investigations should address this problem by stronger experimental control of the learning procedures. For instance by controlling the use of learning strategies. A more direct extension of the keyword mnemonic versus rote memorization of Wang et al. (1992) looking at similarity rather than time delay would be interesting. If strategy is tightly controlled, such an experiment could be informative with respect to the mechanisms behind the between-subjects hypothesis of the current investigation.

Distinction between the neural network and encoding variability accounts could also be investigated by employing eye-tracking techniques. As mentioned earlier, encoding variability can be related to a number of different aspects including meaning, context, components and features. If variability of encoding visual features of stimuli is assumed, an eye-tracking approach would be interesting.

The encoding variability account could also be informed by fMRI and MVPA. The approach of Xue et al. (2010) to investigate encoding variability by comparing pattern similarity across study trials could be extended by including learning speed and levels of similarity. If the disordinal interaction of the present study is replicated for the behavioral data then comparing pattern similarity across trials for different learning speeds and levels of interference should be informative with respect to the encoding variability account. Not only in relation to the hypotheses of the present study, but also for the encoding variability hypothesis in general.

Designing experiments capable of directly informing the neural network account of the hypothesized effect will probably be more challenging. One way this could be approached would be to use the basic design of the present experiments in other learning domains not as dependent on conscious strategies. For instance, a motor task or some kind of implicit learning.

Implications

Even though the results of the experiments reported herein are difficult to directly connect to the underlying mechanisms, the results themselves, however, have some implications that should be noted.

Theoretical implications. The interaction of learning speed and similarity, observed in these experiments, suggests that care should be taken when comparing retention for representations established by different speeds of learning. Interpretation of such a relationship should not be conducted without considering what is going on during the time between acquisition and memory test. This would be like attempting to interpret and generalize main effects when an interaction is known to be present. Such conclusions would only hold as a simple effect valid only for the specific stimuli and context of the situation in which it was produced. This has serious implications for how the results of previous studies, such as those presented in the introduction, should be interpreted. If associations established by fast and slow learning have different vulnerability to interference, it will be insufficient to control for interference by only considering time or number of intervening learning trails between learning and test. The actual stimuli and learning material used must also be considered.

The interaction observed does not directly inform the question of individual differences in forgetting rate for fast and slow learners. However, what it does suggest is that such investigations must take care to specify the assumptions on which their conclusions are based. The classical debate of decay versus interference becomes crucial when making conclusions based on learning speed experiments. Inferences regarding individual differences in forgetting rate must specify whether such conclusions are based on theories of decay, interference, or a combination.

The evidence produced by the present experiments does not directly inform the decay versus interference debate. However, it should be noted that these results emphasizes the role of interference effects in general. Furthermore, one of the suggested mechanisms behind the hypothesized effect was based on neural network simulations, and in fact, such simulations were the source that inspired the present investigation. The current results cannot directly inform the appropriateness of neural networks as models of human cognition, however, they should be enlisted in the evergrowing body of results predicted by such models.

Practical implications. The present investigation was primarily concerned with basic mechanisms of learning and forgetting. However, the parallel to research results on learning strategies suggests some practical implications as well. Contemporary interest in mnemonic techniques is high, evident by the increasing media coverage of memory championships. While the skills of the contestants of such competitions are very impressive, discussions related to the techniques used are rarely balanced. The short-term benefits seem to be obvious, but when such techniques are discussed in an educational context, the possibility of long-term detrimental effects should also be addressed. Research on learning strategies within educational psychology has been concerned with this question for a long time (e.g., Wang et al., 1992; Carney & Levin, 1998). However, because of the pragmatic orientation of such studies, focus has not been on identifying *why* forgetting is different for different time intervals. The present results could provide some clues toward such explanations.

Conclusion

The results of the present investigation suggests that retention of associations established by fast and slow learning are moderated by similarity of intervening learning. Fast learners and associations learned rapidly within subjects were found to be more prone to interference, evident by inferior retention when similarity of the intervening learning was high. This effect was found both between individuals as well as for item-specific learning speed within individuals. This interaction of learning speed and stimuli similarity extends prior research on individual differences in learning and forgetting by emphasizing the possibility of differential interference effects for fast and slow learning. Future studies should further investigate these findings by experimental designs capable of differentiating between the suggested underlying mechanisms of this effect.

References

Appleton-Knapp, S. L., Bjork, R. A., & Wickens, T. D. (2005). Examining the spacing effect in advertising: Encoding variability, retrieval processes, and their interaction. *Journal of Consumer Research*, *32*(2), 266-276.

Atkinson, R. C. (1975). Mnemotechnics in second-language learning. *American Psychologist*, *30*(8), 821-828.

Bakeman, R. (2005). Recommended effect size statistics for repeated measures designs. *Behavior Research Methods*, *37*(3), 379-384.

Barnes, J. M., & Underwood, B. J. (1959). Fate of first-list associations in transfer theory. *Journal of Experimental Psychology, 58,* 97-105.

Bellezza, F. S. (1981). Mnemonic devices: Classification, characteristics, and criteria. *Review of Educational Research*, *51*(2), 247-275.

Booth, J.L., Koedinger, K.R., & Siegler, R.S. (2007). The effect of prior conceptual knowledge on procedural performance and learning in algebra. In D.S. McNamara & J.G. Trafton (Eds.) *Proceedings of the 29th Annual Cognitive Science Society* (pp. 137-142). Austin, TX: Cognitive Science Society.

Bowers, J. S. (2009). On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological Review*, *116*(1), 220-251.

Bowers, J. S. (2010). More on grandmother cells and the biological implausibility of PDP models of cognition: A reply to Plaut and McClelland (2010) and Quian Quiroga and Kreiman (2010). *Psychological Review*, *117*(1), 300-306.

Carney, R. N., & Levin, J. R. (1998). Do mnemonic memories fade as time goes by? Here's looking anew!. *Contemporary Educational Psychology*, *23*(3), 276-297.

Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, *132*(3), 354-380.

Cichy, R. M., Chen, Y., & Haynes, J. D. (2011). Encoding the identity and location of objects in human LOC. *Neuroimage*, *54*(3), 2297-2307.

Delaney, P. F., Verkoeijen, P. P., & Spirgel, A. (2010). Spacing and testing effects: A deeply critical, lengthy, and at times discursive review of the literature. *Psychology of Learning and Motivation*, *53*, 63-147.

Dempster, F. N. (1989). Spacing effects and their implications for theory and practice. *Educational Psychology Review*, *1*(4), 309-330.

Desrochers, A., Gélinas, C., & Wieland, L. D. (1989). An Application of the Mnemonic Keyword to the Acquisition of German Nouns and Their Grammatical Gender. *Journal of Educational Psychology*, *81*(1), 25-32.

Donovan, J. J., & Radosevich, D. J. (1999). A meta-analytic review of the distribution of practice effect: Now you see it, now you don't. *Journal of Applied Psychology*, *84*, 795-805.

Ebbinghaus, H. (1913). Memory: A contribution to experimental psychology (H. A. Ruger & C. E. Bussenius, Trans.). New York: Teachers College. (Original work published 1885).

Eger, E., Ashburner, J., Haynes, J. D., Dolan, R. J., & Rees, G. (2008). fMRI activity patterns in human LOC carry information about object exemplars within category. *Journal of Cognitive Neuroscience*, *20*(2), 356-370.

Ewbank, M. P., Schluppeck, D., & Andrews, T. J. (2005). fMR-adaptation reveals a distributed representation of inanimate objects and places in human visual cortex. *Neuroimage*, *28*(1), 268-279.

Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science*, *6*(3), 205-254.

French, R. M. (1992). Semi-distributed representations and catastrophic forgetting in connectionist networks. *Connection Science*, *4*(3-4), 365-377.

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, *3*(4), 128-135.

Gentile, J. R., Monaco, N., Iheozor-Ejiofor, I. E., Ndu, A. N., & Ogbonaya, P. K. (1982). Retention by "fast" and "slow" learners. *Intelligence*, *6*(2), 125-138.

Gentile, J. R., Voelkl, K. E., Pleasant, J. M., & Monaco, N. M. (1995). Recall after relearning by fast and slow learners. *The Journal of Experimental Education*, *63*(3), 185-197.

Greene, R. L. (1989). Spacing effects in memory: Evidence for a two-process account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(3), 371-377.

Gross, C. G. (2002). Genealogy of the "grandmother cell". *The Neuroscientist*, *8*(5), 512-518.

Hardt, O., Nader, K., & Nadel, L. (2013). Decay happens: The role of active forgetting in memory. *Trends in Cognitive Sciences, 17*(3), 111-120.

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.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 1* (pp. 77-109). Cambridge, MA: MIT Press.

Hintzman, D. L., & Block, R. A. (1973). Memory for the spacing of repetitions. *Journal of Experimental Psychology*, *99*(1), 70-74.

Hintzman, D. L., Summers, J. J., & Block, R. A. (1975). Spacing judgments as an index of study-phase retrieval. *Journal of Experimental Psychology: Human Learning and Memory*, *1*(1), 31-40.

Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, *12*, 35-51.

Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, *96*(16), 9379-9384.

Janiszewski, C., Noel, H., & Sawyer, A. G. (2003). A meta-analysis of the spacing effect in verbal learning: Implications for research on advertising repetition and consumer memory. *Journal of Consumer Research*, *30*(1), 138-149.

Johnston, W. A., & Uhl, C. N. (1976). The contributions of encoding effort and variability to the spacing effect on free recall. *Journal of Experimental Psychology: Human Learning and Memory*, *2*(2), 153-160.

Kahn, I., Davachi, L., & Wagner, A. D. (2004). Functional-neuroanatomic correlates of recollection: Implications for models of recognition memory. *The Journal of Neuroscience*, *24*(17), 4172-4180.

Kyllonen, P. C., & Tirre, W. C. (1988). Individual differences in associative learning and forgetting. *Intelligence*, *12*(4), 393-421.

Lewandowsky, S., & Li, S. C. (1995). Catastrophic interference in neural networks: Causes, solutions, and data. In F. N. Dempster & C. J. Brainerd (Eds.), *Interference and inhibition in cognition* (pp. 329-361). San Diego, CA: Academic Press.

MacDonald, S. W., Stigsdotter-Neely, A., Derwinger, A., & Backman, L. (2006). Rate of acquisition, adult age, and basic cognitive abilities predict forgetting: New views on a classic problem. *Journal of Experimental Psychology: General*, *135*(3), 368-390.

Martin, E. (1968). Stimulus meaningfulness and paired-associate transfer: An encoding variability hypothesis. *Psychological Review*, *75*(5), 421-441.

Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, *58*, 25-45.

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), 419-457.

McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *The Psychology of Learning and Motivation*, *24,* 109-164.

McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, *39*(4), 352-370.

McRae, K., & Hetherington, P. A. (1993). Catastrophic interference is eliminated in pretrained networks. In *Proceedings of the 15h Annual Conference of the Cognitive Science Society* (pp. 723-728). Hillsdale, NJ: Lawrence Erlbaum.

Melton, A. W. (1970). The situation with respect to the spacing of repetitions and memory. *Journal of Verbal Learning and Verbal Behavior*, *9*(5), 596-606.

Nyberg, L., Petersson, K. M., Nilsson, L. G., Sandblom, J., Åberg, C., & Ingvar, M. (2001). Reactivation of motor brain areas during explicit memory for actions. *Neuroimage*, *14*(2), 521-528.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: The MIT Press.

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), 580-590.

Page, M. (2000). Connectionist modelling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, *23*(4), 443-467.

Paivio, A. (1974). Spacing of repetitions in the incidental and intentional free recall of pictures and words. *Journal of Verbal Learning and Verbal Behavior*, *13*(5), 497-511.

Parsons, R. J. (1968). Perceptual selectivity and information processing by fast and slow learners in paired-associate learning. *Perception & Psychophysics*, *4*(3), 144-146.

Plaut, D. C., & Behrmann, M. (2011). Complementary neural representations for faces and words: A computational exploration. *Cognitive Neuropsychology*, *28*(3-4), 251-275.

Plaut, D. C., & McClelland, J. L. (2010). Locating object knowledge in the brain: Comment on Bowers's (2009) attempt to revive the grandmother cell hypothesis. *Psychological Review*, *117*(1), 284-288.

Postman, L., & Greenbloom, R. (1967). Conditions of cue selection in the acquisition of paired-associate lists. *Journal of Experimental Psychology*, *73*(1), 91-100.

Pressley, M., Levin, J. R., Nakamura, G. V., Hope, D. J., Bispo, J. G., & Toye, A. R. (1980). The keyword method of foreign vocabulary learning: An investigation of its generalizability. *Journal of Applied Psychology*, *65*(6), 635-642.

Raaijmakers, J. G. (2003). Spacing and repetition effects in human memory: Application of the SAM model. *Cognitive Science*, *27*(3), 431-452.

Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, *97*(2), 285-308.

Richardson, J. (1971). Cue effectiveness and abstraction in paired-associate learning. *Psychological Bulletin*, *75*(2), 73-91.

Richardson, J. (1972). Stimulus selection in associative learning. In C. P. Duncan, L. Sechrest, & A. W. Melton (Eds.), Human memory: Festschrift in honor of Benton J. Underwood. New York: Appleton-Century-Croft.

Richardson, J. (1973). Effect of speed of learning and degree of learning on cue selection. *Journal of Experimental Psychology*, *98*(2), 396-403.

Richardson, J. (1976). Component selection in paired-associate learning: Research and Theory. *American Journal of Psychology*, *89*(1), 3-49.

Rissman, J., & Wagner, A. D. (2012). Distributed representations in memory: Insights from functional brain imaging. *Annual Review of Psychology*, *63*, 101-128.

Roediger, H. L., & Karpicke, J. D. (2006). The power of testing memory: Basic research and implications for educational practice. *Perspectives on Psychological Science*, *1*(3), 181-210.

Roediger, H. L., Weinstein, Y., & Agarwal, P. K. (2010) Forgetting: Preliminary considera tions. In S. Della Sala (Ed.), Forgetting (pp. 1–22). New York: Psychology Press.

Rosenheck, M. B., Levin, M. E., & Levin, J. R. (1989). Learning botany concepts mnemonically: Seeing the forest and the trees. *Journal of Educational Psychology*, *81*(2), 196-203.

Rumelhart, D. E., McClelland, J. L., & PDP Research Group. (1986). *Parallel distributed processing, Vols. 1 and 2.* Cambridge, MA: The MIT Press.

Sahakyan, L., & Goodmon, L. B. (2007). The influence of directional associations on directed forgetting and interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(6), 1035-1049.

Schoer, L. (1962). Effect of list length and interpolated learning on the learning and recall of fast and slow learners. *Journal of Educational Psychology*, 53(4), 193-197.

Seidenberg, M. S. (1993). Connectionist models and cognitive theory. *Psychological Science*, *4*(4), 228-235.

Shepard, R. N. (1963). Comments on professor Underwood's paper: Stimulus selection in verbal learning. In C. N. Cofer & B. S. Musgrave (Eds.), *Verbal behavior and learning: Problems and processes* (pp. 48-70). New York: McGraw-Hill.

Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: Hippocampal-midbrain encoding of overlapping events. *Neuron*, *60*(2), 378-389.

Shuell, T. J., & Keppel, G. (1970). Learning ability and retention. *Journal of Educational Psychology*, *61*(1), 59-65.

Spiridon, M., & Kanwisher, N. (2002). How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron*, *35*(6), 1157-1166.

Squire, L. R., & Wixted, J. T. (2011). The cognitive neuroscience of human memory since H.M. *Annual Review of Neuroscience*, 34, 259-288.

Storm, B. C., Bjork, E. L., & Bjork, R. A. (2008). Accelerated relearning after retrieval-induced forgetting: The benefit of being forgotten. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(1), 230-236.

Stroud, J. B., & Carter, L. J. (1961). Inhibition phenomena in fast and slow learners. *Journal of Educational Psychology*, *52*(1), 30-34.

Stroud, J. B., & Schoer, L. (1959). Individual differences in memory. *Journal of Educational Psychology*, *50*(6), 285-292.

Tong, F., & Pratte, M. S. (2012). Decoding patterns of human brain activity. *Annual Review of Psychology*, *63*, 483-509.

Underwood, B. J. (1954). Speed of learning and amount retained: A consideration of methodology. *Psychological Bulletin*, *51*(3), 276-282.

Wang, A. Y., & Thomas, M. H. (1992). The effect of imagery-based mnemonics on the long-term retention of Chinese characters. *Language Learning*, *42*(3), 359-376.

Wang, A. Y., Thomas, M. H., & Ouellette, J. A. (1992). Keyword mnemonic and retention of second-language vocabulary words. *Journal of Educational Psychology*, *84*(4), 520-528.

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*, 97-101.

Appendix A

Stimuli used in Experiment 1

L1 English	L1 Norwegian	L1	L2	L2 Norwegian	L2 English
alone	alene	孑	子	barn	child
he	han	他	地	jord	ground
of	av	的	时	tid	time
this	dette	这	还	mer	more
and	og	和	知	vite	know
soldier	soldat	士	土	bonde	peasant
the	den	所	听	høre	listen
day	dag	天	夫	mann	husband
white	hvit	白	日	sol	sun
write	skrive	写	与	gi	give
mile	mil	哩	理	årsak	reason
always	alltid	永	水	vann	water
simple	enkel	卞	下	neste	next
large	stor	万	方	vei	direction
reside	bo	住	隹	fugl	bird
department	avdeling	司	可	kan	can
see	se	看	春	vår	spring
home	hjem	家	象	elefant	elephant
door	dør	戶	戸	familie	family
hang	henge	垂	重	tung	heavy
ask	spør	问	向	til	to
not	ikke	没	设	hvis	if
buy	kjøpe	买	实	ekte	real
open	åpen	开	并	allianse	union
already	allerede	已	己	har	has
wait	vent	等	第	først	first
hair	hår	毛	手	hånd	hand
age	alder	代	付	betale	pay
close	lukk	合	台	stasjon	station
each	alle	各	名	navn	name
military	militær	军	车	maskin	machine
tool	verktøy	具	真	sant	true
flat	flat	平	乎	nær	near
right	høyre	右	石	stein	rock
sprinkle	strø	洒	酒	vin	wine
sleep	sove	眠	眼	øye	eye
shell	skall	甲	申	utvide	expand
high	høy	卬	卯	tidlig	early
easy	lett	易	昜	lys	bright
in	inn	入	人	person	person

Appendix B

Stimuli used in Experiment 2

L1 English	L1 Norwegian	L1	L2	L2 Norwegian	L2 English
alone	alene	子	子	barn	child
he	han	他	地	jord	ground
this	dette	这	还	mer	more
and	og	和	知	vite	know
soldier	soldat	士	土	bonde	peasant
write	skrive	写	与	gi	give
always	alltid	永	水	vann	water
simple	enkel	卡	下	neste	next
large	stor	万	方	vei	direction
door	dør	戶	戸	familie	family
hang	henge	垂	重	tung	heavy
ask	spør	问	向	til	to
buy	kjøpe	买	实	ekte	real
already	allerede	已	己	har	has
wait	vent	等	第	først	first
hair	hår	毛	手	hånd	hand
each	alle	各	名	navn	name
right	høyre	右	石	stein	rock
sleep	sove	眠	眼	øye	eye
easy	lett	易	昜	lys	bright