

Passing the Test: Improving Learning Gains by Balancing Spacing and Testing Effects

Hedderik van Rijn (D.H.van.Rijn@rug.nl)

Department of Psychology & Department of Artificial Intelligence
Grote Kruisstraat 2/1
Groningen, 9712 TS The Netherlands

Leendert van Maanen (leendert@ai.rug.nl) & Marnix van Woudenberg

Department of Artificial Intelligence,
P.O. Box 407
9700 AK Groningen, The Netherlands

Abstract

Where the spacing effect promotes longer intervals between facts that needs to be memorized, the testing effect argues for intervals that are short enough to recall the facts. As the ease by which facts are memorized differs greatly between students, an individual assessment of how well certain facts are represented in memory is required to successfully balance spacing and testing effects. We present a model that adapts itself to the abilities of the student, and show in a real-world experiment that this model outperforms other approaches to spacing.

Keywords: spacing-effect; testing-effect; subsymbolic model tracing; cognitive model.

Introduction

The last couple of years have seen a renewed interest in applying insights from fundamental memory research in real-world settings. One of the most visible lines of work are studies to the application on the spacing effect. The spacing effect, first described by Ebbinghaus (1913/1885) at the end of the 19th century, is the positive effect on factual recall that is observed when study trials are temporally separated. Thus, the probability of recall of facts learned in a spaced sequential order (e.g., abcabcabc or abc-break-abc-break-abc) is higher than the probability of recall of facts that are learned massed (e.g., aaabbbccc). The consequence of this finding is that the presentation sequence of a to-be-memorized list of facts partly determines how well these facts will be recalled on a later test: all items on a list that presents the items with optimal spacing will be recalled better than items on a list that presents the items as many times as the first list, but massed instead of spaced.

This observation was central to a lot of applied research in the '60s and early '70s of the last century. Using the possibilities provided by digital computers, scientists tried to construct optimal schedules. Although some of this work has stood the test of time from an applied or commercial point of view (e.g., the Pimsleur and Leitner methods are still available commercially), the methods used by these early systems are relatively simple and the learning gains often did not outweigh the extra investment associated with using these systems. This led to a decline in applied research on the spacing effect, although over the decades, more fundamental research on this effect has thrived (see for reviews Dempster, 1988, Cepeda, Pashler, Vul, Wixted, & Rohrer,

2006). Only recently has attention again shifted to using algorithms to determine the optimal schedule for learning (Wozniak & Gorzalanczyk, 1994, Pavlik, 2007, Pavlik & Anderson, 2008).

A second finding that has a potentially large effect on how on optimal sequence has to be constructed, is the testing effect that can be described as: "If students are tested on material and successfully recall or recognize it, they will remember it better in the future than if they had not been tested" [but merely studied the same material] (Roediger & Karpicke, 2006, p.249, see also Carrier & Pashler, 1992). As it is generally assumed that memory decays over time, and that decay is in some way a function of time, an increased interval between successive presentations makes it more likely that an item cannot be recalled, thus failing the benefit of the testing effect (c.f., the inverse u-shape often observed when the performance on a test is plotted as a function of the interval between two presentations, Cepeda, Vul, Rohrer, Wixted, & Pashler, 2008).

When it comes to computing an optimal presentation sequence for fact learning, spacing and testing have different interests. For the spacing effect, increased spacing is theoretically preferred. But for the testing effect, small to no spacing would theoretically provide best results since longer intervals increase the chances of having forgotten a fact, and thus decrease probability of recall. One of the aims of the study reported here is to reconcile these seemingly conflicting requirements.

An interesting observation in almost all work on the spacing effect is that the "optimal schedule" is defined as the schedule that reaches a rather high-level, almost noble goal: the best performance (often defined as the highest probability of recall) over a longer timeframe. Although this is of course what the goal of all learning *should be*, the goal of learning in a real-world situation is often more pragmatic: passing the next day's test by studying for a limited, often more or less fixed amount of time. So, although the results of more than a century of spacing results can be used for the real-world situation of having to learn a lot of vocabulary word pairs for a foreign language test that is scheduled a couple of weeks or months in advance, these results do not necessarily tell us anything about the pragmatic goal of learning: What method should a student use to learn a set of 20 vocabulary word pairs for a potential test tomorrow,

knowing that, because of soccer practice, favorite TV-shows and other homework, all he or she has is 15 minutes to spare?

Note that this real-life situation differs quite a bit from typical experimental setups: First, to prevent contamination of the results by different levels of prior knowledge, the learning materials are often selected in such a way that it is to be expected that none of the participants brings any prior knowledge (by either learning sequences of nonwords, e.g., Ebbinghaus, 1913/1885, very obscure facts, e.g., Cepeda et al, 2008, or word-pairs from languages previously unstudied and selected to be from different language-families, e.g., Japanese-English, Pavlik & Anderson, 2008). In contrast, when learning for a vocabulary test, most students bring additional knowledge to the learning session from earlier experiences with that language. Second, in most studies the list of word pairs presented to the participants is much longer than the 10 to 30 words that typically have to be learned in a single learning session. Third, the retention interval (defined as the time between the final test on the learned materials and the last study of the materials) in most studies is less than a day (221 out of 254 studies reviewed in Cepeda et al, 2006, used an interval less than a day). Fourth, where many experimental studies aim for finding a general law that describes the effects of different types of spacing on performance in general, the goal of an individual student is not striving for the best performance of a larger group, but for the optimal results on his or her test. As the speed and ease by which vocabulary is learned differs greatly between individuals (see Baddeley, 2003, p832), settings that are optimal for the group as a whole might not be the optimal settings for an individual. These differences are less substantial with respect to the spacing effect than with respect to the testing effect. That is, irrespective of the individual expertise in vocabulary learning, the spacing effect predicts that increased spacing provides better scores. However, with respect to the testing effect, individual differences greatly determine the probability of recall of a particular item. Since successful recall is associated with better learning gains, it is important to account for individual differences in such a way that facts are presented before they cannot be recalled anymore.

To test whether the general findings associated with spacing and testing effects hold when these issues are taken into account, we ran an experiment that closely mimics everyday learning contexts. In this experiment, four classes of approximately 15-year old pre-university level pupils were asked to memorize Dutch translations of 20 French words that were selected from a lesson that would not be discussed until one week after the experiment. Each word pair was presented first in a study trial of 5 seconds in which both the French and the Dutch word were presented. Immediately following the study trial, participants were presented with a test trial of that same word. During a test trial, only the French word was presented, and the participant had to type the Dutch translation. (If the participant did not respond in time, or an incorrect answer was given, the study trial was presented to refresh the participant's memory.) After the

initial presentation, the next trial was scheduled on the basis of one of four algorithms.

Algorithm 1 was based on a flashcard strategy: the 20 words were clustered in 4 sets of 5 words which were presented individually until all words in the set had been responded to correctly once. After all sets had been presented, the sequence was started anew. **Algorithm 2** is an implementation of the spacing method proposed by Pavlik and Anderson (2005), which will be discussed below. **Algorithm 3 and 4** are adaptations of the original Pavlik and Anderson algorithm in that the model that is used to determine the optimal sequence is dynamically adapted on the basis of the observed performance of the student while taking the testing effect in account. Before turning to these algorithms, we will first discuss Pavlik and Anderson's spacing model and how this model can be applied to provide an optimal learning sequence.

Pavlik & Anderson's Spacing Model

The spacing model proposed by Pavlik and Anderson (2005, referred to as the PA model) is based on the work of Anderson and Schooler (1991). In this work, Anderson and Schooler demonstrated that the "availability of human memories for specific items shows reliable relationships to frequency, recency, and pattern of prior exposures to the item" (Anderson & Schooler, 1991, p.396). Based on their and later work, the following formula was proposed to express the availability (or activation) A of a certain item i at a certain time (t) as a function of prior encounters:

$$A_i(t) = \sum_{j=1}^n (t - t_j)^{-d_j}$$

According to this equation, which has become central to all memory related models created in the ACT-R cognitive architecture (Anderson, 2007), all previous encounters ($t_1..t_n$) of the item i contribute to its current activation. However, the older an encounter (t_j represents the time of encounter j), the smaller the contribution of that encounter to the total activation. The speed of this decline is expressed by $-d_j$, the decay parameter. Although initially $-d_j$ was assumed to be variable for different encounters j (and Anderson and Schooler, 1991, even provided an equation to account for some spacing effects but downplayed its importance by noting that "its exact form is a bit arbitrary", p.407), it quickly became a parameter that was treated as a constant ($d=.5$, see for example Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004) as different values for different encounters did not add much explanatory power for most tasks to which this equation was applied. However, in contrast to the original work of Anderson and Schooler, in none of these later tasks was spacing a factor of importance. To account for a broader range of spacing phenomena, the PA model reintroduced individual decay values for individual items.

Pavlik and Anderson proposed to relate the decay values for the individual encounters to the activation of that particular item at the time of the encounter. As recently presented items have a high activation, the second encounter of an item presented twice in quick succession will be associated with a high decay value. Therefore, the long-term influence of this item will be small as its activation will decay

quickly. On the other hand, an encounter of an item of which the last presentation was longer ago (and therefore has a lower activation) will receive a lower decay value, resulting in more long term impact on the activation of that item. The proposed equation calculates the decay, d , for encounter j of item i by calculating the activation of that item (A_i) at the time of encounter j .

$$d_{ji} = ce^{A_i(t_j)} + \alpha$$

In this equation, alpha represents the decay intercept. This intercept is the minimum decay for an encounter that will also be used as decay value for the first encounter. The decay scale parameter c determines the relative contribution of the activation dependent component. Pavlik and Anderson (2003, 2005, 2008, Pavlik, 2007) have shown in a series of studies that these equations account for a wide range of spacing-related learning phenomena.

In both the PA model and in ACT-R, the activation of a fact determines both the probability of recall of that fact and the latency associated with recalling that fact. For the probability of recall, the activation of the fact is compared to the retrieval threshold while taking into account the noise that is associated with declarative memory. If the activation of a fact is higher than the retrieval threshold, that fact can be recalled. However, if the fact is below the retrieval threshold, it is unavailable for further processing. The latency of a retrieval at a certain point in time (t) is a function of the activation (A_i) of the fact that is retrieved:

$$L_i(t) = Fe^{-A_i(t)} + \text{fixed time}$$

In this equation, F is a scaling factor and “fixed time” refers to the time cost of all non-fact-retrieval processes required in giving the answer.

Applying the Spacing Model

In Pavlik and Anderson (2008), the spacing model is used to actively determine the optimal sequence for learning a list of Japanese-English word-pairs. In this paper, Pavlik and Anderson do not take into account the testing effect, but instead determine the optimal sequence by balancing the positive effects of increased spacing intervals on probability of recall with the negative effects increased intervals have on the retrieval latency. That is, although increased spacing is always positive in the PA model, it is, because of more decay, detrimental for the retrieval latency. In the 2008 paper, the additional time required for a retrieval if the presentation of a word-pair is shifted forward in time is weighted against the predicted increase in probability of recall. This results in a series of relative complex formulae to determine the learning gains of test-trials and study-trials.

An alternative approach is to determine the optimal sequence on the basis of the activation of the word-pairs in relation to the retrieval threshold. That is, if we assume on the basis of the combination of spacing and testing-effects that the time between two encounters is optimal *just before* the activation of the fact drops below the retrieval threshold, an optimal sequence can be determined on the basis of the activation of all facts.

Algorithm 2: Default PA

On the basis of the notion discussed above, the default PA model can be used to determine the optimal spacing sequence: as soon as a fact is about to fall below the retrieval threshold, it has to be presented again. If no previously presented fact is below the threshold, a new fact can be introduced. More precisely, as it could be that a fact drops below the retrieval threshold while another fact is being tested, the algorithm computes the activation of all facts 15 seconds ahead to determine whether to introduce a new fact or present a previous one. If all facts have been introduced, the fact with the lowest activation is selected for presentation. The performance of this algorithm is highly dependent on the accuracy of the internal activation representations, which are in turn dependent on the choice of parameter values. Although the PA model has been tested on a large number of studies, no fixed set of parameter settings have emerged yet. The values for the decay scale (c) range (Pavlik & Anderson, 2005, 2008) from 0.143 to 0.495, and for the decay intercept (alpha) from 0.058 to 0.300. The threshold parameter is typically set at -0.704. As these parameters have been fit to experiments with longer study session than used in the current experiment, we explored the effects of different settings on the resulting sequences. As a too low threshold results in extended spacing (in our explorations, sometimes all word-pairs were presented before the first word-pair was repeated), we decided to raise the retrieval threshold to -0.500. Following similar reasoning, the decay intercept and the decay scale were set at .25. With respect to the latency equation, we decided against separate estimations for F and the “fixed time”. In the 2008 model, F is set at a value larger than 1 (1.29) indicating an enhanced effect of A_i on the latency. At the same time, using a “fixed time” diminishes the effect of A_i on the latency. Therefore, we set F to 1, and the “fixed costs” to 0.

Using the default PA algorithm, we can create an optimal schedule. However, this schedule will be similar for all participants: if the first word-pair is repeated after 5 trials because it will drop below the retrieval threshold within 15 seconds, this holds for all participants. Obviously, this does not match real performance profiles: some participants will have a higher overall performance level than other participants, but it might also be that some words are recalled better by participant A, but a different set of words is recalled better by participant B. Of course, the default PA model described above and the standard ACT-R equations have only been fit to aggregate behavioral data. Luckily, because each time an item is presented the learner provides us with additional behavioral data, we can use this data to fine-tune the internal representation.

Subsymbolic Model Tracing

In the traditional model tracing account (Anderson, Boyle, Corbett, & Lewis, 1990), the behavior of a student is matched against all knowledge available in a tutoring system. For example, if a student has shown accurate performance in a number of subtraction problems in which carrying is required, the knowledge in the tutoring system that represents carrying is marked as mastered. Thus, the tutoring

system keeps an accurate representation of all knowledge the student has mastered by updating the internal representation each time new behavioral information becomes available. The behavior that the learner displays can similarly be used to update the subsymbolic activation of facts.

Given that each time a student has to answer a test trial both accuracy and latency information becomes available, we can, in principle, use this information to determine what the current activation of the retrieved chunk is. If we know the latency and therefore the activation at the time of encounter j , and we also know the latency/activation at the time of encounter $j-1$, we can calculate what the decay for encounter $j-1$ should have been. By this rationale, we can minimize the difference between the predicted activation and the observed latency and use the behavioral data of the student to update our model that represents the state of the student.

However, given the general assumption that the retrieval process is inherently noisy, using this direct relation might be problematic when the response is fast. That is, when $t_j - t_{j-1}$ is relatively long and the latency for t_j is short because of a temporal boost in activation due to noise, the calculated decay for t_{j-1} will be very low (or even negative). As a very low decay results in facts that are predicted to be highly active over a very long period of time, this temporal noise-boost will ruin the scheduling of the fact. Therefore, we have chosen not to use the outcome of the algorithm described here directly, but instead change the d_{j-1} with a fixed, small amount in the direction indicated by the mismatch between predicted activation and observed latency (c.f., hill-climbing search algorithms).

Algorithm 3: Threshold-based Adaptation

Given the issues related to the noisy observations, it might not be wise to use the more fine-grained subsymbolic model tracing method described above as this might result in overfitting. To minimize the chances of overfitting, a coarser algorithm might prove beneficial. Therefore, Algorithm 3 adapts the PA model by only modifying the decay parameter for a certain encounter when at test the word-pair cannot be correctly recalled. As the system always presents word-pairs of which it estimates that they are still above the retrieval threshold, the failure to recall indicates that the estimated activation was too high. Thus, the decay for that particular item should be higher, which is reflected in increasing the alpha parameter with 0.01.

Algorithm 4: Latency-based Adaptation

The threshold-based adaption algorithm focussed on maximizing the testing-effect. Each time a fact cannot be recalled, its decay is increased, ensuring that it will be presented with shortened spacing in subsequent trials. Although this provides the optimal effect from successful retrievals, this algorithm neglects the importance of the spacing effect as it would never result in increased spacing. However, where a failure to retrieve a fact is a marker of a lower than expected activation, a too fast response is a marker of a higher than expected activation. Thus, the latency-based adaptation of Algorithm 4 extends the threshold-based adaptation algorithm by comparing the expected latency with the

observed latency. To prevent overfitting a threshold-based approach is used to determine if an activation-based mismatch is observed. That is, the decay intercept will only be changed if the expected and observed latency differ by more than 0.5 seconds. Instead of a constant modifier, the decay intercept is changed according to:

$$\Delta\alpha = \max(0.01, \frac{\text{observed} - \text{expected}}{1000})$$

where observed and expected are the latencies expressed in seconds.

Experiment

Procedure The outline of the procedure have been detailed earlier. Study trials, consisting of a French-Dutch word pair, were presented for 5 seconds. After each initial study trial, a test trial of the same word-pair was presented. During a test trial, only the French word was presented and students had 15 seconds to reply by typing in the correct Dutch translation. After pressing Enter, students were presented with a 2-second feedback screen stating “Correct”, “Incorrect” or “Almost correct” (which was given if the Levenshtein-distance to the correct answer was smaller than 3). After an incorrect or almost correct trial, participants are presented with another study trial. The four algorithms determined which word pair to present next. After 15 minutes, the participants were thanked for their participation and were left in the impression that the experiment was done. After the learning session on Day 1, all words were tested by means of a traditional paper-and-pencil test on Day 2. The post-test was graded in the typical way in which such a test is graded: each error deducts one point from the maximum score of 10. Participants were naive with respect to the experimental manipulation and did not know that they would be tested on Day 2.

All participants were tested at the same time in a classroom equipped for computer-supported education. Each participant operated his or her own computer.

Materials A list of 20 words was compiled for each class separately. All words were selected from a chapter that was scheduled to be covered in class, but was not discussed yet.

Participants Ninety-one pre-university-education level students (all students of four 3rd year HAVO/VWO classes) of approximately 15 years of age participated, of which 85 took part in both tests. Participants were semi-randomly distributed over conditions to ensure that in each class an equal number of participants used each algorithm. All participants were instructed that their results would be stored anonymously and that the results would not be communicated to their school or teachers on an individual level.

Results

Of the 85 students who took part in both sessions, six students were removed from further analyses because they did not respond in more than 5% of all trials and gave a number of answers that did not fit the instructions (e.g., “I’m bored”, or names of rock bands). Four participants were removed because their performance in terms of correct responses




Figure 1: Average grades on French per algorithm condition




Figure 2: Raw scores on post-test per algorithm condition

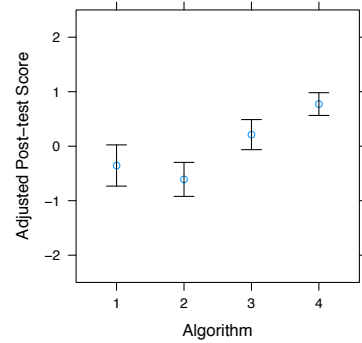


Figure 3: Post-test scores adjusted for covariates mentioned in text

All errorbars depict standard errors.

during the learning session deviated more than 2 standard deviations from the average of their group. One participant was removed because of scores on the final test that deviated more than 2 standard deviations from the average score for his or her group. This leaves 74 participants, 18 in the flashcard condition, and 19 participants in each of the three spacing conditions.

Covariates As we tested participants in a domain in which they have a significant amount of prior knowledge, it is important to control for potential differences in prior knowledge between groups. Hereto, we analyzed the students' school grades (graded on a theoretically linear scale from 0 to 10, with a 6 representing the grade required to pass that class) for French, see Figure 1 ($F(3,73)=1.27, p=0.29$). Although this effect is far from significant, the value of the F-statistic is larger than we hoped for. Therefore, we decided to include the grades for French as covariate in all subsequent analyses.

Given that we limited the amount of time to learn 20 word-pairs and the algorithm determined when a new word pair was introduced, not all participants might have seen all 20 word pairs in the non-flashcard conditions (algorithm 2 to 4). This did indeed turn out to be the case in all three conditions. The average number of word pairs presented to the participants was 19.5, 19.6 and 19.8 for the default PA, the threshold-based and the latency based conditions respectively. Although these differences (when compared to the 20 words seen by the students in the flashcard condition) fail to reach significance (ANOVA $F(3,69)=2.6, p=0.057$, post-hoc pairwise t-test with pooled standard deviations: flashcard vs default PA algorithm, $p=0.08$, all other comparisons $p > .1$), this does give the flashcard-based condition an advantage when comparing scores on the post-test, as some participants in the other conditions will not have seen all word pairs. Therefore, we also included the number of words seen by the student as covariate in subsequent analyses.

To account for possible effects associated with the session in which the

study was run or peculiarities of a particular class, another factorial covariate was included representing group.

Post-test Scores Figure 2 shows the raw scores on the post-test, and Figure 3 shows the scores on the post-test adjusted by the covariates French grade, group and number of words seen.

Analyzing the data presented in Figure 3 shows that the algorithm has a significant effect on the post-test scores ($F(3,70)=4.19, p=0.009$). Testing the individual effects by conducting pairwise comparisons using t-tests with pooled standard deviation and Benjamini and Hochberg's (1995) p-value adjustment method showed that students in the latency-adaptation group, Algorithm 4, score significantly higher than students in the flashcard ($p=0.032$) or in the PA model ($p=0.010$) group. None of the other comparisons reached significance ($p > 0.100$).

Adaptions The observed differences between the more static PA model and the latency adaptation condition suggests that the adaptations resulted in different decay patterns for different participants. Figure 4 shows the average estimation of the alpha parameter associated with the last encounter per word-pair. As can be seen, different participants required different alphas, with, for example, participant 1 and 4 requiring relative low decay values and participant 13 requiring a very high decay value. If these three participants would have been set at the average alpha (0.259), the estimated activation for participant 13 would be too high, resulting in many retrieval failures - and violating the testing-effect constraints. At the same time, participants 1 and 4 would have had a too low estimated activation, resulting in a sequence with too low spacing, violating the spacing constraints.

Furthermore, note that although fast responses could have resulted in an alpha value below the starting value of .25, all participants have a higher average alpha.




Figure 4: Effects of adaption on latency per participant

Discussion

The current study set out to answer two questions. The first was to test whether the results obtained in the scientific literature on the spacing effect would also hold in the more real-life case of learning a small set of vocabulary items in a small period of time. The second question was to examine whether the learning gain would improve when the algorithms that construct the learning sequence take individual differences into account. With respect to the first goal, the significant difference between the flashcard and the latency adaption conditions illustrates that a learning sequence that is based on an algorithm that takes spacing and testing-effects into account outperforms a more traditional flashcard sequence. However, *only* the condition that optimizes the sequence on the basis of individual latency differences outperforms the flashcard condition, answering the second question.

It is striking to see that the traditional PA spacing condition scores - in absolute terms - worse than the flashcard condition. This result for the traditional spacing condition might be caused by the parameter settings chosen for this study. It might be that with an alternative parameter setting the PA model also outperforms the flashcard condition. However, it is difficult to come up with the settings in parameters required. The first candidate for change would be the retrieval threshold, as in most PA studies, the threshold is set at -0.704 instead of -0.5. However, lowering the threshold would increase the spacing between two presentations of the same item. As the difference in adjusted post-test score between the PA algorithm and the threshold adaption algorithm is close to reaching significance ($p=0.107$, two-sided, but one could argue that this difference should be tested one-sided), and the sole difference between these conditions is that the individual decay settings are possibly increased (which has a similar effect as *increasing* the retrieval threshold), it is difficult to imagine how decreasing the retrieval threshold would improve the performance. With respect to changes in the parameters involved in calculating d_{ji} , it is most likely that these changes would benefit the other algorithms as well. Thus, although changes in the parameter settings might diminish the gap between the different spacing algorithms, it is hard to imagine how the traditional PA model would outperform the alternative. It is interesting to note that Pavlik and Anderson (2008, p.102) discuss a very similar approach they call "performance tracking" and mention that this method will "add considerable power". However, they conclude that this approach will make scheduling much more complex.

We have shown in this study that performance tracking is possible, but also that adapting the sequence to the characteristics of individual learners improves learning gains considerably, even if the learning session takes only 15 minutes.

Acknowledgments

Parts of this work has been conducted in the context of Van Woudenberg's Master thesis (Van Woudenberg, 2007). The authors would like to thank the students and teachers of het Belcampo College, het Gormarus College and het Werkman College for participating in the study reported here.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111 (4), 1036–1060.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2 (6), 396–408.
- Baddeley, A. (2003). Working memory: looking back and looking forward. *Nature Reviews Neuroscience*, 4 (10), 829–839.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B*, 289–300.
- Carrier, M., & Pashler, H. (1992). The influence of retrieval on retention. *Memory and Cognition*, 20, 633–633.
- 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.
- Cepeda, N. J., Vul, E., Rohrer, D., Wixted, J. T., & Pashler, H. (2008). Spacing effects in learning: A temporal ridge-line of optimal retention. *Psychological Science*, 19, 1095–1102.
- Dempster, F. N. (1988). The spacing effect. *American Psychologist*, 43, 627–634.
- Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology*. Translated by Henry A. Ruger and Clara E. Bussenius (1913).
- Pashler, H., Rohrer, D., Cepeda, N., & Carpenter, S. (2007). Enhancing learning and retarding forgetting: Choices and consequences. *Psychonomic Bulletin & Review*, 14 (2), 187.
- Pavlik Jr, P. I., & Anderson, J. R. (2008). Using a model to compute the optimal schedule of practice. *Journal of Experimental Psychology: Applied*, 14 (2), 101–117.
- Pavlik Jr, P. (2007). Understanding and applying the dynamics of test practice and study practice. *Instructional Science*, 35, 407–441.
- Pavlik Jr, P., & Anderson, J. (2003). An ACT-R model of the spacing effect. In *Proceedings of the Fifth International Conference of Cognitive Modeling*, Bamberg, 177–182.
- Pavlik Jr, P., & Anderson, J. (2005). Practice and forgetting effects on vocabulary memory: An activation-based model of the spacing effect. *Cognitive Science*, 29 (4), 559–586.
- Roediger, H. L., & Karpicke, J. D. (2006). Taking memory tests improves long-term retention. *Psychological Science*, 17 (3), 249–255.
- Van Woudenberg, M. (2008). *Optimal word pair learning in the short term: Using an activation based spacing model*. Unpublished master's thesis, University of Groningen.
- Wozniak, P. A., & Gorzelanczyk, E. J. (1994). Optimization of repetition spacing in the practice of learning. *Acta Neurobiologicae Experimentalis*, 54 (1), 59–62.