Optimizing Practice Scheduling Requires Quantitative Tracking of Individual Item Performance

Luke G. Eglington¹ & Philip I. Pavlik Jr¹

¹Institute for Intelligent Systems
University of Memphis

**This is a preprint of an article published in *npj Science of Learning*.
The final authenticated version is available online at:
<http://doi.org/10.1038/s41539-020-00074-4>**

Abstract

Decades of research has shown that spacing practice trials over time can improve later memory, but there are few concrete recommendations concerning how to optimally space practice. We show that existing recommendations are inherently suboptimal due to their insensitivity to time costs and individual- and item-level differences. We introduce an alternative approach that optimally schedules practice with a computational model of spacing in tandem with microeconomic principles. We simulated conventional spacing schedules and our adaptive model-based approach. Simulations indicated that practicing according to microeconomic principles of *efficiency* resulted in substantially better memory retention than alternatives. The simulation results provided quantitative estimates of optimal difficulty that differed markedly from prior recommendations, but still supported a desirable difficulties framework. Experimental results supported simulation predictions, with up to 40% more items recalled in conditions where practice was scheduled optimally according to the model of practice. Our approach can be readily implemented in online educational systems that adaptively schedule practice, and has significant implications for millions of students currently learning with educational technology.

Introduction

There are large potential benefits to applying findings from the cognitive science of learning to educational contexts. Two especially promising findings are spacing and retrieval practice (testing). Spacing practice trials across time has been shown to benefit memory [1, 2]. Testing has also been shown to benefit memory, relative to restudying [3, 4]. Taken together, spaced retrieval practice has been shown to provide substantial benefits to later memory [5]. Successfully integrating these findings into educational technology could benefit millions of students that currently learn in online educational systems. However, implementation of these findings has been elusive - research has been unclear regarding exactly how much and when to practice specific information. We demonstrate how standard methods to uncover optimal practice schedules are confounded by fundamental characteristics of memory, and that any single conventional schedule (see Figure 1) applied to a set of items will be inherently suboptimal. We provide evidence that adaptive schedules lead to superior performance, better accommodate item and learner variability, and are implementable with a mixture of computational memory models and relatively simple economic decision rules.

Figure 1. Several conventional spacing schedules.

Research on Spacing Schedules

Researchers have explored the effects of various spacing intervals on memory [6-8], frequently while enforcing fixed trial durations and numbers of trials. Figure 1 depicts several popular spacing schedules compared in the literature: the inter-trial interval (ITI) could be a fixed value (first row of Figure 1), gradually increase over time (second row), or decrease over time (third row). Those three cases in Figure 1 have the same average interval, but the distribution of the spacing could influence performance and difficulty. For instance, the longer initial retention interval for the uniform condition (Figure 1) may increase potential learning, but also entails a higher risk of retrieval failure. More generally, there is a tradeoff between the potential gain from more spacing and the risk of retrieval failure. This tradeoff is relevant to a popular theoretical explanation of the

spacing effect named *study-phase retrieval*, which states that retrieving memories of previous exposures to an item strengthens the memory trace [9, 10]. Retrieval difficulty is thought to influence the benefit - more difficult retrieval is thought to better strengthen a memory, but is riskier since it may fail.

There is also empirical evidence that practice can *slow* forgetting [11], which implies that as practice accumulates retrieval becomes easier. This finding of slowed forgetting, in combination with study-phase retrieval theories, implies that gradually increasing difficulty (via spacing) may balance difficulty and lead to better memory via an expanding schedule. Surprisingly, research findings have been more mixed than may have been predicted by the theories and empirical results described above. For instance, Landauer and Bjork [12] found that an expanding schedule of test trials without feedback resulted in better memory at a final test than both uniform and contracting schedules. In contrast, Karpicke and Roediger [13] showed that providing feedback could lead to a benefit of uniform scheduling. Recently, Kang et al. [14] found that expanding intervals provided higher average recall probability across multiple sessions, but equal memory to uniform intervals at a final test.

Heuristic Schedules

Some heuristic schedules that adapt to item-level differences have been evaluated, such as dropping an item from practice once it has been answered correctly a specific number of times [15]. These methods offer the possibility to shift practice time from items that are easier (have been recalled successfully) to those that are more difficult.

Using vocabulary word-pairs for learning materials, Pyc and Rawson compared the memorial benefits of conventional spacing schedules (e.g., a uniform schedule) to those that dropped items from practice after being successfully recalled. They found that a drop schedule in which word-pairs were spaced by 23 intervening pairs and were dropped after one successful retrieval resulted in superior final test memory than conventional schedules. However, variation in item difficulty and participant performance make it unlikely that the number of intervening items or number of correct answers to allow before dropping (1) would be generally optimal. Drop heuristics are also silent regarding when an item should be reintroduced after being dropped from practice.

A promising approach that values efficiency is the Region of Proximal Learning framework (RPL). RPL recommends that students should choose to practice items that are “..slightly beyond the individual’s grasp” [16]. Relatedly, Tullis et al. [17] showed that allowing students to choose what to study provided significant learning benefits. A challenge for many of these frameworks is that it is not clear when previously practiced items should be reintroduced. If left up to the student, then their subjective judgments would need to be well-calibrated. Allowing student choice may also introduce additional time costs: the candidate study items would need to be presented in some fashion for one to be selected. A related challenge with RPL, common to many approaches, is one of implementation. Time-cost is considered important and used to justify the practice of easier items [16], but difficulty and time costs are not quantified which makes it unclear how to apply the principles broadly. Additionally, whether easier items are more efficient

depends on the learning gains from successes and the informativeness of feedback when failures occur. Considering time cost could reveal in some cases that easy items are preferable (when success causes strong learning and feedback is more time-consuming), but in others that more difficult items should be practiced instead (if feedback after failures is informative and relatively quicker).

Practice Difficulty and Time Costs

Many proposed practice schedules manipulate difficulty in search of improved learning outcomes. However, there is a frequently under-appreciated consequence of varying difficulty that bears directly on the educational value of practice schedules. The most important negative consequence of difficulty is the time needed to engage with it, since students and educators from a practical perspective are limited in the time they can spend being instructed or instructing. Because of this reality, gains in learning are typically valued in relation to their cost, and so costly learning methods should be relatively less valuable and attractive than faster methods. Considering both cost and gain leads to practice decisions that are sensitive to the value of the item (in terms of learning gains per second). If practice is too difficult, learners may take longer to recall when correct, and they will more consistently fail, which will frequently be time consuming due to review costs. This will lead to less learning per unit of time spent practicing.

The complexity of scheduling practice is increased when we consider that there is not one dependent measure, learning gain, but rather two dependent measures: learning

gain and learning time cost. Both are objectively important to the student or educator and relevant to determining optimal schedules of practice. Although the number of trials (and sometimes their durations) are preset in many experiments, in educational contexts there are no such necessary constraints. Fixing trial durations may serve to mask the increased time it can take to recall harder items or decreased time to recall easier items, which has important implications in a natural learning situation.

Microeconomics provides the conceptual tools to consider the gain from a trial along with its cost (in time); maximizing gain and minimizing costs is a common concern in economics research. In short, determining the optimal practice schedule is a fundamentally economic problem: find the maximally efficient level of difficulty that provides the most gains in learning per unit of time. To motivate the experiment and subsequent simulations, below we describe the important relationship between practice and latency we use to compute efficiency. We then further elaborate on the economic principles that motivate designing schedules based on efficiency, in contrast with conventional practice schedules that frequently have focused on gain unconditional on time cost.

Practice and Response Time

There is strong evidence that as practice accumulates, the time it takes for a participant to respond decreases (response time, henceforth RT). Correct trials are also frequently faster than incorrect trials, especially if feedback is required after errors [18]. As Kahana and Loftus [19] describe, correctness and RT are typically considered in

isolation even when they can be mutually informative. As we will show below, their relationship is important when determining a schedule of practice for a finite study period, and use of only one measure can be misleading. The reliable relationship between practice and RT has been shown in many contexts and paradigms [20, 21]. In particular relevance to the present work, many researchers have shown strong relationships between practice and RT with declarative concepts. Anderson, Fincham, and Douglass [22] demonstrated this consistent relationship with participants learning declarative concepts and rule application. Anderson [23] modelled RT as an exponentially decaying function of memory activation (Equation 1a). Newell and Rosenbloom [24] have shown how RT can follow a power function of prior practice attempts (Equation 1b).

$$(1a) \quad RT = Fe^{-a} + fc \qquad (1b) \quad RT = BN^{-c}$$

In Equation 1a, F is a scalar, a is memory activation, and fc is the RT asymptote. In Equation 1b, B is the difference between RT on the first trial and the asymptote, N is the number of practice trials, and c is a parameter that determines how quickly RT approaches asymptote.

The decrease in RT as practice accumulates is highly relevant to contexts in which efficiency is desirable. This importance remains regardless of whether the relationship follows a power law or an exponential decay curve, the critical point is that the *efficiency* of practice (in terms of time) is related to when and how much practice has already occurred. The importance of the efficiency of practice is mentioned in research concerned with improving learning outcomes with study techniques such as testing and spacing, although RT does not usually influence interpretation when comparing conditions [15, 25-27]. To optimize a schedule of practice gain relative to time spent, the relationship between practice latency and gain must be included in the analysis. Longer

spacing intervals may increase the effect size of spacing, but if a wide interval takes substantially longer overall than a narrower interval, it may be a suboptimal use of a learner's time. For instance, Yan, Eglington and Garcia [28] found that a very wide spacing interval provided comparable (or slightly higher) memorial benefits than a narrower spacing schedule. However, response times during practice were significantly longer in the wider spacing schedule, resulting in the wider spacing interval being slightly *less* efficient per second of practice.

Although the primary measure of interest in many studies on spacing is final test performance, correctness *during* practice is also critical because it is an important determinant of efficiency. For instance, while feedback after failed retrieval attempts increase the benefit of testing [29], this feedback takes time. However, If a participant is correct, feedback may be unnecessary [30]. Thus, in the context of vocabulary learning and other paradigms where processing and responding to the stimulus are fast, correct trials are often a faster and relatively more efficient use of a learners' time.

In sum, if the efficiency of a practice attempt is defined as *the gain in retrievability at some final test divided by the time cost of the current practice*, there is strong theoretical reasoning and empirical evidence that the efficiency of an individual practice attempt will vary as practice accumulates for most learning tasks, depending on the number of prior practice attempts and their temporal spacing.

Economic Efficiency of Practice

When using a model of practice to schedule learning activities the goal is to use model predictions to make pedagogical decisions that maximize learning. One approach would be to choose the item which will provide the most gain (e.g., in the probability of recall) at a final test. However, there are several ways to measure gain with different shortcomings. For instance, percent gain could be maximized [31], logit gain, or an efficiency score [32] could be composed from either as a measure of gain. In economics, this choice of the gain metric is a choice about how to measure utility (or value) in the domain. Utility in learning represents some way to order our preferences for learning events (practice trials) such that we can compare which learning events are more useful than other learning events. Adapting this term to education, we are seeking a measure of learning value that we can use to compare different possible outcomes of instruction in terms of learning gain to guide practice scheduling decisions.

However, choosing a gain metric without considering time costs can result in very inefficient practice. This inefficiency is because if maximal percentage gain in recall probability is desired for each trial, more difficult (lower probability items) will be chosen, despite the high rates of error which will increase the average time per trial. In practice, selection focused on higher error rates (to maximize gain in terms of probability) tends to be only effective when time costs are ignored and inefficient otherwise [18, 33].

Considering both variable cost and valuation leads to an alternative formulation of gain - the *efficiency* of practice. Computing the efficiency of practice is an inherently economic calculation in which the expectancy of the choices' outcomes are relevant. Learning

tasks are frequently tests where the attempt can fail or succeed, and each of these outcomes has different consequences. Calculating efficiency requires a measure of learning utility and a measure of the cost, typically in terms of time (see Equation 2).

$$U = (p) \frac{\text{correct gain}}{\text{correct cost}} + (1 - p) \frac{\text{incorrect gain}}{\text{incorrect cost}}$$

Equation 2. An economic interpretation of practice utility. p is the probability of correctness. Gains and costs for correct and incorrect trials can be quite different (and nonlinear), and thus practice difficulty is highly relevant to optimizing practice scheduling.

Maximal efficiency as the goal of scheduling practice

If the success or failure of a trial influences efficiency (incorrect trials can be more time consuming), at what probability of recall (difficulty) should items be practiced to maximize efficiency? In the present study, we aimed to find this optimal efficiency threshold (OET) via simulation, which requires a model that can accurately estimate the probabilities of recall and gain from practice and spacing. We subsequently tested the results of the simulation with an experiment. In the following sections, we describe candidate models as well as how they can be used to accommodate our efficiency-oriented approach.

Model-Based Practice Scheduling

Several excellent computational models of memory have been proposed [34-36]. One of the primary difficulties has been to account for the effect of repeated practice, spacing, forgetting (interference), and the cross-over interaction between practice spacing and the retention interval between the practice session and the final test [36, 37]. Several computational models described below track the most critical patterns found in studies manipulating practice and spacing.

Pavlik and Anderson [18, 38] attempted to account for the complex interactions associated with spacing practice over time by assuming decay to be a function of base-level activation. Higher activation at the time of practice lead to faster forgetting, such as when practice is massed together, and slower forgetting when practice was spaced. This adjustment allowed them to account for many aspects of practice and spacing. Lindsey et al. [39] introduced a model of spacing and forgetting, and also demonstrated how it could be used in an applied setting. Finally, recent work by Walsh et al. [35] improved upon prior models of spacing (i.e. Pavlik and Anderson) by accounting for accelerated relearning after extensive forgetting. A powerful aspect of their model was that average elapsed average spacing of practice since a practice trial was treated as a weighted decaying function. Their model compared favorably with alternative models.

Computational models of memory can account for many well-known patterns found in spacing studies, and can accurately predict the probability of recall. There is great interest in optimally scheduling practice, but most practice recommendations are insensitive to time costs and item and student differences. This oversight means that suggested practice schedules are likely to be inherently suboptimal. Recent advances in computational modeling allow us to move past conventional schedules because their sensitivity to time costs and individual differences allows precise estimation of efficiency.

Although time efficiency is critically relevant for students, it is only occasionally discussed in the context of scheduling practice (e.g., Metcalfe & Kornell [16]). It has rarely been included in calculations of optimal practice scheduling according to a model. For instance, Khajah, Lindsay, & Mozer [40] simulated the effects of using difficulty thresholds to decide practice and found that practicing items close to being forgotten (probability of recall = .40) resulted in superior memory performance. This recommendation is similar to Bjork's [41] argument that items should be practiced when they are about to be forgotten. However, if a student's time is a valuable resource, such recommendations by Bjork and Khajah et al. are contradicted by research showing that information is recalled substantially faster at higher probabilities [20, 42]. This aforementioned recall probability-latency relationship implies that while a policy such as $p = .40$ may improve memory it may be inefficient relative to a policy with a higher probability decision rule (e.g., $p = .80$). In short, we agree with the policy of guiding practice based on a predicted probability threshold, but we believe that the time cost of practice at these probabilities needs to be considered when determining an optimal policy.

In sum, conventional spacing schedules can be suboptimal relative to an adaptive scheduling algorithm due to variable item difficulty, the slowing of forgetting as practice accumulates, and the interaction between difficulty and response time. Response time is especially relevant if total study time is constrained, but trial duration is not. We also highlighted the possibility of an optimal amount of difficulty to maximize learning, based on arguments by Hebb [43], Bjork [41], and Pavlik & Anderson [18]. Tracking probability

of retrieval (our operationalization of difficulty) throughout learning was also shown to be achievable, both with our methods and given the existence of several impressively accurate models of memory that account of practice attempts and spacing [38, 40, 44, 45]. The following simulations combine these ideas by evaluating how scheduling according to a model and difficulty thresholds compared to conventional practices schedules and heuristics. An important aspect of our approach was to also simulate response time, which led to different conclusions than prior research. Finally, we tested the predictions of the simulation with experimental conditions intended to the relative benefits of valuing efficiency versus valuing discrepancy reduction (focusing on more difficult trials).

Results

In the following simulations, simulated students practiced according to various schedules of practice. Correctness probability and response times were estimated by models with parameters fit to a dataset in which participants practiced Japanese-English word pairs in one 22 minute session, and were tested in a second session that was 2 minutes, 1 day, or 3 days later. In that dataset participants learned the word-pairs at eight levels of spacing (ranging from 0 intervening items (massed) up to 16 intervening items (spaced), and 3 levels of repetition (2, 4 or 8 trials). Additional details of that dataset and model fitting procedures are described in the supplemental materials.

Each simulated student completed a practice session and a test 3 days later. The primary constraints were A) the *total* duration of the practice session (22 minutes, see

methods) and B) that feedback was provided when a retrieval attempt was unsuccessful (which lasted four seconds). We show how traditional practice schedules (e.g., uniform spaced schedule) as well as heuristic methods fare [15]. We also show how conventional schedules perform when trial duration is not fixed, but is determined by the speed of retrieval (faster if successful). Finally, we show the effects of scheduling practice based on different Optimal Efficiency Thresholds (OETs) that schedule practice according to the probability of correctness.

In all OET conditions, a simulated student practiced whichever item closest to (but less than) the recall probability threshold was practiced. If all items were above the threshold, the item closest to the threshold was practiced. After each trial, the model would update the probabilities for all items (based on the outcome of the trial as well as elapsed time), and the process would repeat until the time ran out.

The primary goals of the simulations were to a) recreate typical patterns of results found in prior work, and b) evaluate the effects of using a computational model to schedule practice at different OETs that may balance the benefit of spacing with the time costs associated with reduced recall probabilities due to forgetting over longer intervals. We found that utilizing a computational model of memory in conjunction with economic principles of efficiency provided superior performance than conventional spacing schedules or heuristics.

Simulation Results

Figure 2. Simulated final test performance after a 3-day delay. There were 200 simulated students per condition. Schedules on the left denote the probability threshold (OETs) for different simulated conditions. 0, 15, and 30 denote spacing in terms of trials (0 being massed). Text on bottom axis shows which simulations fixed trial duration or allowed them to be determined by the latency model (i.e., correct trials being faster than incorrect). Drop1 denotes a heuristic in which items were dropped from practice after successful retrieval, until all items were retrieved and the process restarted. Error bars denote ± 1 SEM.

Figure 2 depicts many of the critical simulation results. All spaced conditions were superior to the massed condition. Conditions with more spacing also provided a larger benefit than those with less spacing. These results are broadly consistent with what is found in studies comparing conventional spaced and massed conditions [2]. The conventional schedules in which trial duration was determined by the latency model were better than if trial duration was fixed; more trials were possible, especially if simulated students were successful more often. As a result, simulated student skill was highly correlated with the number of trials experienced in conventional schedules without fixed trial durations (see Figure 6). The best performing conditions were those that adaptively selected which item to practice on each trial according to model predictions and an OET. The best OET in this simulation was $p = .94$. It is important to

note that this does not mean items were always practiced at this probability, especially early in practice (see Figure 3).

The benefit of this approach was apparent in many of the OET conditions. OETs of .8 or greater had an average of 39% better memory at final test than the next best conventional schedule with a maximum of 55% improvement. This result implies scheduling according to item difficulty and prior practice history while considering efficiency provides superior performance, within the same timeframe, as conventional schedules. Note the achieved probability of the conventional schedule in Figure 3. On average, probabilities remained fairly low because item difficulty and simulated student ability played no role in deciding what should be practiced next. As a result, items were frequently too difficult. The optimal zone of difficulty we found supports prior suggestions that some difficulty is beneficial [41, 43], with the practical benefit that our method quantified the amount of difficulty, which can in turn be leveraged in combination with a model. Interestingly, the optimal zone that we found indicates imposing less difficulty than previously suggested [40, 41]. This is due to our methodology considering the time costs of difficulty.

Figure 3. Rolling recall probability (window = 5 trials) of practiced items across 5 conditions. Note that an OET doesn't dictate the actual probability of recall for practiced items, but is a threshold determining what is practiced (closest item underneath threshold). All conditions had the same total time to complete as many trials as possible. Rolling mean computed up to the number of trials N completed by 95% of simulated students in each condition. Gray, blue, green, and black dashed lines denote the respective OET thresholds.

Inspection of item-level model behavior provides more examples of the adaptivity of the efficiency-based approach (model + OET). Figure 4 depicts the strong relationship between item difficulty on the x axis (item intercept) and average spacing, and how it differed between high and low OET simulations. Easier items (higher intercept) experienced wider spacing intervals. This supports findings by Metzel-Baddeley & Baddeley [46] in which they showed that applying narrower spacing to harder items benefitted later memory. As practice progressed, spacing for individual items also widened, as shown in Figure 5. This adaptivity means difficulty adjusted as items were learned, due to the items probability of recall keeping it above the OET (and thus out of

contention for practice) for longer and longer intervals of time. Note the positive trend in Figure 5, this indicates an expanding schedule naturally develops within the adaptive scheduling condition. Given the reasonable assumption that items vary in difficulty, non-adaptively conventional schedules are likely to provide practice schedules that are too easy for some items, and too difficult for others (and thus inefficient).

Figure 4. The relation between item difficulty and achieved spacing in OET40 and OET94 conditions. In both conditions, easier items (positive intercepts) have wider spacing, while harder items (negative intercepts) have narrower spacing. The increased spacing in OET 40 leads to worse performance. Lines depict a linear regression fit.

Figure 5. The average number of intervening trials since the last attempt as a function of the number of prior attempts in OET 40 and OET 94. In OET 94, as practice accumulates for an item, spacing interval tended to increase. The opposite was true for OET 40, spacing intervals contracted. Each condition plotted to the number of attempts that at least 95% of simulated students attained in the condition (6 and 10 for OET40 and OET94, respectively).

Simulated student-level skill also influenced the behavior of the adaptive model behavior. In the conventional schedule in which trial duration could vary (and thus the number of trials), the skill of the simulated student (here as the intercept) was strongly correlated with the number of trials (see Figure 6). Thus, “better” simulated students got more opportunities to practice. In contrast, in the adaptive conditions, there was no correlation between simulated student intercept and the number of trials they experienced. The adaptivity of the model meant that difficulty was kept consistent across ability levels and item difficulties. Lower ability simulated students would receive

relatively more practice on easier items than higher ability simulated students. In other words, the adaptive model + OET is fairer, on top of leading to superior recall.

Figure 6. The relationship between simulated student intercept and number of practice trials in two self-paced conditions, an optimal threshold condition (OET 94) and a conventional spaced practice schedule (self-paced with average spacing of 15 trials). Lines indicate linear regression fits. The strong interaction highlights a consequence of the lack of adaptivity in the conventional schedule that leads to better simulated students receiving more practice trials. In contrast, adaptive scheduling led to similar numbers of practice opportunities across various simulated student intercepts.

Experimental Results

This following experiment was intended to test several related predictions of the simulation. First, we expected high OETs to lead to better memory retention than both conventional schedules and low OETs. However, we also expected OETs beyond the optimal value (e.g., $OET = .98$) to lead to worse performance on a delayed test due to shorter-spacing and overpractice of a subset of items. We had participants practice in one of six different conditions. Five were adaptive ($OETs = .40, .70, .86, .94, .98$) and were chosen to test the hypothesis implied by the simulation that a difficulty threshold of

.94 was optimal, with additional conditions .40, .70, .86 and .98 to facilitate estimating the shape of the optimality curve with real data. The nonadaptive condition was intended as an additional control (referred to as S15¹), and repeated items every 15 trials on average. Trial duration was self-paced, except for feedback after incorrect answers that lasted 4 seconds. OET40 was intended to test an alternate theory of optimal practice in which practice was directed towards more difficult items [33].

As can be seen on Figure 7, final memory test performance in the OET conditions followed a skewed inverted-U shaped pattern as predicted in the simulation. A one-way non-parametric ANOVA (Kruskal-Wallis) indicated a significant difference across conditions in average performance on the final test, $\chi^2 = 16.5$, $p < .01$. Post hoc comparisons indicated this difference was driven by the OET86 and 94 conditions having significantly higher recall performance than the OET40 (ps .017 and .023 respectively), OET98 ($ps < .011$), or VS15 conditions ($ps < .01$). There was not a significant difference between OETs 94, 86, and 70, $ps > .17$. OET 70 was not significantly different from other conditions, $ps > .07$. OET86 was numerically higher than OET94 (M s .48 and .47 respectively), and highest overall, but was not significantly different than OET94, $p = .64$. We also compared performance across conditions with a mixed-effects model with random intercepts for items and participants. This analysis indicated the same pattern of statistically significant differences. Including condition as a factor significantly improved the model relative to a model with only random effects, $\chi^2 =$

¹ Initial piloting indicated that maximal spacing (practicing all other items before repetition) was very difficult for participants. We reduced the spacing to make the task more achievable and less frustrating. There is also evidence that non-maximal spacing within a practice session can be better for retention [2], and thus our adjustment may have strengthened the control condition.

18.72, $p = .002$. Final test performance of both OET86 and OET94 was significantly higher than OET40, OET98, and VS15, $Z_s > 2.83$, $ps < .005$.

In addition to predicting an advantage of high OETs over lower OETs and nonadaptive schedules, the simulation predicted a skewed inverted U relationship between OET and memory retention at final test. As a basic test of this prediction we compared the fit of a linear relationship between OET condition value (as a numeric predictor) and average final test performance of each participant to a model with an additional quadratic term. A likelihood ratio test indicated that the model with an additional quadratic term was a significantly better fit ($\chi^2 = 23.47$, $p < .001$), and importantly predicted that performance would eventually decrease as the OET approached 1 (maximally easy). In other words, some difficulty is desirable. The model used to schedule practice did fit the data well, McFadden's pseudo- $R^2 = .37$. McFadden's scores from .2 to .4 are considered evidence of a good fit [47]. There was a strong correlation between predicted and actual average performance on the final test, $r = .775$, $p < .001$. An analysis of the sensitivity and specificity of the model also indicated a good fit, $AUC = .871$, 95% CI = .868–.874.

The OET94 condition led to approximately 40% higher memory retention than more conventional scheduling (VS15), as well as the OET40 condition which served as a test for a different approach in which higher difficulty items are practiced more often (e.g., a discrepancy reduction approach). The simulation was reasonably successful at predicting the relative ordering and differences among conditions, but underestimated performance overall (e.g., experimental OET94 was $\sim .47$, but was closer to .4 in the

simulation). This discrepancy was more pronounced in conditions with higher per-trial difficulty, such as OET40. This discrepancy could be due to a larger gain from failures or a larger spacing effect than estimated from our model and the previous dataset.

Figure 7. Experiment results with overlaid original simulation predictions. †, ¥, and * denote the condition being significantly different than OET40, OET98, and S15, respectively. Error bars represent +/- 1 SEM.

Discussion

The present research used a computational model to simulate the outcomes of 6 conventional and 22 model-scheduled practice conditions that used different hypothetical OETs. OET94 was predicted to provide the best recall performance at a final test for learning Japanese-English vocabulary word pairs, with several other high OETs providing comparable results. An experiment confirmed that high OETs (.86 and .94) conferred significantly better memory retention at a final test than conventional

scheduling, and were also superior to adaptive scheduling that focused practice on more difficult items (OET40) or easier items (OET98). The general conclusion is that given a correct OET, model-scheduled practice can allow selection of items for practice in a way that is optimal for learning. Our results also provide evidence that practice time is highly likely to be misallocated in all conventional practice schedules that are insensitive to participants' history of practice and the difficulty of the items. Similarly, targeting too low or high an OET with an adaptive schedule may be inefficient.

The OET methodology offers several benefits. For one, easier items will go beyond the decision rule threshold sooner and (temporarily) be excluded from practice. Other harder items will then be more likely to be practiced. The OET method also includes a natural mechanism for when new items may be introduced - when all other items are currently above the OET. The issue of when to introduce new items is critical when evaluating model-based scheduling algorithms. As practice accumulates, inter-trial spacing increases between repetitions in OET-guided scheduling conditions. Thus model-based scheduling naturally leads to an expanding schedule. The critical difference that distinguishes the conventional expanding schedule from the model-based schedule is that in the conventional schedule the exact interval between repetitions is identical for all items regardless of difficulty. Here our expanding schedules were adaptive based on the dynamic model predictions at the item level.

Implementation

One rebuttal to the proposed model-based approach may be that it requires too much information to be useful - specifically the model parameters. However, it is quite rare for content to be taught for the first time to a student. There is frequently prior data from which to estimate model parameters and intercepts of the items themselves. The difficulty is in whether the dataset contains the necessary conditions from which to fit the model (e.g., variety in spacings and repetitions), or if it is biased in some way that precludes meaningful model fitting. These potential issues can be overcome by embedding experimental conditions within educational systems [45] or running controlled experiments during the educational system design process. Our experiment provided evidence that adaptive practice based on difficulty estimates from a model along with difficulty thresholds (OETs) can provide substantial memorial benefits. This finding provides a strong example to illustrate our point that scheduling practice according to a model and an efficiency threshold (OET) can provide meaningful benefits. While more data across a wider variety of spacing intervals and learning contexts will be helpful to strengthen these arguments, there is reason to believe these methods will generalize, at least to tasks that have a similar structure, where failure has substantially large time costs than success and where forgetting is a substantial effect.

Limitations and Future Directions

In the present study we focused on word-pair learning as a simplifying assumption. Such materials are common in cognitive psychology research on spaced practice schedules, and allow easier interpretation of results due to the independence among the items. Of course, learners do need to learn vocabulary and sets of facts, but those materials differ in ways that justify testing the OET approach with other materials. For

one, the speed and learning gains of successes and necessity of feedback with paired associates means that a high OET will tend to perform best. But with materials in which failures provide larger gains, or are faster, lower OETs may be preferable (see Figure 8). Additionally, another important property of other educational materials is that they contain critical interrelations that are important for the participant to learn. Complete computational models of memory and spacing will need to account for such interrelations, and model possible learning transfer among related concepts. An important extension of the present research will be to account for transfer among items to improve adaptive scheduling of educationally relevant learning materials and the possible resultant changes in scheduling and OET selection.

Figure 8. The original simulation (in black) as well as two additional simulations in which the efficiency of failures is increased. An alternative simulation in red shows performance if feedback time cost for failures was set to 0 seconds (instead of 4). Another simulation in blue shows performance if failures provided 25% more learning gains than successes (but time cost was same as original). Optimal threshold for

paired associate learning remains high for all of them, but the relative benefits over lower OETs is clearly influenced by gains from failures and time cost. S15 (repetitions every 15 trials, trial duration not fixed) was also simulated under the three sets of assumptions and serves as a point of comparison.

A further improvement could be to adjust model parameters based on prediction errors [48]. Relatedly, recording students' preferred study targets may be helpful for improving the model, there is evidence their metacognitive judgments can positively improve practice selection [17]. In both cases, additional experimentation would be necessary to determine how these improvements could be introduced. However, this limitation to the adaptivity of the model doesn't detract from the general conclusion that scheduling practice based on the most *efficient* threshold is superior to any conventional schedule, just that the model to estimate probability could be adjusted to adapt to its own performance.

Readers may have inferred that sometimes items would be practiced fairly far from the intended OET, typically early in practice. The efficiency of trials early in practice (or for items that are especially difficult) could be improved by adaptively manipulating the strength of retrieval cues. For instance, Fiechter & Benjamin [49] demonstrated how adaptively changing cue strength can benefit memory (e.g., cue - Targ__ instead of cue - ____). Adaptive cueing could be integrated into the model+OET framework by estimating the probability of recall for each item with different cue strengths. This would change the decision rule from selecting among one probability per item to each item having multiple probabilities due to varying cue strengths. A newly introduced item would have strong cues (e.g., cue – Targ__), gradually weakening until “cue - ____” was the only cue under the OET for that item.

Finally, features of the adaptive learning context (e.g., variable numbers of trials and lower difficulty) may meaningfully influence performance and need to be studied further. For instance, in high OET conditions, there is likely to be less interference, because new items are gradually introduced, reducing the average difficulty of intervening practice. In other words, at any given moment most items that could interfere with retrieval tended to be more well-learned as a function of OET, which was not necessarily the case in traditional practice schedules and also not true of our original parameterization experiment. The model we used did not account for this varying interference during practice, and may explain why our simulation underestimated the benefit of high OETs, similar to results of Pavlik and Anderson [18] which showed better learning of fixed schedule items under conditions of reduced difficulty of the surrounding adaptively scheduled practice.

Conclusion

We have emphasized throughout that we are not arguing for the superiority of a particular computational model of memory, or even for a particular model to be used for scheduling practice. There are several that have been shown to accurately model spacing and practice effects in various learning contexts [34, 35, 38, 50]. Therefore, in our view, accurately estimating recall probability is not the primary issue. Rather, our focus has been on how these computational models allow scheduling to be sensitive to individual differences and *efficiency*. We focused on using the model to estimate recall probability, given a history of practice, and on using those estimates to practice items closest to a target threshold of optimal difficulty (or probability of correctness). Our

argument differed from prior research in how we believe an optimal threshold should be chosen, and how the efficacy of practice should be evaluated. Considering time cost in our simulations (making them sensitive to efficiency of practice) led to different conclusions than prior recommendations - the relationship between latency and recall probability, coupled with the large time costs of being incorrect, meant that practicing paired-associates at an OET of .40 was highly likely to be less efficient than an OET of .90. Our experimental data supported our arguments - practicing paired associates near an optimal high OET led to significantly higher memory retention than alternative OETs as well as better than a conventional spaced practice condition.

We also want to emphasize that we are not arguing that a high OET is always preferable. In fact, it will always be the case that if there is any spacing effect or advantage to difficulty, as the OET approaches 1, efficiency will decline. A 1 OET can only be efficient if learning is optimal with perfect correctness (i.e., if Skinner was correct). A higher OET will be best when gains from successes are equal to or greater than failures while also being faster. This is true for many learning materials (especially facts), but not necessarily all of them. Situations in which failures provide more learning gains or are notably faster will result in lower recommended OETs.

The critical takeaway of the present work is that if an accurate computational model is used, adaptively scheduling practice with model predictions according to an *optimal efficiency threshold* (OET) will confer superior memory relative to conventional practice schedules. The exact value of that OET will vary across learning materials and types of

feedback, but the general takeaway is that the efficiency of a practice event must be considered except for cases where time costs of success and failure are exactly equal. This issue is discussed infrequently, but our simulation and results demonstrate how important this consideration can be for scheduling practice. The benefit of our approach relative to a conventional spaced practice condition was meaningful (cohen's $d = .64$) and comparable in magnitude to that of the spacing effect itself [51] (cohen's $d = .42$).

We also provided evidence that variable difficulty of items and skills of students means that applying one schedule to all items (a common method in prior research) is likely to be inherently suboptimal. Finally, shifting the focus away from fixed schedules and trial durations and toward an adaptive framework offers new exciting avenues of research, such as modifying the component parts of the adaptive framework (models for predicting correctness and RT, estimating OETs) to accommodate learning more complex materials that contain inter-item relations.

Methods

Simulation model fitting

For the simulations we fit a model intended to predict probability of recall (see equation 3), inspired by PPE [35] and Base4 [50]. The model was fit to a dataset in which participants practiced Japanese-English word pairs with several spacing intervals and retention intervals (see supplemental materials for details). The model was fit with intercepts for participants and items². There are many candidate models that would

² See supplemental materials for details, including parameter values.

probably account for the general patterns found in the present data. However, we want to emphasize that our motivation is not to find the best computational model of memory, but rather to investigate how best to *use* such models. We believe the general conclusions of the results below would extend to other computational memory models. A model for response time was also fit. The response time model allowed the total number of trials within a fixed duration to vary according to correctness, given that correct answers would be faster than incorrect answers³.

$$y = \beta_1 a^{-d_s} N_s s^c + \beta_2 a^{-d_f} N_f s^c$$

Equation 3. The general form of the model to track memory over time. There are two parts, meant to allow differential contributions of successes and failures. Parameter a represents the time since first practice, with d representing a decay parameter (separate for successes and failures). N represents counts of prior practice (separate for successes and failures). Parameter S is average spacing in between practice attempts, with curvature parameter c .

The response time model $ce^{-a}+fc$ had two free parameters, c and fc that were estimated via maximum likelihood. Parameter a was the log odds prediction from the correctness model described above [23]. In other words, the response time model sought to translate the correctness model predictions into response time. This model was solely intended to track response time for correct answers (response time for incorrect answers is typically estimated as a constant [52]).

Simulation Methods

Each run of a simulation represented 2 sessions of a single student studying 30 word pairs for 22 minutes, followed by a test 3 days later. The total duration and number of items was chosen to match the duration and practice characteristics of typical

³ We also ran simulations in which response time for correct answers was set to the median RT found for correct answers in the initial dataset (faster than failures), rather determined by the latency model. The findings were highly similar, suggesting a latency model may not be required in some cases. More details are in the supplemental materials.

paradigms in which participants study approximately 30 items, practicing them 4 times each, with 7 second trial durations for initial study trials, 11 seconds for subsequent tests with feedback (7+4 seconds), 0.5 second ISI, for a total of 1320 seconds (22 minutes). We believe simulating a scenario that is common to the literature on this topic makes it easier for researchers to compare our findings and conclusions to prior literature. 200 students were simulated for each condition. Item and student intercepts were sampled from their respective distributions estimated from fitting the model on the experimental dataset. As a concrete example, to compute the probability of success on a given trial for each item, the computational model would predict recall probability with the (simulated) history of practice for each item for the student as well as the intercept for that particular item and student. The correctness of a trial was determined probabilistically using the prediction from the model (e.g., if the model predicted 68% probability of recall, the trial was correct with 68% probability).

Simulated students that were in “optimal” conditions practiced whichever item was closest (but less than) the target OET for their condition (e.g., 90%). For instance, if an item was practiced at 89%, its post-practice estimated probability would go beyond the optimal value (90%), and it could not be practiced again until it decayed to below 90%. If all items had a predicted recall probability greater than the OET, the item with the closest probability was chosen. Figure 9 shows a time course of item probabilities for a student practicing in the 90% optimal condition, with one item highlighted to show the consequences of scheduling practice according to the computational model and the .90 OET.

Figure 9. A plot of the estimated probability of all items during optimized practice for one simulated student (at OET = .90). The blue line indicates the probability of one item across the sessions. Squares indicate when that item was chosen to be practiced. Gray lines represent other items (squares indicating when other items were practiced were omitted to avoid clutter). The black dashed line denotes the OET.

Optimal Probability (OET) Conditions

These conditions differed according to their target probability threshold at which to practice (or the decision rule). There were 22 optimal probability conditions depicted here, in intervals of .05 from .2 to .8, and intervals of .02 from .8 onward to .98. We simulated across the range to illustrate the efficiency curve, the particular intervals were simply intended to help the reader gauge the general pattern. For example, a student in the .80 condition would practice whichever item was closest to (but less than) .8 probability. Items above the optimal value were not practiced until they decayed below the threshold unless all items were above the threshold, in which case whichever item was closest to .8 would be practiced.

Conventional Conditions

Given that our focus is on practice of word-pairs or other similarly independent facts, we simulated schedules common to that literature. Schedules could be simulated either with a fixed trial duration (common to the spacing literature) or a variable trial duration in which trial duration was determined by the latency model. One adaptive heuristic schedule was included (Drop-1), in which an item was dropped from practice after being recalled (if all items were answered correctly before the time elapsed, the process restarted). Thus, in fixed trial duration conditions the number of trials was also fixed, while in variable duration conditions the number of trials could vary. If a student was successful more often, and thus faster, they could complete more trials within the maximum allotted time.

The critical comparison was among the conventional variable duration spacing conditions, Drop-1, and the OET conditions, because in these three types of conditions successes are faster and may lead to additional practice attempts. The question is whether employing an OET (e.g., the rule “practice whichever item is nearest but below .90 recall probability”) improves the efficiency of practice above and beyond conventional spacing schedules that do not use a model.

Simulating Latency

Although computational models do not necessarily differentiate between successes and failures (counts of all attempts are frequently used instead), outcomes still have important consequences for trial duration. Most importantly, failure trials may require

feedback, whereas successful trials may not. Given that failures (with feedback) and successes (with or without feedback) can provide the same memorial benefit [53], successful trials can be more efficient per-second. In simulated “optimal” conditions, trial duration was decided based on success and failure, and thus also influenced the total number of completed trials for a student within the allotted time.

The duration for the first trial for each simulated student was set to 7 seconds, to represent an initial study trial. Subsequently, if the condition allowed the simulated student to proceed at their own pace, the duration of correct trials were estimated with the latency model described above plus 500ms to simulate the correctness feedback. Incorrect trial durations were set to be the median trial duration for incorrect trials found in the experimental data, plus four seconds feedback (8.98 seconds total). In simulated conditions with fixed trial durations (common in many experimental designs), the duration was fixed at 11 seconds (7 seconds to recall, plus 4 seconds feedback). There was a 1 second inter-stimulus-interval (ISI) between trials in all conditions. The values chosen for trial duration and ISI are typical for spacing experiments [45] .

Experiment Methods

Materials

Materials were Japanese-English word pairs used in the parameterization dataset (see supplemental materials). Each participant was assigned a randomly selected subset of 30 word pairs to practice.

Participants

322 Participants were recruited via Amazon Mechanical Turk (Mturk) with the help of TurkPrime [54]. Data from 25 participants were excluded from analysis for reporting that they had at least moderate ability to read the Japanese language. 6 participants were excluded for showing no evidence of completing the task (e.g., timing out on all trials). Participants were excluded at similar rates across the conditions. We aimed for approximately 50 participants for each condition. The one exception was OET98, in which we aimed to collect fewer participants (20) because simulations indicated that participants would overpractice a small set of the items and thus there would be very little variance. In other words, we were very confident the condition would be suboptimal, and we collected data in it solely to concretely demonstrate inefficiency of extremely high OETs. Ultimately, the Ns per condition were 59, 53, 57, 50, 56, and 16 for S15 and OETs 40, 70, 86, 94, and 98, respectively. The number of participants was chosen based on effect sizes observed in past research on adaptive practice scheduling [18, 33, 55]. These prior findings indicated the potential for a medium to large effect size (Cohen's d of .56 to more than 1). We powered our experiment to detect an effect size of $d = .55$. Using the BUCSS statistical package [56], which corrects for publication bias and sample uncertainty, we estimated we had approximately 89% probability to detect a significant difference among conditions.

Method

Participant data collection was completed using the MoFacTS system, which is software designed to allow practice scheduling according to a model as well as more traditional approaches [57]. All participants completed two sessions. In the first session, after providing informed consent, participants studied 30 word pairs for 22 minutes. Twenty-

two minutes was chosen to approximately match previous research on learning word pairs that typically involves studying about 30 word pairs a few times each [13, 28]. Twenty-two minutes also ensured 4 exposures to each item during the first session given maximum trial durations. For all participants in all conditions, the first trial for each item was a 7 second study trial in which the complete Japanese-English word pair was presented. Participants could not proceed before 7 seconds had elapsed. On subsequent trials, participants were provided with cue (e.g., “awa - _____”) and had up to 7 seconds to begin typing in an answer. If a participant began typing before 7 seconds elapsed, the timer reset to seven seconds so they could complete typing. This accommodation rarely extended the trial duration beyond 7 seconds (~3% of trials). On test trials, feedback was presented once the participant typed in an answer or once 7 seconds elapsed. If a participant typed the correct answer, feedback was a 500ms “correct!” message. If they answered incorrectly, the correct Japanese-English word pair was presented for four seconds. After either type of feedback, there was a 500ms inter-trial delay and then the next trial began. After completing the first session, all participants completed a brief survey asking them for demographic information as well as their knowledge of Japanese vocabulary. Participants were informed that their answers would not influence their payment. Data from participants that reported knowledge of Japanese vocabulary prior to the experiment were discarded. The conditions differed in terms of how practice was scheduled during the first session, described below.

Three days later an email was automatically sent to all participants inviting them to return for a final test. Participants were allowed up to 48 hours from when they received the email to complete the final test. During the final test, participants were tested for their memory of all Japanese-English word pairs selected for them during the first session. The order in which word pairs were tested was randomized, and feedback was presented in the same way as in the first session (i.e., a brief “correct” message if correct, and corrective feedback if incorrect).

Model

The practice scheduling model was highly similar to that fit to the parameterization dataset and used to schedule practice in the simulation. The only difference was that due to not having participant intercepts (due to having not collected data from these new participants yet), the participant intercept was excluded. The previously estimated parameters were fixed (e.g., spacing and decay parameters), and the model was refit with only item intercepts (because we do have previous data for all word pairs from the previous experiment).

Adaptive Scheduling Method

If participants were assigned to an adaptive scheduling condition, they practiced according to one of five difficulty thresholds - .40, .70, .86, .94, or .98. This approach proceeded as follows: for each trial, the probability that each item could be recalled was estimated by the model (the difficulty). This estimate was determined by the model using the prior history of practice for that item for that participant (including counts of successes, failure, average spacing, and elapsed time since those attempts), or if there

was no prior history just the item intercept was used. Whichever item was closest to (but less than) the difficulty threshold was chosen to be practiced. If all items were above the threshold, the item closest to the threshold was chosen. After each trial, the model re-estimated the recall probabilities for all items, and the process continued until the allotted time elapsed.

Conventional Scheduling Method

If participants were assigned to the conventional scheduling condition, their practice was separated into two consecutive blocks, with 15 items practiced in each block. In each block, items were repeated approximately every 15 trials (maximally spaced within block). The exact ordering was slightly jittered to reduce the probability that participants could predict the presentation of upcoming items. We chose this spacing interval due to evidence in initial pilot studies that maximal spacing within one block (i.e., repeating every 30 trials) was exceedingly difficult. Maximal spacing has been shown to be suboptimal in prior research as well [2]. The implementation of initial study trials, test trials, and feedback (when incorrect) was the same as in the difficulty threshold conditions.

Data Availability

The experimental data are available at <https://osf.io/d9rms/> and <https://datashop.memphis.edu/>.

Code Availability

Simulation script (with embedded instructions) is available at <https://osf.io/d9rms/>.

Contributions

Both authors contributed to the research design and interpretation of results.
L.G.E. developed the simulation and analysis code.

Competing Interests

The authors declare no competing interests.

References

1. Ebbinghaus, H., *Memory: A contribution to experimental psychology*. 1913, New York: Teachers College, Columbia University.
2. Cepeda, N.J., et al., *Distributed Practice in Verbal Recall Tasks: A Review and Quantitative Synthesis*. Psychological Bulletin, 2006. **132**(3): p. 354-380.
3. Rowland, C.A., *The effect of testing versus restudy on retention: A meta-analytic review of the testing effect*. Psychological Bulletin, 2014. **140**(6): p. 1432-1463.
4. Pan, S.C. and T.C. Rickard, *Transfer of test-enhanced learning: Meta-analytic review and synthesis*. Psychological Bulletin, 2018. **144**(7): p. 710-756.
5. Kang, S.H.K., *Spaced Repetition Promotes Efficient and Effective Learning*, in *Policy Insights from the Behavioral and Brain Sciences*. 2016. p. 12-19.
6. Young, J.L., *Reinforcement-test intervals in paired-associate learning*. Journal of Mathematical Psychology, 1971. **8**(1): p. 58-81.
7. Glenberg, A.M., *Influences of retrieval processes on the spacing effect in free recall*. Journal of Experimental Psychology: Human Learning & Memory, 1977. **3**(3): p. 282-294.
8. Cull, W.L., *Untangling the benefits of multiple study opportunities and repeated testing for cued recall*. Applied Cognitive Psychology, 2000. **14**(3): p. 215-235.
9. Bjork, R.A., *Institutional impediments to effective training*. Learning, remembering, believing: Enhancing human performance, 1994: p. 295-306.
10. Hintzman, D.L., *Judgment of frequency versus recognition confidence: repetition and recursive reminding*. Mem Cognit, 2004. **32**(2): p. 336-50.
11. Carpenter, S.K., et al., *The effects of tests on learning and forgetting*. Memory & Cognition, 2008. **36**(2): p. 438-448.
12. Landauer, T.K. and R.A. Bjork, *Optimum rehearsal patterns and name learning*, in *Practical Aspects of Memory*, M.M. Gruneberg, P.E. Morris, and R.N. Sykes, Editors. 1978, Academic Press: New York. p. 625-632.
13. Karpicke, J.D. and H.L. Roediger, *Expanding retrieval practice promotes short-term retention, but equally spaced retrieval enhances long-term retention*. Journal of Experimental Psychology-Learning Memory and Cognition, 2007. **33**(4): p. 704-719.
14. Kang, S.K., et al., *Retrieval practice over the long term: Should spacing be expanding or equal-interval?* Psychonomic Bulletin & Review, 2014. **21**(6): p. 1544-1550.
15. Pyc, M.A. and K.A. Rawson, *Examining the efficiency of schedules of distributed retrieval practice*. Memory & Cognition, 2007. **35**(8): p. 1917-1927.
16. Metcalfe, J. and N. Kornell, *The dynamics of learning and allocation of study time to a region of proximal learning*. Journal of Experimental Psychology: General, 2003. **132**(4): p. 530-542.
17. Tullis, J.G., J.L. Fiechter, and A.S. Benjamin, *The efficacy of learners' testing choices*. J Exp Psychol Learn Mem Cogn, 2018. **44**(4): p. 540-552.

18. Pavlik, P.I. and J.R. Anderson, *Using a model to compute the optimal schedule of practice*. J Exp Psychol Appl, 2008. **14**(2): p. 101-17.
19. Kahana, M.J. and G. Loftus, *Response time versus accuracy in human memory.*, in *The nature of cognition*, R.J. Sternberg, Editor. 1999, MIT Press. p. 323-384.
20. Anderson, *Interference: The relationship between response latency and response accuracy*. Journal of Experimental Psychology: Human Learning & Memory, 1981. **7**(5): p. 326-343.
21. Madigan, S., J. Neuse, and U. Roeber, *Retrieval latency and "at-risk" memories*. Memory & Cognition, 2000. **28**(4): p. 523-528.
22. Anderson, J.R., J.M. Fincham, and S. Douglass, *Practice and retention: A unifying analysis*. Journal of Experimental Psychology: Learning, Memory, & Cognition, 1999. **25**(5): p. 1120-1136.
23. Anderson, J.R., *ACT - A simple theory of complex cognition*. American Psychologist, 1996. **51**(4): p. 355-365.
24. Newell, A. and P.S. Rosenbloom, *Mechanisms of skill acquisition and the law of practice*. Cognitive skills and their acquisition, 1981. **1**: p. 1-55.
25. Rawson, K.A. and J. Dunlosky, *Optimizing schedules of retrieval practice for durable and efficient learning: how much is enough?* J Exp Psychol Gen, 2011. **140**(3): p. 283-302.
26. Putnam, A.L., V.W. Sungkhasettee, and H.L. Roediger, 3rd, *Optimizing Learning in College: Tips From Cognitive Psychology*. Perspect Psychol Sci, 2016. **11**(5): p. 652-660.
27. Storm, B.C., R.A. Bjork, and J.C. Storm, *Optimizing retrieval as a learning event: when and why expanding retrieval practice enhances long-term retention*. Mem Cognit, 2010. **38**(2): p. 244-53.
28. Yan, V.X., L.G. Eglinton, and M.A. Garcia, *Learning Better, Learning More: The Benefits of Expanded Retrieval Practice* Journal for Applied Research in Memory and Cognition, 2020.
29. Roediger, H.L., et al., *Benefits of testing memory: Best practices and boundary conditions.*, in *Current issues in applied memory research*, G.M. Davies, Editor. 2009, Psychology Press. p. 27-63.
30. Pashler, H., et al., *When does feedback facilitate learning of words?* Journal of Experimental Psychology: Learning, Memory, and Cognition, 2005. **31**(1): p. 3-8.
31. Atkinson, R.C., *Ingredients for a theory of instruction*. American Psychologist, 1972. **27**(10): p. 921-931.
32. Smallwood, R.D., *A decision structure for teaching machines*. 1962, Cambridge: MIT Press.
33. Atkinson, R.C., *Optimizing the learning of a second-language vocabulary*. Journal of Experimental Psychology, 1972. **96**(1): p. 124-129.
34. Mozer, M.C., et al., *Predicting the optimal spacing of study: A multiscale context model of memory*, in *Advances in Neural Information Processing Systems*, Y. Bengio, et al., Editors. 2009, NIPS Foundation: La Jolla, CA. p. 1321-1329.
35. Walsh, M.M., et al., *Mechanisms underlying the spacing effect in learning: A comparison of three computational models*. Journal of Experimental Psychology: General, 2018. **147**(9): p. 1325-1348.
36. Benjamin, A.S. and J. Tullis, *What makes distributed practice effective?* Cogn Psychol, 2010. **61**(3): p. 228-47.
37. Bahrack, H.P., *Maintenance of knowledge: Questions about memory we forgot to ask*. Journal of Experimental Psychology: General, 1979. **108**(3): p. 296-308.
38. Pavlik, P.I., Jr. and J.R. Anderson, *Practice and forgetting effects on vocabulary memory: an activation-based model of the spacing effect*. Cogn Sci, 2005. **29**(4): p. 559-86.

39. Lindsey, R.V., et al., *Improving Students' Long-Term Knowledge Retention Through Personalized Review*. Psychological Science, 2014. **25**(3): p. 639-647.
40. Khajah, M.M., R.V. Lindsey, and M.C. Mozer, *Maximizing Students' Retention via Spaced Review: Practical Guidance From Computational Models of Memory*. Topics in Cognitive Science, 2014. **6**(1): p. 157-169.
41. Bjork, R.A., *Memory and metamemory considerations in the training of human beings.*, in *Metacognition: Knowing about knowing*, J. Metcalfe, Editor. 1994, MIT Press: Cambridge, MA. p. 185-205.
42. Rickard, T.C., *Bending the power law: A CMPL theory of strategy shifts and the automatization of cognitive skills*. Journal of Experimental Psychology: General, 1997. **126**(3): p. 288-311.
43. Hebb, D.O., *Drives and the C.N.S. (conceptual nervous system)*. Psychol Rev, 1955. **62**(4): p. 243-54.
44. Tabibian, B., et al., *Enhancing human learning via spaced repetition optimization*. Proceedings of the National Academy of Sciences of the United States of America, 2019. **116**(10): p. 3988-3993.
45. Walsh, M.M., et al., *Evaluating the Theoretic Adequacy and Applied Potential of Computational Models of the Spacing Effect*. Cogn Sci, 2018. **42 Suppl 3**: p. 644-691.
46. Metzler-Baddeley, C. and R.J. Baddeley, *Does adaptive training work?* Applied Cognitive Psychology, 2009. **23**(2): p. 254-266.
47. McFadden, D., *Quantitative methods for analyzing travel behavior of individuals: Some recent developments*, in *Behavioural travel modelling*. 1977, Croom Helm. p. 279-218.
48. Sense, F., et al., *An Individual's Rate of Forgetting Is Stable Over Time but Differs Across Materials*. Topics in Cognitive Science, 2016. **8**(1): p. 305-321.
49. Fiechter, J.L. and A.S. Benjamin, *Techniques for scaffolding retrieval practice: The costs and benefits of adaptive versus diminishing cues*. Psychon Bull Rev, 2019. **26**(5): p. 1666-1674.
50. Pavlik Jr, P.I., L.G. Eglington, and L.M. Williams, *Generalized Knowledge Tracing: A Constrained Framework for Learner Modeling*. 2020 arXiv preprint arXiv:2005.00869, arXiv.
51. Donovan, J.J. and D.J. Radosevich, *A meta-analytic review of the distribution of practice effect: Now you see it, now you don't*. Journal of Applied Psychology, 1999. **84**(5): p. 795-805.
52. Pavlik Jr., P.I. and J.R. Anderson, *Using a model to compute the optimal schedule of practice*. Journal of Experimental Psychology: Applied, 2008. **14**(2): p. 101-117.
53. Kornell, N. and K.E. Vaughn, *How Retrieval Attempts Affect Learning*. 2016. p. 183-215.
54. Litman, L., J. Robinson, and T. Abberbock, *TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences*. Behav Res Methods, 2017. **49**(2): p. 433-442.
55. Mettler, E. and P.J. Kellman, *Adaptive response-time-based category sequencing in perceptual learning*. Vision Research, 2014(0).
56. Anderson, S.F., K. Kelley, and S.E. Maxwell, *Sample-Size Planning for More Accurate Statistical Power: A Method Adjusting Sample Effect Sizes for Publication Bias and Uncertainty*. Psychol Sci, 2017. **28**(11): p. 1547-1562.
57. Pavlik Jr., P.I., C. Kelly, and J.K. Maass, *Using the mobile fact and concept training system (MoFaCTS)*, in *Proceedings of the 13th International Conference on Intelligent Tutoring Systems*, A. Micarelli and J. Stamper, Editors. 2016, Springer: Switzerland. p. 247-253.