



Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Phases of learning: How skill acquisition impacts cognitive processing



Cognitive Psychology

19

Caitlin Tenison*, Jon M. Fincham, John R. Anderson

Department of Psychology, Carnegie Mellon University, United States

ARTICLE INFO

Article history: Accepted 10 March 2016 Available online 25 March 2016

Keywords: Skill acquisition Cognitive modeling ACT-R fMRI

ABSTRACT

This fMRI study examines the changes in participants' information processing as they repeatedly solve the same mathematical problem. We show that the majority of practice-related speedup is produced by discrete changes in cognitive processing. Because the points at which these changes take place vary from problem to problem, and the underlying information processing steps vary in duration, the existence of such discrete changes can be hard to detect. Using two converging approaches, we establish the existence of three learning phases. When solving a problem in one of these learning phases, participants can go through three cognitive stages: Encoding, Solving, and Responding. Each cognitive stage is associated with a unique brain signature. Using a bottom-up approach combining multi-voxel pattern analysis and hidden semi-Markov modeling, we identify the duration of that stage on any particular trial from participants brain activation patterns. For our top-down approach we developed an ACT-R model of these cognitive stages and simulated how they change over the course of learning. The Solving stage of the first learning phase is long and involves a sequence of arithmetic computations. Participants transition to the second learning phase when they can retrieve the answer, thereby drastically reducing the duration of the Solving stage. With continued practice, participants then transition to the third learning phase when they recognize the problem as a single unit and produce the answer as an automatic response. The duration of this third learning phase is dominated by the Responding stage.

© 2016 Elsevier Inc. All rights reserved.

* Corresponding author at: 5000 Forbes Ave., Pittsburgh, PA 15213, United States. *E-mail address:* ctenison@andrew.cmu.edu (C. Tenison).

http://dx.doi.org/10.1016/j.cogpsych.2016.03.001 0010-0285/© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Across domains, practice is acknowledged to have transformative effects on performance. Several models of skill acquisition propose different explanations for how practice reduces the duration and increases the accuracy of task performance. These models differ in whether they attribute speedup either to discrete changes in the cognitive processes executed to solve a problem or to a greater efficiency of the same processes. In this paper, we address both explanations within the context of modeling the speed up in mathematical problem solving. In particular, we combine cognitive modeling with new methods of analyzing fMRI data to understand the detailed changes that take place as participants transition from the first time they solve a novel problem, to the point at which they automatically recognize a problem's solution.

1.1. Models of skill acquisition and practice related speedup

As people practice solving a problem, the time it takes to solve that problem decreases. Previous research has focused on understanding the nature of this speed up. In their classic paper, Newell and Rosenbloom (1981) observed that performance tends to speed up as a power function of the amount practice, highlighting what has been called the 'Power Law of Practice'. In their paper describing these effects, performance improvements were explained as chunking of cognitive processes into fewer processes (Newell & Rosenbloom, 1981). Subsequent research has refined this characterization of practice by investigating whether the speedup is really best fit as a power function (e.g. Heathcote, Brown, & Mewhort, 2000), and by examining to what degree the speedup might reflect changes in the strategies used for solving the problems (e.g. Delaney, Reder, Staszewski, & Ritter, 1998).

Unlike Newell and Rosenbloom's work, the Race model, which is part of the Instance Theory (Compton & Logan, 1991; Logan, 1988, 2002), described the learning mechanism for practicerelated speedup as involving a quick shift from computation to retrieval followed by a power-like speed up in the retrieval. According to the Race model, each time a problem is practiced it becomes encoded in memory; then, when a participant sees the problem again, each of the previously encoded instances independently races to generate the response and the fastest process 'wins the race'. As the number of instances in which a participant retrieves the answer increases, the speed of the winning retrieval will increase as well. This model not only predicts a power-law speedup with practice, but also a decrease in the variability in latency with practice.

In contrast to the Race model, a number of studies have suggested the computation process may have its own speedup before the shift to retrieval. Delaney et al. (1998) argued that both computation and retrieval are strategies that have their own power-law speedup. Rickard's Component Power Law (CMPL) model (Bajic, Kwak, & Rickard, 2011; Bajic & Rickard, 2009; Rickard, 1997) focuses on two sources of speedup—the discrete switch from computation to retrieval and the speed up within the computation and retrieval phases. Initially Rickard focused on speed up in retrieval but he has also suggested that faster associative retrieval may be responsible for speed up in the computations used to solve a problem.

Fitts and Posner (1967) were the first to propose a three-phase model of skill acquisition that consists of a cognitive phase, an associative phase, and an autonomous phase. Anderson (1982, 1987) later related this to an early version of a cognitive architecture, the ACT theory. He suggested that the transition from the cognitive to the associative phase reflected a transition from computation of the answers to declarative retrieval of learned answers, and that the transition to the autonomous phase was produced when retrieval was replaced by a production rule that directly produced the response to the stimulus.

The task in this paper is modeled within modern ACT-R theory (Anderson, 2007; Anderson et al., 2004). ACT-R models specify the full time course of processing that occurs while performing a task, from perceptual encoding through response generation, differentially leveraging specific functional modules over the course of a task. A significant advantage of modeling at this level of detail is that it allows us to relate a model of skill acquisition to changes in brain activation by identifying the major

cognitive systems used at different stages of a task. ACT-R models for a number of tasks have addressed the transition from computation to retrieval and the impact of practice on performance. Anderson and Lebiere (1998) described a model of an alpha-arithmetic experiment by Zbrodoff (1995), which has become a standard in the ACT-R tutorial (Bothell et al., 2009). The major speedup is produced by a discrete transition from computation to retrieval although there is some further speed up in retrieval of the answer with practice. That model did not learn production rules for directly mapping stimuli to responses (and so did not reach the autonomous phase of skill acquisition); later models, however, added this functionality. Such item-specific production rules have played a major role in a number of ACT-R models where specific items have been repeated a large number of times (Anderson, Taatgen, & Byrne, 2005; Taatgen & Anderson, 2002; Taatgen & Lee, 2003), and will do so here as well.

1.2. Models of practice and the brain

The models of skill acquisition previously discussed emerged from research in which cognitive processing was inferred from reaction time. A separate, but related field of research on the impact of practice on brain function has largely neglected theory concerning how cognitive processing changes with practice. Many studies on the impact of practice occur over long periods of time, such as hours or days (e.g. Lee et al., 2012; Tang & Posner, 2014), using contrasts between trained and untrained tasks to observe gross structural changes (Lövdén, Wenger, Mårtensson, Lindenberger, & Bäckman, 2013), as well as changes in functional activation (Patel, Spreng, & Turner, 2013). In a review of this literature, Kelly and Garavan (2005) acknowledge the importance of considering the impact of practice on strategic changes in how the task is performed. Although Kelly and Garavan wrote their review a decade ago, this sentiment continues to be echoed by others (Reber, 2013; Tang & Posner, 2014). In this section, we examine two different cognitive models of skill acquisition that have emerged to explain the impact of practice on neural activation.

In their 'Scaffolding framework', Petersen, Van Mier, Fiez, and Raichle (1998) study the impact of practice on task-specific demands. When participants first encounter a novel task, a set of brain regions categorized as scaffolding regions are used to perform the task. These regions vary depending on the task demands. With practice, associations become stored in memory and a different set of storage brain regions are involved in the task. While this framework shares many similarities with the Instance theory (Logan, 1988, 2002) and the CMPL (Rickard, 1997, 2004), the authors point out the difficulty in using neural activations to prove models such as their own. The scaffolding framework is one of a family of frameworks researchers have pieced together to interpret changes in neural activation that occur during skill acquisition (e.g. Kelly & Garavan, 2005; Patel et al., 2013). While these studies often mention the previously described models of skill acquisition (e.g. Anderson, 1982; Fitts & Posner, 1967; Logan, 1988; Rickard, 1997), these are most often used to describe practice effects in broad strokes rather than to test predictions about learning related change. In the current paper, we propose a method that addresses some of the limitations raised by Petersen et al. (1998) when using neuro-science to understand skill acquisition.

With the dual processing theory, Schneider and Chein (2003) makes the distinction between control and automatic processes. Unlike the scaffolding framework, the dual processing theory can be simulated using the computational model, CAP2, developed by Schneider and Chein. When first introduced to a task, the participant performs the task using controlled processes that are limited to the sequential execution of cognitive processes. This initial process is modeled with CAP2, a system that is much like an abbreviated version of ACT-R. Eventually, with practice, controlled processes switch to automatic processes that simulate the reflex like automization phase described by Fitts and Posner (1967). These automatic processes are modeled in the CAP2 system with a connectionist network and allowed for parallel processing. Unlike the Instance theory, CMPL, or three-phase model, which each all distinguish computation from retrieval processes, both processes occur in the control process arm of CAP2. The distinction between computation and retrieval is not emphasized by Schneider and Chein. This theory is used in several studies to interpret neural changes from pre- to post-training (Albouy et al., 2012; Chein & Schneider, 2005; Luu, Tucker, & Stripling, 2007). The dual process theory, however, is unsuitable for addressing the questions surrounding skill acquisition that we address in this paper as it does not address the cognitive changes that occur prior to the shift to automization (Chein & Schneider, 2005).

Much of the prior work exploring the impact of practice on neural activation forgoes an exploration of the continuous changes associated with practice in favor of pre- and post-training contrasts (for reviews: Kelly & Garavan, 2005; Patel et al., 2013). While models we discussed here exist to explain the general distinction between novel and well-practiced tasks, there is a gap in this literature in understanding how this shift occurs. Although these changes may be obscured in traditional analysis (Kelly & Garavan, 2005; Patel et al., 2013; Poldrack, 2000) there is evidence that, even over the course of an initial pre-training scan, changes in neural activation attributable to learning the stimuli can be found (Petersen et al., 1998). Our approach offers an alternative for both identifying when shifts in cognitive processing occur and also for building hypotheses concerning the expected impact cognitive changes have on neural activation.

1.3. Pyramid problems and the use of hidden Markov models

This paper merges two efforts that have used hidden Markov models to understand problem solving and skill acquisition by parsing sequences of episodes into distinct states. The two efforts are focused on different time scales. Anderson and Fincham (2014a, 2014b) were concerned with the cognitive stages that occur within a particular problem-solving trial. Tenison and Anderson (2015) on the other hand, focused on between-trial changes: the learning phases that occur over an entire experimental session. The current effort combines these two approaches, providing a characterization of changes that span three orders of magnitude, from seconds to an hour.

Anderson and Fincham (2014a, 2014b) used multi-voxel pattern analysis (MVPA) and hidden semi-Markov models (HSMMs) to develop and test ACT-R models. Using this approach, they parsed fMRI data collected during mathematical problem-solving episodes into multiple stages defined by distinct brain signatures, and tested model predictions about the length of these stages. They were, however, only concerned with problem solving on the first encounter with a problem and not on subsequent repetitions of that problem. Tenison and Anderson (2015), using a similar experimental task, focused on how problem solving changed across multiple practice opportunities. Using the latency data, they found evidence for three phases of learning, which they tentatively associated with the cognitive, associative, and autonomous phases proposed by Fitts and Posner (1967). Furthermore, fMRI data collected during this experiment was qualitatively consistent with this three-phase characterization. Here we will use the Tenison and Anderson (2015) learning paradigm, but apply the HSMM–MVPA methods used by Anderson and Fincham (2014a, 2014b).

In both the experiments of Anderson and Fincham (2014a, 2014b) and Tenison and Anderson (2015), participants solve *pyramid problems*, which are presented with a dollar symbol as the operator—e.g., 8\$3 = X. Pyramid problems involve a *base* ("8" in this example) that is the first term in an additive sequence and a *height* ("3" in this example) that determines the number of terms to add. Each term in the sequence is one less than the previous—so 8\$3 = 8 + 7 + 6 = 21. Anderson and Fincham (2014b) compared problems with positive single digits for base and height (*regular problems*) to problems that required participants to extend their knowledge (*exception problems*). Exception problems contained features such as problems with negative or variable bases (e.g. -9\$4 = X, X\$4 = X). Using the HSMM–MVPA technique, they identified and modeled 4 cognitive stages in the solution of exception problems—an Encoding stage, a Planning stage, a Solving stage, and a Responding stage. When applying this approach to regular problems, they identified only 3 stages, discovering that regular problems were solved without a Planning stage.

The current study focuses on regular pyramid problems, and the impact of practice on the cognitive stages participants go through when solving these problems. Fig. 1 shows a swimlane representation of the activity of Anderson and Fincham's (2014b) ACT-R model for the solution of the regular problem 7\$4 = X, which took the model 10.91 s to solve, close to the participant's average time of 11.24 s on this problem. The ACT-R cognitive architecture is organized into modules, each of which processes a different type of information and has its own buffer to hold information. Fig. 1 shows when the Visual, Imaginal, Retrieval, Manual, and Metacognitive modules are active in solving the task. Periods when different subgoals were active are labeled above the swimlane. During the Encoding period the 5



Fig. 1. A swimlane representation of an ACT-R model whose subgoals correspond to the bracketed stages. The boxes in a row represent when a module is engaged.

terms in the expression are separately encoded, requiring heavy engagement of the Visual module and the Imaginal module to build up a mental representation of the problem. The Encoding period is terminated when the Metacognitive module selects the strategy for solving the regular problem. The Solving subgoal involves the iterative addition of 7 + 6 = 13, 13 + 5 = 18, and 18 + 4 = 22. This requires heavy engagement of the Retrieval module to process relevant facts and the Imaginal module to update problem representations as the computations are performed. The Manual module dominates during the Responding subgoal when the digits in "22" are keyed out followed by an enter key. The Manual module is also involved during the Solving stage, where it is used in finger counting to keep track of the number of terms added.

The representation of the problem solving activity in ACT-R acts as a hypothesis that can be tested by exploring each participant's brain activation while problem solving. The module activity can be organized into subgoals that reflect different stages in solving the problem. For instance, in the model of the regular problem 7\$4 = X depicted in Fig. 1, there are three stages: Encoding, Solving, and Responding, Anderson and Fincham developed the HSMM-MVPA approach illustrated in Fig. 2 to find convergence between the predictions of cognitive models such as Fig. 1 and the participant's neuroimaging data. This method simultaneously used MVPA to infer the mean brain activation pattern in a stage and HSMM to infer how a trial was parsed into stages. Each cognitive stage is associated with an expected whole brain pattern of activation referred to as that stage's brain signature. The brain signatures on individual trials were treated as samples from a normal distribution around the mean brain signature for that stage. The durations of a cognitive stage were treated as samples from a Gamma distribution for that stage. One of the important products of the estimation process for each trial is stage occupancy profiles (see Fig. 2). These are the probabilities that specific scans are in specific stages. The duration of a stage within a trial was estimated as the product of the duration of a scan and the sum of the probabilities that each scan on that trial was in that stage. Finally, Anderson and Fincham (2014a) linked the ACT-R model (e.g. Fig. 1) to the HSMM by assuming that the duration of a subgoal in ACT-R corresponded to the mean duration of the corresponding stage.

Tenison and Anderson (2015) did not use fMRI data to address the problem of partitioning individual trials into stages; rather, they used problem-solving latency data to address the issue of partitioning a sequence of trials into learning phases (i.e., computation, retrieval, autonomous). Fig. 3A shows the HMM state structure that Tenison and Anderson used to address this ambiguity. There was a different HMM state for each number of trials an item could spend in a learning phase. This allowed the HMM to track speedup within a learning phase due to practice. On any trial, the participant could transition to a faster HMM state within that learning phase or transition to the first HMM state within the next learning phase. Fig. 3B shows, for one item practiced 36 times, the inferred probability that it is in a particular learning phase, noted by the different line types. Fig. 3C shows the problem solving times on which those inferences were based. As can be seen in Fig. 3B, there is some ambiguity about when Learning phase 1 ends and Learning phase 3. Fig. 3C also shows the expected latency for each trial, assuming that trial was in its most probable Learning phase. As that figure



Fig. 2. Illustration of HSMM–MVPA producing a 4-stage solution with 3 conditions. Input to the HSMM–MVPA are scans organized within trials. Each scan consists of the 20 PCA scores. Parameters are estimated for each stage: 20 PCA means that define the brain signature and distribution parameters for each condition-stage combination. Also calculated are the stage occupancies (probabilities that the participant is in each stage on a particular scan).

illustrates, the majority of the speed up seemed to be between learning phases and not within learning phases.

1.4. The current study

In the current study, as in Tenison and Anderson (2015), participants were given a set of Pyramid problems, which they practiced 36 times over the course of the experiment. The goal of the current study is to extend the HSMM–MVPA approach developed by Anderson and Fincham (2014a) to identify what the cognitive stages were within each of the learning phases identified by Tenison and Anderson. Therefore, rather than Tenison and Anderson's approach of using latency to guide the estimation of an HMM and then using fMRI to examine the qualitative structure of the resulting states, this paper uses fMRI to guide the estimation of an HSMM (Fig. 4), and then latency data and ACT-R modeling to examine the properties of the resulting states. To reliably use fMRI for this more ambitious goal, we used the same task as in Tenison and Anderson but doubled the sample size. Ultimately, by identifying the cognitive stages within each learning phase, our goal is to provide a more nuanced picture of the changes in problem solving that occur with practice.



Fig. 3. The HMM analysis in Tenison and Anderson (2015): (A) The HMM state structure used to keep track of amount of learning within three learning phases. For a particular item practiced by a single individual in (B) we show the probability that a particular item was in each of the learning phases on 36 trials. Phase 1 probability is indicated by the solid line, Phase 2 as the dotted line, and Phase 3 as the dashed line. In (C) we show the observed latencies for the practiced item modeled in (B) and the expected latency of that item if assigned to the majority phase on each trial as suggested by (B).



Fig. 4. The structure of the HSMM–MVPA used in the current study: Horizontally we show how an MVPA analysis similar to the one represented in Fig. 2 is used to identify distinct cognitive stages and an HSMM is used to model the duration of time spent in each cognitive stage. Vertically, we illustrate that differences in estimated duration of each cognitive stage drive the transition to new learning phases.

2. Methods

2.1. Participants

Forty university students (21 women; mean age 22/SD 2.7) participated in the study. The first 20 participants served as a basis for Tenison and Anderson (2015). Participants gave informed written consent and received monetary compensation for their participation. All participants were right-handed and had normal or corrected-to-normal vision. The university ethics board approved the study.

2.2. Materials and procedure

The current study only considers regular pyramid problems as described in the introduction. Regular pyramid problems follow the form of base\$height, where the base indicates the first term in the additive sequence, and the height determines the number of terms to be added together (e.g., 8\$4 = 8 + 7 + 6 + 5). Problem selection criteria were identical to Tenison and Anderson (2015) and included problems with heights of 3, 4, and 5. The bases varied from 4 to 11 with the restriction that the minimum base for a given height was that height plus one.

Prior to entering the scanner, participants were taught how to compute pyramid problems. The investigator explained the role of base and height in constructing iterative additions to solve the problem, but rather than presenting "4\$3 = X" as in Anderson and Fincham (2014a, 2014b), we only presented "4\$3"—which means less visual encoding is required. After each problem was solved, the participant received correctness feedback and instruction on how to solve the problem. This feedback was visually represented as the chain of additions required to calculate the answer (e.g., 4\$3 = 4 + 3 + 2 = 9). None of the problems used in the practice session were used in the scan session. Participants also received practice on using a numeric keypad that they would use in the scanner to input the solutions to the problems.

Over the course of the experiment, each participant practiced solving three problems, one of each height (i.e., 3, 4, and 5) repeated 36 times. Repeated problems were chosen for repetition with nearly equal frequency across the 40 participants. The remaining 18 unique problems were presented twice, once in the first half of the experiment and once in the second half (approximately 30 min passed between first and second exposure). We will refer to the three practice problems repeated 36 times as repeated problems and to mirror this distinction we refer to the remaining eighteen problems as non-repeated problems despite the fact these problems were seen twice. In previous work using this design, we used a concurrent strategy assessment and found that only 3% of these non-repeated problems were reported as retrieved, and we observed no increase in reports of retrieval between first and second half of the experiment (Tenison & Anderson, 2015). These results suggest that the repetition did not impact problem-solving and it is possible that participants did not notice the repetition in the second half of the experiment. We divided the experiment into 6 blocks of 24 problems during which each repeated problem occurred 6 times and there were 6 non-repeated problems. Pyramid problems were presented on the screen following a 2-s fixation period. Once the problem appeared on the screen, the participant was allowed a maximum of 30 s to indicate knowledge of a solution by pressing the return key on the numeric keypad. After pressing 'return', participants had 5 s to input a solution using the keypad and then press the return key. At the end of each problem solving trial, a 1-back task was presented onscreen for a randomly selected time (between 6 and 12 s) to prevent reflection on the previous problem and allow the hemodynamic response of the brain to return to baseline. In the 1-back task, participants were asked to judge if the letter in the center of the screen was the same as the previous letter seen. Further details of the experimental task and protocol are provided in Tenison and Anderson (2015). Next, we focus on the special processing of the fMRI data required for the HSMM–MVPA analysis.

2.3. fMRI data acquisition

Images were acquired using gradient echo-echo planar image acquisition on a Siemens 3T Verio Scanner using a 32 channel RF head coil, with 2 s. repetition time (TR), 30 ms. echo time, 79° flip angle,

and 20 cm. field of view. The experiment acquired 34 axial slices on each TR using a 3.2 mm thick, 64×64 matrix. This produces voxels that are 3.2 mm high and 3.125×3.125 mm². The anterior commissure-posterior commissure line was on the 11th slice from the bottom scan slice. Acquired images were pre-processed and analyzed using AFNI (Cox, 1996). Functional images were motion-corrected using 6-parameter 3D registration. All images were then slice-time centered at 1 s and co-registered to a common reference structural MRI by means of a 12-parameter 3D registration and smoothed with an 6 mm full-width-half-maximum 3D Gaussian filter to accommodate individual differences in anatomy. As a step of dimension reduction and to accommodate variations in anatomy over participants that may not be dealt with in co-registration, we aggregated the original voxels (3.2 mm high and 3.125×3.125 mm²) in each slice into larger 2 × 2 voxel regions within a slice. There are 12,481 such regions. Some of these regions show an excess of extreme values for some participants, probably reflecting differences in anatomy; these were regions mainly on the top and bottom slices, as well as some regions around the perimeter of the brain. Eliminating these regions resulted in keeping 9973 regions.

The blood-oxygen-level dependent (BOLD) response is calculated as the percent change from a linear baseline defined from first scan (beginning of fixation before problem onset) to last scan (beginning of fixation before next problem). The BOLD response is delayed from the brain activity that initiated it and builds up over many seconds. The BOLD response is deconvolved with a hemodynamic response function (HRF) to produce an estimate of the underlying activity signal at a moment in time. The HRF in SPM is the difference of Gammas: $Gamma(6, 1) - \frac{Gamma(16, 1)}{6}$ (Friston, 2011). A Wiener filter (Glover, 1999) with a noise parameter of .1 was used to deconvolve the BOLD response into an inferred activity signal on each scan.

The output of this process is an estimation of activity in 9973 regions during every 2-s scan. However, the 9973 regions are not independent because of high correlation among regions. To decorrelate the data and extract the meaningful independent sources of information, we performed a spatial principal component analysis (PCA) of the voxel activity where each voxel is treated as a variable that varies over scans, trials, and participants. As in Anderson and Fincham (2014a, 2014b), we focused on the first 20 PCA components, which account for 47% of the total variance in the data, which we assume captures mostly systematic effects in the data. These 20 component scores are approximately 20 independent normally distributed variables. We standardize them to have mean 0 and variance 1.

3. Results

Table 1 presents the summary statistics by condition (height, repetition) and repeated measures ANOVAs for solution times for corrects, and keying times for corrects. For problem solving accuracy, we use a mixed-effects logistic regression with a random intercept for each participant. All three measures show large and highly significant effects of repetition. The repetition effect on keying is noteworthy because there is no difference in the number of characters that have to be entered for non-repeated and repeated problems. There are also highly significant effects of height on problem solving latency, as well as strong interactions of height with repetition. Much of the interaction is driven by the decrease in accuracy and increase in solving time for non-repeated items of increasing heights. A different pattern is observed in repeated items. Similar to non-repeated problems, height 3 repeated items are solved significantly faster than height 4 items (t(39) = 6.03, p < .005)¹ and with greater accuracy according to a Wilcoxon signed rank test(V = 105.5, p < .005). The repeated items, however, show worse performance on height 4 items than height 5 items, which is the opposite direction of the effect for non-repeated items. While small, the difference between height 4 and height 5 is significant both for accuracy (V = 139, p < .05) and solving time (t(39) = 2.59, p < .05).

Because repetition had a major effect on keying time, in further analyses we combined the solving and keying time as the total duration of problem solving, and the fMRI analyses focused on the scans that combine the solving and keying periods. Items in which participants timed out while solving or recording the answer were included in our analyses. A total of .14% of items were not solved within

¹ All *t*-tests are 2-tailed.

Table 1

Basic performance measures shown on the left side of the table. On the right side of the table we show the results of an ANOVA for solving and responding times and the results of a mixed-effects model of accuracy scores.

		Height				d.f	F	p-Value
Solving (s)	Non-repeated Repeated	4.7 2.2	7.5 2.9	9.5 2.5	Repetition (R) Height (H) R \times H	1, 39 2, 78 2, 78	139.5 61.8 39.4	<.001 <.001 <.001
Responding (s)	Non-repeated Repeated	2.1 1.6	2.2 1.7	2.2 1.4	Repetition (R) Height (H) $R \times H$	1, 39 2, 78 2, 78	74.2 1.9 1.5	<.001 .16 .22
					_	Slope (S.E.)	z-Value	
Accuracy	Non-repeated Repeated	83.9% 96.2%	74.8% 92.6%	70.3% 95.1%	Repetition (R) Height (H) R × H Intercept	1.1 (.26) .11 (.08) .33 (.12) -3.3 (.22)	4.4 1.3 2.8 -15.1	<.001 .2 <.005 <.001

the 30 s provided for solving, and a total of 2.4% of solutions were not keyed in within the 5 s allotted. We also chose to include both correct and incorrect items under the assumption that participants were attempting to execute the same algorithms for both correct and incorrect solutions. Fig. 5 shows the latency as a function of amount of practice (which for repeated problems means the number of times that problem has been repeated while for non-repeated problems it means how many non-repeated problems have been seen). Purely for descriptive purposes, it also shows best-fitting power functions of the form $\alpha n^{-\beta}$ where *n* is the number of repetitions. While the speedup for non-repeated items is small, the coefficient of a linear regression of non-repeated latency on exposure reflects a significant speed up (t(39) = 4.79, p < .0001). Both repeated and non-repeated items start out showing a large effect of height, but for the repeated items the height differences compress rapidly with continued exposures. There is an interesting crossover between height 4 and height 5 repeated items. Height 5 items are slower on first presentation by 3.38 s (t(39) = 3.62, p < .001), but by the second half of the experiment they average .59 s faster (t(39) = 3.57, p < .001). The difference between height 3 and 5 repeated problems disappears in the second half of the experiment (in the second half, height 3 are 2.97 s and height 5 are 2.88 s, t(39) = 0.60, p < .5).

The remainder of this Results section consists of three parts. First, we will describe how we used the HSMM–MVPA method to provide a bottom-up characterization of the structure of the task and the phases of learning. Second, we will enrich this characterization using the top-down guidance of an ACT-R model of the task. Third, with the theoretical interpretation in place, we will investigate the learning mechanisms that are responsible for both transitions between learning phases and learning within cognitive stages.

3.1. Bottom-up stage analysis: cognitive stages within learning phases

Our analysis of the data will track the cognitive stages that participants go through within different learning phases. As we described in the introduction, Tenison and Anderson (2015) have found evidence that participants go through three learning phases as these problems are practiced. Anderson and Fincham (2014a) found evidence that during a single trial on these items participants go through three cognitive stages in the solution of a Regular problem: Encoding, Solving, and Responding. Fig. 4 illustrates how to put these two analyses together into a single HSMM. Across 36 trials of a repeated problem, we assume that participants progress through some number of learning phases. At the trial level within any particular learning phase, they transition through the three cognitive stages within the same learning phase for the next problem or graduate to the Encoding stage of the next learning phases. The brain signature for any particular cognitive stage is assumed to be constant across all learning phases; however, the durations of the stages can vary across learning phases, with the expectation that they will decrease and some may effectively shrink to zero. We assume durations of stages are



Fig. 5. Latency of problems by height at different points in the experiment. Data are plotted separately for repeated and non-repeated problems. The points represent the observed mean times and the lines are the best fitting power functions of the form αn^{-B} .

sampled from Gamma distributions. As in Anderson and Fincham (2014b), we estimated just one scale parameter per cognitive stage for each learning phase (i.e., if there are N learning phases, we estimated $3 \times N$ scale parameters), but we also estimated a separate shape parameter to vary for each of the 21 problems (i.e. $3 \times 21 \times N$ shape parameters if we assume N learning phases).²

The estimation process involved fitting the data from both the repeated and non-repeated items. All problems started in Learning phase 1 and repeated problems could graduate to a later phase with some probability after each trial. The non-repeated items (i.e. the 18 items seen twice) were coded as 36 unique trials and therefore were only estimated to be in Learning phase 1. In addition to estimating brain signatures and stage durations, we estimated this graduation probability between each pair of successive learning phases, assuming that graduation probability applied equally for all problems. For each item solved by each participant the model generates (a) the likelihood of belonging in each learning phase and (b) the estimated time spent in each cognitive stage within those phases. This avoids issues associated with averaging learning data (Haider & Frensch, 2002; Heathcote et al., 2000; Myung, Kim, & Pitt, 2000). We conducted a leave-one-participant-out-cross-validation (LOOCV) to determine the number of learning phases that best described the data. This involved fitting the data to 39 participants and using the parameters to calculate the likelihood of the data from the 40th. Fig. 6 shows the results in terms of the average participant log likelihood gain over the simplest case of a 1-Learning phase model. The results are particularly clear; three learning phases improves the predictions for all 40 of the participants over fewer learning phases, while additional learning phases offer no further improvement. This confirms the conclusion of Tenison and Anderson (2015) that there are three phases of skill acquisition.

Fig. 7 compares the brain signatures for the three cognitive stages of problem solving identified by Anderson and Fincham (2014b) (Fig. 7a) with the brain signatures for this experiment (Fig. 7b), which were estimated independently. The correspondence is quite striking. Specifically, the mean correlation among corresponding stages between the experiments is .92 (compared to the mean between-experiment correlation of .78 for non-corresponding stages across experiments and .82 for the mean correlation between non-corresponding stages within both experiments).³ The mean absolute difference in voxel activations for corresponding stages between the experiment sis small, averaging a 0.11% difference (compared to the between-experiment difference of 0.30% for non-corresponding stages

 $^{^2}$ This can be conceived of a stage involving different numbers of steps for a problem (shape parameter) that each take a fixed time (scale parameter), but the real motivation is to limit the number of free parameters.

³ These statistics are for the 8245 regions the two studies have in common.



Fig. 6. Performance of different numbers of learning phases in LOOCV.



Fig. 7. The brain signatures of the 3 stages involved in solving regular pyramid problems. (a) Anderson and Fincham (2014b); (b) Current experiment. Dark red reflects 1% above baseline while dark blue reflects 1% below baseline activity. The *z* coordinates for a brain slice (radiological convention: image left = participant's right) is at x = y = 0 in Talairach coordinates. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and 0.16% between non-corresponding stages in the first experiment and 0.19% in the second experiment).

For each problem, we obtained a weighted average of the estimated time spent in each cognitive stage within each learning phase. This average was computed as the sum of stage occupancies in that stage (Fig. 2) divided by the probabilities that the problems were in that learning phase (Fig. 3B). Fig. 8 shows the resulting mean times spent in each stage organized by height. In order to measure the effect of height, we performed a mixed-effects ANOVA where we treated the 21 problems as random effects and the cognitive stages and learning phases as fixed effects. Of the 9 cognitive stages only one shows



Fig. 8. Estimated mean time in each cognitive stage during each learning phase.

a significant effect of height and it shows it quite strongly: the Solving stage in Learning phase 1 (F (2,18) = 106.96). Every height 3 problem has an estimated duration in this stage shorter than any height 4 problem, which has an estimated duration shorter than any height 5 problem. Thus, even though this estimation process has no information about the height of problems or about the interpretation of the phases, it finds that the effect of this height is on Solving stage and only in Learning phase 1. The Solving stages for Learning phases 2 and 3 are estimated to be strikingly brief (.07 and .03 s), in contrast to the times for the Solving stage in Learning phase 1. While the Encoding stage does not show an effect of height, it shows a steady decrease with learning phase: 3.26 s in Phase 1, 2.44 s in Phase 2, and 0.54 s in Phase 3. Pairwise comparisons of Phases 1 and 2, and Phases 2 and 3 are quite significant (t(20) = 8.94, p < .0001 and t(20) = 17.59, p < .0001). The Responding stage also decreases across the learning phases (2.68, 2.32, and 2.20 s), but only the decrease from Learning phase 1 to 2 is significant (t(20) = 3.28, p < .005).

The passage through the learning phases is the major factor driving the decrease in times that we see in Fig. 5. However, more is happening than just a decrease in time. The proportion of time in different cognitive stages is dramatically changing—it is roughly equally divided among the three cognitive stages in Learning phase 1, the Solving stage nearly disappears in Learning phase 2, and the Encoding stage shrinks dramatically in Learning phase 3. We now turn to an ACT-R model of what is happening in these three learning phases.

3.2. Top-down theory-driven analysis: identity of cognitive stages and learning phases

Fig. 9 shows the ACT-R models for performance in the three learning phases, in each case solving the problem 8\$4. As mentioned in the introduction, a standard model for speedup is learning the answer and retrieving it, and this is what drives the transition from Learning phase 1 to Learning phase 2. A further phase of learning can happen within ACT-R when a new production rule is learned that directly maps the stimulus to the response without any retrieval. This is what is driving the change from Learning phase 2 to Learning phase 3. We will now describe the details of the ACT-R processing in each of these learning phases.

1. *Learning phase 1:* The model for Learning phase 1 is the minimal adaption of the model in Anderson and Fincham (2014b) (Fig. 1) to correspond to the current task. Only three characters are presented, rather than the five in Anderson and Fincham (2014b), and so the visual encoding is reduced. On the other hand, we assume that during the Encoding stage participants check problems to see if they have a stored answer. The long retrieval in the Encoding stage reflects the failed effort to retrieve an answer. The Solving stage has identical structure to Anderson and Fincham. The Responding stage differs from Anderson and Fincham only in that there is an additional key press to begin the keying of the answer.



Fig. 9. A swimlane representations of the engagement of ACT-R modules during the three learning phases when solving 8\$3.

- 2. *Learning phase 2:* The major difference between Learning phase 1 and Learning phase 2 is that the answer is retrieved and there is no need for a Solving stage. The Encoding stage is briefer because the retrieval process succeeds in this stage before it would have timed out in Learning phase 1. In addition, we assume that the participants now skip the encoding of the dollar-sign operator, and rather just encode base and height. In contrast to what might be concluded from the analysis in Fig. 8, there is no shortening of the Responding stage in the model.
- 3. *Learning phase 3:* In the third learning phase, the problem is recognized as a single unit and a production maps this onto the response, resulting in a very short Encoding stage. The Responding stage is a bit faster because there are no longer visual checks of whether the correct keys are actually struck.

The times of the cognitive stages are determined by how long the individual modules take. We used the timing parameters from Anderson and Fincham (2014b). The only additional parameters estimated were the time for failed retrieval (1.35 s), which influences the time of the Encoding stage in the first learning phase, and the time of the successful retrieval (0.95 s), which influences the time of the Encoding stage in the second learning phase. Fig. 10 shows the ACT-R times for the stages according to the models for each learning phase. It shows large effects of height only for the Solving stage of the first learning phase because height controls the number of additions that need to be performed in that stage.

As in Anderson and Fincham, these times can form the basis for an HSMM by constraining the HSMM to have a Gamma distribution of times in a stage for a problem with the mean time of the ACT-R model. Then, we estimate just a single scale parameter *a* per cognitive stage for each learning phase. Given the model's mean time *t* for each stage of each problem, we can directly compute the shape parameter for a stage as t/a. This contrasts with the model fit in procedure used in the HSMM–MVPA technique where we also estimate a shape parameter per stage per problem. Thus, rather than estimating 3 stages \times 3 learning phases \times 21 problems = 189 distributional parameters, we just estimate 7 scale parameters (there is no Solving stage for Learning phases 2 and 3). In LOOCV, this model-constrained HSMM outperforms the unconstrained HSMM–MVPA model, fitting 28 of the 40 participants better than the 3-stage unconstrained HSMM (significant by a sign test, *p* < .005). Its mean log-likelihood advantage is 6.0 per subject. The brain signatures of corresponding stages for the unconstrained and the model-constrained fit are essentially identical (an average .997 correlation)



Fig. 10. ACT-R predictions of mean time in each cognitive stage during each learning phase.

for the 20 PCA scores that define these brain signatures). So the advantage of the model-constrained approach in LOOCV is that it avoids the overfitting of the many time parameters.

While there is a strong correlation between the estimated model times shown in Fig. 10 and the bottom-up estimates shown in Fig. 8, they are not identical. The estimation of so many parameters in Fig. 8 is problematic in that, in principle, many different combinations of parameters can yield nearly identical fits. The top-down constraints of ACT-R serve to indicate which of these nearly equivalent fits will generalize better.

While the HSMM–MVPA analysis uses whole brain activation, the ACT-R model makes predictions for activity in specific brain regions (Anderson, 2007). As discussed by Anderson and Fincham (2014b), these brain regions are selected because they are consistently active across diverse tasks that share similar demands on specific ACT-R modules. Fig. 11 shows the brain mappings of the 5 modules used in our model (Fig. 9). We used the stage occupancy probabilities from the HSMM model to estimate brain activity in each of these regions for the three cognitive stages in Learning phase 1, the two cognitive stages in Learning phase 2, and the combined Encoding and Responding stage in Learning phase 3.⁴ The two stages in Learning phase 3 were combined because the Encoding stage was too brief for accurate estimation. Fig. 11 compares these brain activation measures derived from the HSMM–MVPA model against predictions from the proportion of time a module was engaged in a stage according to the ACT-R model. To get the predictions we regressed these modules proportions against the activity measures to obtain an intercept (baseline activity in the region) and a slope (how activity in the region scaled with activity in the module). While we estimated these parameters to get the predictions displayed in Fig. 11, the correlation between the activity measures and module proportions provide parameter-free measures of the correspondence. Except for the Posterior Parietal Cortex (PPC) these correlations are greater than .9. The PPC correlation is negative, but the activity is quite high during all stages. As can be confirmed from Fig. 9, the Imaginal module is engaged at a relatively constant rate during all stages. Thus, there is not much predicted variance to correlate with. The other modules all have distinctive patterns of engagement across stages which is confirmed by the activation in the corresponding regions: The Retrieval module is most engaged during the Encoding stage of Learning phase 1 and 2 where retrieval is engaged, and then in Solving stage of Learning Phase 1 when addition facts need to be retrieved. The Retrieval module is least engaged in the Response stages. The Metacognitive module is only engaged in Encoding stages of Learning phases 1 and 2 when the decision to compute or use retrieval is being

⁴ We created a design matrix (Friston, 2011) that had 10 variables associated with each scan of each problem. There were 6 stage regressors and 4 binary variables indicating whether or not a scan occurred during fixation, feedback, repetition-detection, or problem-solving. The stage regressors are just the stage occupancy probabilities that the scan was in that cognitive stage and learning phase (see Fig. 2). These were then convolved with the SPM hemodynamic function and regressed against activity in that region to get an estimate of activity associated with each variable. The values plotted in Fig. 11 are the estimates associated with the 6 stage regressors.



Fig. 11. Average activation in predefined brain regions as a function of cognitive stage within learning phase and the predictions based on module activity: *Rostrolateral prefrontal cortex (RLPFC – Metacognitive module)*: a 12.8 mm. (high) by $15.6 \times 15.6 \text{ mm}^2$ region centered at Talairach coordinates -34, 47, 8; *Lateral inferior prefrontal cortex (LIPFC – Retrieval module)*: a 12.8 mm. (high) by $15.6 \times 15.6 \text{ mm}^2$ region centered at Talairach coordinates -34, 47, 8; *Lateral inferior prefrontal cortex (LIPFC – Retrieval module)*: a 12.8 mm. (high) by $15.6 \times 15.6 \text{ mm}^2$ region centered at Talairach coordinates -43, 23, 24 spanning Brodmann Areas 9 and 46; *Posterior Parietal Cortex (PPC – Imaginal module)*: a 12.8 mm. (high) by $15.6 \times 15.6 \text{ mm}^2$ region centered at Talairach coordinates -23, -63, 40 spanning Brodmann Areas 9 and 46; *Motor – Manual module*: a 12.8 mm. (high) by $15.6 \times 15.6 \text{ mm}^2$ region centered at Talairach coordinates -42, -19, 50, involving Brodmann Areas 2 and 4 at the central sulcus; *Fusiform – Visual module*: a 9.6 mm. (high) by $12.5 \times 12.5 \text{ mm}^2$ region centered at Talairach coordinates -42, -60, -8 in Brodmann Area 37.

made. The Manual module⁵ is only engaged during the Responding stages; and the Visual Module is engaged in all stages but the Solving stages. The correspondence is not perfect: except for the Rostrolateral prefrontal cortex (RLPFC) and Lateral inferior prefrontal cortex (LIPFC), the predicted values fall sufficiently outside of the standard error of the estimates as to be significant. Nonetheless the correspondences are striking and support our interpretation of the cognitive stages.

3.3. Analysis of learning

The fits we have considered so far have focused on performance within the three learning phases and have not considered the learning factors that are driving the transitions between phases or the possible learning within a learning phase. We can obtain from the HSMM–MVPA probabilities, prob

 $^{^{5}\,}$ The counting in the Solving stage is associated with the left hand and right motor region.



Fig. 12. Average probability of a trial being in each phase on each of 36 repetitions of a problem. In red we show the model prediction of the probability of learning phase given the transition probabilities. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(i,k), of an item being in learning phase *i* on any of the 36 trials *k*. These probabilities will allow us to explore a number of issues relevant to learning.

Fig. 12 shows the average probabilities, prob(i,k), that for each height the repeated problems are in one of the learning phases over the 36 trials.⁶ It also shows the predictions from the model given the transition probabilities are the same for all heights. These are derived from the estimated per-trial transition probabilities between phases (.202 from Phase 1 to Phase 2 and .094 from Phase 2 to Phase 3). Table 2 shows the mean number of trials in each phase (summed probabilities over the 36 trials). There are significant differences among the number of trials spent in Learning phase 1 for the three heights (*F* (2,78) = 9.05, *p* < .0005) and Learning phase 3 (*F*(2,78) = 8.73, *p* < .0005), but only a marginal effect in Learning phase 2 (*F*(2,78) = 3.10, *p* < .1). Pairwise *t*-tests confirm significant difference between height 3 and 4 (t(39) = 3.63, *p* < .001 for Learning phase 1; t(39) = 2.58, *p* < .05 for Learning phase 3) and between height 4 and 5 (t(39) = 3.07, *p* < .005; t(39) = 4.17, *p* < .0005) but not between heights 3 and 5 (t(39) = .23; t(39) = 1.49). Basically, participants take longer to leave the first learning phase for height 4 problems and consequently get to the last learning phase later and so, spend fewer trials in it. This resulted in the longer average latencies for height 4 problems (Table 1 and Fig. 5), but no difference in estimated stage occupancies in Learning phases 2 and 3 in Fig. 8.

Our interpretation of the probability of transitioning between Phase 1 and Phase 2 is that it reflects probability of memorizing the answer. If this interpretation is correct, then height 4 problems might be harder to memorize even if one did not know their mathematical interpretation. To test whether or not this was true, we ran an Amazon Mechanical Turk experiment with 100 participants who were just asked to memorize paired associates consisting of a stimulus like 8\$4 and a response like 26, without any mathematical explanation (see Appendix A for experimental details). As material, we used the pyramid problems from this experiment. Focusing on the 65 participants who performed between 20 and 80 percent correct, the mean recall probabilities were .63, .37, and .63 for heights 3, 4, and

⁶ These are estimated from the ACT-R HSMM.

Table 2

Number of trials in each phase for repeated problems and model predictions.

	Mean number of trials				
	Phase 1	Phase 2	Phase 3		
Height 3	3.6	10.4	22.0		
Height 4	6.8	11.7	17.5		
Height 5	3.8	7.8	24.4		
Model	5.6	10.9	19.6		

 $5.^7$ The mean difference between height 4 and the other two was highly significant (t(64) = 9.86, p < .0001). We speculate that the reason height 3 and height 5 problems are easier to memorize is because they overlap multiplications facts—n\$3 is always similar to a multiple of 3 (e.g., $8\$3 = 21 = 7 \times 3$), n\$5 is always similar to a multiple of 5 (e.g., $8\$5 = 30 = 6 \times 5$), but n\$4 is not a multiple of 4 (e.g., 8\$4 = 26). Whether this speculation is correct or not, the Mechanical Turk results support the interpretation of the transition from Phases 1 to 2 as a matter of remembering the answer.

3.3.1. Contrasting the theoretical assumptions of learning with the practical assumptions of our model

While hidden Markov modeling offers dynamic programing solution for identifying cognitive stages and learning phases involved in skill acquisition, there are assumptions which underlie the semi-hidden Markov model we used in this paper. The Markov property, which is central to the model we use, assumes that each state is dependent only upon the previous state. This assumption informs additional assumptions governing both transitions between phases and durations within cognitive stages. We investigate here the degree to which the assumptions of HSMM theory and the assumptions built in the structure of our model are met in this paper's HSMM-MVPA model. The Markov assumption, which underlies our estimation software, implies that the future behavior of the model depends only on what learning phase it is in on the current trial and does not depend on how long it has been in that phase. This assumption is not true of the ACT-R model or a number of other models of skill acquisition. It was also not true of the Tenison and Anderson (2015) model (Fig. 3), where there was a different HMM state to keep track of the number of trials in each learning phase; however, we can use the results that we have obtained to test for possible deviations from the assumptions in the current model. As we have already seen with respect to the differences in latency and learning rate for problems of different heights, just because an assumption is built into the estimation process (such as the assumptions that behavior across heights is uniform), we are not prevented from gathering evidence from the estimation process that it is wrong. We will use the estimated prob(i,k) to explore three issues where the HSMM assumptions are at odds with common assumptions about skill acquisition.

3.3.1.1. Issue 1: Sigmoidal phase transitions. The Markov property assumes that on every trial there is a constant probability p of transitioning into the next Learning phase. If p is the probability of transitioning from Learning phase 1 to Learning phase 2, the probability that the transition from Learning phase 1 to Learning phase 2, the probability that the transition from Learning phase 1 to Learning phase 2 will occur on trial n is

$$p_{12}(n) = p \times (1-p)^{(n-1)}$$
(1)

This exponential learning function is not what many models of learning would imply, including the ACT-R model. Repeated encounters with a memory in ACT-R increase the strength of the memory as a logarithmic function and when the memory exceeds a threshold it can be recalled. The probability of being able to recall an item on the *n*th trial (i.e., advancing to Learning phase 2) is

⁷ We chose this focus to avoid participants who were not trying or perhaps copying answers. For all participants the probabilities were .72, .56, and .71 (t(99) = 7.69, p < .0001).

$$p_{12}(n) = \operatorname{normal}(\log(n), \tau, \sigma) \tag{2}$$

19

where τ is the threshold and σ represents the variability of individual items.⁸ The exponential function implies that the first trial will be the trial of maximum transition while the sigmoidal function allows for the possibility that the trial of maximum transition can be later. Transition between Learning phases 2 and 3 is driven by production compilation. The newly formed production has to build up enough utility to be preferred over the older production. Because of variability in utility, the transition probability for using a new production follows a similar sigmoidal function (Fig. 5.10 in Anderson (2007)). To test between the exponential (1) and sigmoidal (2) transition functions, we examined whether our estimated state occupancies better fit the exponential equation or the sigmoidal equation. While we have indicated a basis for the sigmoidal equation in ACT-R, the assumption is not unique to ACT-R. So we are really testing two general theories about the likelihood of transition between learning phases.

We fit the exponential learning Eq. (1) and the sigmoidal learning Eq. (2) to the probabilities, prob (i,k), of being in the three phases. To make the parameters of the sigmoidal and the exponential models equal we set σ in the sigmoidal equation to 1. For each repeated problem the HSMM delivers 108 prob(i,k)'s of being in one the three phases *i* on one of the 36 trials *k*. If $p_{12}(n)$ is the theoretical probability of transitioning between Learning phase 1 and 2 after n trials in Learning phase 1 and $p_{23}(n)$ is the theoretical probability of transitioning between Learning phase 2 and 3 after n trials in Phase 2, then the probability of the estimated prob(i,k)'s for an item is

$$\sum_{i=1}^{36} \sum_{j=i+1}^{36} \left[p_{12}(i) \times p_{23}(j-i) \prod_{k=1}^{i} prob(1,k) \prod_{k=i+1}^{j} prob(2,k) \prod_{k=j+1}^{36} prob(3,k) \right]$$
(3)

The exponential and sigmoidal equations differ in their calculation of $p_{12}(n)$ and $p_{23}(n)$. We fit these models to the data allowing for separate parameter estimates for heights 3, 4, and 5. Table 3 shows the parameters estimated for the two models. These parameters again confirm the difficulty of transitioning out of Learning phase 1 when practicing height 4 problems (either in the lower transition probability for the exponential equation or the increased threshold for the sigmoidal equation). Fig. 13 contrasts the estimated probabilities of transitioning between phases according to the two models, averaging over the three heights. As can be seen, the maximum trial of transition is delayed according to the sigmoidal equation relative to the first trial prediction of the exponential equation. The sigmoidal equation provides slightly better fits in terms of log likelihood for all three heights. Since the number of parameters for the two models is equivalent we can double the total log likelihood difference to get a Bayesian Information Criterion (BIC) difference of 12.7, which provides strong support of the sigmoidal equation as the best fitting model. Given that we were using state occupancies estimated from a Markov model, which embodies the exponential assumption, the fact that a sigmoidal function does better suggest there is a delay in the trial of maximum transition.

3.3.1.2. Issue 2: Within-phase speedup. Another assumption of the Markov property is that the duration of all trials in a learning phase should have the same average duration. However, a number of models reviewed in the introduction, including ACT-R, allow for within-phase speedup. In the ACT-R model, there should be speed up in both the first and second learning phases since both involve retrieval of new declarative information. In Learning phase 1 this occurs during the Solving stage where mathematical retrievals are required in order calculate the pyramid value, while in Learning phase 2, this occurs during the Encoding stage where participants practice retrieving the newly learned information about the value of a specific pyramid problem. Tenison and Anderson (2015) fit models to latency data that had within-phase speedup or no speedup and found evidence for shallow within-phase speedup as illustrated in Fig. 3C. They did not address whether separate phases might speed up and others not, but rather tested an assumption of uniform within-phase speedup versus none.

⁸ In ACT-R the variability σ in this equation both reflects variability in items and participants, and variability in trial-to-trial activation for a single item. The latter allows for the possibility that an item may regress from a later learning phase to an earlier learning phase as discussed previously. In this analysis we assume all of the variability is among items.

Fit of exponential and sigmoidal learning equations.

Table 3

20

	p for 1–2 transition	p for 2–3 transition	Log Likelihood
Exponential equation			
Height 3	.23	.07	-257.7
Height 4	.15	.05	-287.3
Height 5	.28	.09	-251.2
Total			-796.3
	au for	au for	Log
	1–2 transition	2–3 transition	likelihood
Sigmoidal equation			
Height 3	.41	2.3	-254.5
Height 4	1.3	2.7	-286.6
Height 5	.8	2.2	-248.8
Total			-789.9



Fig. 13. Estimated probability of transitioning between the learning phases according to the two learning equations.

To address this issue we fit 9 separate power functions of the form $\alpha n^{-\beta}$ for the 3 × 3 combinations of learning phase and height. We have no strong commitment to the power function; it is just being used as representative of a negatively accelerated speed up.⁹ For each item for each individual on each trial *n*, we measured the prediction error as the sum of the squared differences between the predicted and observed times for each phase weighted by the probability that the item was in that phase on that trial:

$$\sum_{i}^{3} prob(i,n) (A_{ij} n^{-b_{ij}} - t(n))^{2}$$
(4)

where A_{ij} is the scale and b_{ij} is the exponent associated with phase *i* and height *j*. With 9 combinations of phase and height and 2 parameters there are 18 potential parameters to estimate, but there proved to be considerable potential for parameter simplification. Table 4 summarizes the search process as we went from the full model to simpler models. The first line gives the full model with all 18 parameters, which has a residual sum of squares (RSS) of 15,738. From this we can calculate the BIC measure as

$$BIC = n \times \log\left(\frac{RSS}{n}\right) + k \times \log(n) \tag{5}$$

⁹ Tenison and Anderson (2015) could not discriminate between exponential and power speedup for their data but did find evidence for a power law in a reanalysis of data from Rickard (1997). The fits of a power law to both datasets estimated a 0 intercept as in the power function we are using here.

Table 4	
Parameter estimation for speed up modeling.	

Model comparison	Number of parameters	Residual sum of squares	BIC
Full model	18	15,738	5735.7
1 Exponent per phase	12	15,894	5727.9
Exponent only in Phase 1	10	15,911	5716
3 Scales only for Phase 1	6	15,925	5685.3

Table 5

Parameters of best-fitting speed up model.

Phase 1				Phase 2	Phase 3
Exponent	Height 3	Height 4	Height 5	Factor	Factor
0.077	9.4	11.8	14.0	4.8	2.7

where *n* is the 4320 observed latencies and *k* is the number of model parameters. The lower BIC values reflect better model fits, so as the number of parameters a model uses increases the better the fit must be for those parameters to be justified. Estimating a single exponent per learning phase leads to a reduction of 6 parameters and a decrease in the BIC measure of 7.8. The question of whether there is speed up in a phase comes down to the question of whether the exponent for that phase is non-zero. We did a search of the 8 (i.e., 2^3) possible restrictions of the exponent to zero for phases and the best fitting model only estimated a learning exponent for the first phase. The closest competitors also estimated an exponent either for the second phase or third phase but both had a BIC measure 6.1 larger. The model with just an exponent for the first learning phase has a BIC measure more than 10 less than the model that assumes exponents for all phases, and almost 60 less than a model that has no learning exponents. It turns out that we can further simplify the model by assuming that the scale factor varies with height for only Learning phase 1. The parameter estimates of the final model are given in Table 5. The Phase 1 factor in Table 5 increases about 2 s with each height, which corresponds to the magnitude of the effect of height on the Solving stage for Learning phase 1 (see Fig. 10). The exponent for the first learning phase, .077, is rather small but it does imply a speedup of more than a second during the average time that a problem is in Phase 1. So, even with an HSMM estimation that assumes a constant time in a stage, there is evidence for speed up in the first stage. We should note that the failure to find evidence for speed up in the other phases does not mean that there was not speedup in these phases, only that it was not substantial enough to justify the estimation of the additional parameters.

3.3.1.3. Issue 3: Regression to earlier learning phases. In the Markov model we employed, once a learning phase has been reached the model either stayed in that phase or graduated to a later learning phase. There was not the possibility of regressing back to an earlier learning phase. This is in contrast to many models of learning that assume retrieval failures and successes can alternate, with probability of success increasing with practice (Averell & Heathcote, 2011; Rubin, Hinton, & Wenzel, 1999). Similarly in ACT-R, while a newly formed production rule increases in its probability with usage, it remains possible to regress back to using an older procedure. Depending on the steepness of the learning curve, such backward regressions can be common or rare. To assess whether there are backward regressions, we need to independently assess what phase a participant has reached on a trial and whether they have regressed on that trial.

The calculation of phase occupancy probabilities prob(i,k) can be decomposed into two components that allow a separation of what the history implies about the phase of an item and what phase the item is actually in on a trial. First, the fMRI signal on trial k by itself provides some information about the phase on that trial: *signal*(k,i) will denote the probability of phase given that signal.¹⁰ Second, the signals on all prior and subsequent trials, in conjunction with the transition probabilities,

¹⁰ The HSMM–MVPA analysis calculates the log conditional probability of the imaging data given each state. This can be converted into the probability that is signal(k,i) by assuming each state is equally likely.

provide evidence as to what the phase should be independent of the signal on that trial: context(k,i) will denote the probability from the context. With three learning phases we can write:

$$prob(i,k) = \frac{context(i,k) \times signal(i,k)}{\sum_{i=1}^{3} context(i,k) \times signal(i,k)}$$
(6)

The two quantities, context(i,k) and signal(i,k) can be calculated independently will allow us to assess whether there has been a regression to an earlier trial.

There is a fairly strong relationship between the context probabilities and the signal probabilities. Fig. 14 shows how the signal probabilities for the three learning phases vary as a function of the context probabilities. While this is aggregated data, the individual trial probabilities are also quite correlated (r = .718, .544, and .743 for learning phases 1–3). The signal probabilities tend to be less extreme (standard deviation of .33 across trials) than the context probabilities (standard deviation of .43). This implies that the net force of the context of a trial is more informative about its phase than the signal. In some cases, this is particularly obvious. For instance, on the first trial an item must be in Learning phase 1. Thus the context probability for Phase 1 on this trial is 1 in these cases while the average signal probability is 0.6.

We looked for cases where there was strong disagreement between the context and the signal, specifically situations in which the context strongly implied that the participant was in a later learning phase than was implied by an equally strong signal. These are cases where it is likely that participants regressed to a previous learning phase. For instance, there are 27 trials (.6% of all trials) where the context probability of being in Learning phase 1 is less than 0.1 but the signal implies a probability greater than .9. One might suspect these of being cases where the participant has regressed back to calculating the answer. Similarly, there are 68 trials (1.6% of all trials) where the contextual probability of being in the Autonomous Learning phase 3 is greater than .9 and the signal probability is less than .1. One might suspect these are cases where the participant has regressed from the Autonomous learning phase. Even if there were not regressions, we would expect a few random occurrences of discordant cases like this (while the probabilities are extreme, they are not 0 and 1). However, in Appendix B we



Fig. 14. The values are the average context and signal probabilities for each decile of the context probability. They are based on the context and signal probabilities for each trial.

show that these frequencies are greater than would be expected by chance from the context and signal probabilities. Thus, despite the assumption of no backwards transitions we used to do our modeling, we have evidence that participants can slip back a phase for a single trial, as would be expected by many learning models.

3.3.2. The impact of the Markov assumption and our modeling decisions

While the violations we have identified of the Markov model are not particularly large, one can ask why we premised our estimation on a model that embodies such assumptions. The answer is that it enabled us to use the HSMM machinery to track the trajectory of individual problems and estimate parameters. The average repeated problem involved anywhere from 41 to 165 scans (mean 76.3) over the 36 repetitions. There are 9 cognitive stages (3 stages for each learning phase). If we had to consider each possible interpretation of the history of a problem that covered N scans over its repetitions we would be considering at 9^N interpretations per problem. The Markov assumption, which is the source of the violation of Issues 1 and 2, is that future behavior only depends on current state. This enables efficient computation of the implications of all of these interpretations without separately computing each.

Beyond violations of the Markov assumption, we introduced the assumption that there was no regression to previous Learning phases when building our model (Issue 3). We could have put backward transitions into an HSMM, but we judged that, in combination with the Markov assumption, this would lead to a very implausible model of human learning. The probability of a backward transition cannot change with duration in the phase, so if added to the model it would be just as likely to have a failed recall trial after the first successful recall trial, as after a sequence of 10 successful recall trials. Also, once having transitioned backward, it might take a good number of trials to get back to the original phase because the return would be governed by the original transition probability. Another possibility would be to have a different HMM state for each combination of numbers of trials in different phases, as in Tenison and Anderson (2015), but including backwards transitions in this model would have led to an explosion of phases and parameters. The current approach seemed the best method for tracking the specific phase history of each problem, which was our major goal. This tracking enables us to identify new detailed features of the learning process and has not prevented us from identifying places where the simplifying assumptions in our analysis machinery were not exactly right.

4. Discussion

This research examines the underlying learning phases and cognitive stages that occur when participants practice a novel procedure. On average, their solution times decrease from about 12 to about 3 s over the course of 36 practice opportunities (Fig. 6). The current study uses an fMRI-guided approach to come to the same conclusion as the Tenison and Anderson (2015) findings that used a latency-guided approach. Both approaches find that participants go through three qualitatively distinct learning phases. The more detailed analysis of trial structure in this paper has allowed us elucidate finer details of the structure of the three learning phases. In the first learning phase, where calculation dominates, participants solved problems by moving through three cognitive stages: Encoding, Solving, and Responding, with the Solving stage involving a sequence of additions determined by the problem's height. On average (Fig. 6), it takes them 5–6 practice opportunities before participants can then recall the solution the next time the problem is seen. Once this happens there is a transition to the second learning phase, where participants skip the Solving stage and transition directly from Encoding to Responding. By the end of the experiment, about 85% of the problems are in the third learning phase, where a simple stimulus-response rule results in a further substantial shortening of the Encoding stage. The theoretical analysis suggests that problems in this final phase are solved in about 2.5 s with most of the time devoted to executing a 4 character keying sequence to enter the answer (Fig. 9).

A major signature of the transition from Learning phase 1 to Learning phase 2 is the drop out of factors reflecting computational complexity, such as height of the problem. There is not the same clear behavioral signature in the transition from Phase 2–3. It was here that imaging data were particularly

important, showing the change in the nature of the processing. While state identification depended on whole brain activity, we examined the activation patterns in Fig. 11 and compared Learning phase 2 (by averaging Encode and Respond) with Learning phase 3. There is more LIPFC activity in Learning phase 2 (.53% vs. .31%, t(39) = 2.18, p < .05), more RLPFC activity in Learning phase 2 (.35% vs. .00%, t(39) = 2.8, p < .05), but more motor activity in Learning phase 3 (.70% vs. 1.03%, t(39) = 3.84, p < .0005). These differences are predicted from the differing mixes of ACT-R modules involved in these learning phases.

The transition from Phase 2 to Phase 3 raises questions about the roles of familiarity and recall during what we labeled the Encoding stage. Previous research on encoding processes suggests that as items become more familiar early frontal activation diminishes in strength (Curran, 2000). Such work agrees with our findings that the involvement of frontal regions during Encoding stage diminishes with practice. Furthermore, fMRI studies have found reduction in the activation of the left dorsal lateral prefrontal cortex (Reas & Brewer, 2013) and left lateral prefrontal cortex (Cabeza, Locantore, & Anderson, 2003) as retrieval becomes easier; however, the diverse functions performed by these regions make it difficult to know the exact role of these regions in memory recall.

As reviewed in the introduction, both Delaney et al. (1998) and Rickard (1997) have proposed models that involve two learning phases. Rickard and colleagues (Bajic et al., 2011; Bajic & Rickard, 2009; Rickard, 2004) associated the first phase with computation and the second phase with retrieval. Across a range of novel tasks, researchers have shown that fitting calculation and retrieval to separate power functions provides a better model of individual skill acquisition than fitting a single function across both strategies (Bajic et al., 2011; Bajic & Rickard, 2009; Delaney et al., 1998; Rickard, 1997, 2004). We found relatively weak learning in the first phase and no learning in the second or third phases. The ACT-R theory would predict second-phase speedup as the retrieval was practiced, but not speed up in the third phase. The ACT-R model of phase 2 would suggest however, that retrieval of the answer is only about 18% of the solution time in the second learning phase (Fig. 9), and most of the time is taken up in parsing the problem and executing the answer. The major speedup after entering Learning phase 2 comes from the later transition to Learning phase 3. It is possible that the evidence seen in these previous efforts for second phase power-law learning could really have reflected this third phase transition.

Like previous work on understanding the learning processes driving skill acquisition (Bajic et al., 2011; Bajic & Rickard, 2009; Haider & Frensch, 2002; Heathcote et al., 2000; Myung et al., 2000) our effort has focused on following the trajectory of individual items. If these items are going through multiple phases rather than a continuous speed up, the challenge becomes identifying the phase a problem is in on a particular trial. Past efforts have relied on measurements like subject reports to make the distinction between calculation and retrieval (Rickard, 2004). While participants probably have some ability to identify when they are engaged in an activity like repetitive addition, we suspect that it is largely beyond introspective powers to discriminate between retrieval and direct recognition. Adding imaging data to latency data facilitates making such discrimination; however, whatever the sources of evidence, one cannot know with certainty what phase a problem is in on a particular trial. We gained evidence about the underlying phases by aggregating the *evidence* produced by our different models and measures across problems (which is quite different than averaging the *data* across problems). Future research on skill acquisition would benefit from considering the changes in information processing involved in practicing a task alongside questions of individual and item level differences.

In addition to making advances in the understanding of skill acquisition, this work also puts forth a method for more thoroughly studying the neural basis of skill acquisition. Prior work studying practice effects on brain activation has been frequently limited to pre- and post-training study designs, and consequently the resulting theories and models have focused on the distinction between static cognitive states (Chein & Schneider, 2005; Petersen et al., 1998; Schneider & Chein, 2003). While these models are not necessarily counter to the ACT-R model which informed this study, they do not focus on the finer-grain analysis that we use to consider the impact of practice on each trial. Future work would benefit from using the bottom-up and top-down approaches discussed in this paper to better understand how learning changes cognition and neural activation. Building and testing theories to describe the correspondence between these changes is important piece in the larger puzzle of understanding the neural basis of cognition.

While our results are specific to the current task, they can be seen as part of a movement to replace the continuous behavioral characterizations that have dominated much of psychology with mixtures of qualitatively distinct processes. For instance, individual differences that are often measured on continuous ability scales might reflect different mixtures of distinct processes. There is considerable current effort (e.g. Scheibehenne, Rieskamp, & Wagenmakers, 2013) to develop rigorous methods to identify such qualitative mixtures. Within the context of such efforts, our research can be seen as a push to elaborate the sequential structure of learning with a similar mixture analysis.

Acknowledgments

This work was supported by the National Science Foundation grant DRL-1420008, a James S. McDonnell Scholar Award, and by Carnegie Mellon University's Program in Interdisciplinary Education Research funded by the U.S. Department of Education grant R305B090023. We would like to thank Jelmer Borst and Christopher MacLellan for their comments on the paper, and Shawn Betts for his help running the Mechanical Turk study. The analyses and models in this paper can be obtained at http://act-r.psy.cmu.edu/?post_type=publications&p=19047.

Appendix A. Mechanical turk experiment

An analysis of the data from our 40 fMRI participants indicated differences in the memorability of the three repeated problems. Height 4 problems took longer to transition from Learning phase 1 to Learning phase 2, than either height 3 or height 5 problems. We considered two possible explanations: (1) participants may develop alternative strategies for solving problems of different heights, or (2) that the numbers sequences of the height 4 problems were less memorable than the other heights. The height 3 and height 5 problems both share some resemblance to multiplication facts. For example, 8\$5 = 30 which overlaps with $6 \times 5 = 30$ and 9\$3 = 24 overlaps with $8 \times 3 = 24$. In order to test these explanations we ran two studies on Amazon Mechanical Turk in which participants answered repeated and non-repeated pyramid problems. Although we did not collect demographic information for this sample of 120 participants, previous mathematical problem solving research (Lee, Betts, & Anderson, 2015) conducted on Mechanical Turk included a diversity of age (28 years with a standard deviation of 6) and education (roughly 10% masters, 40% 4-year, 40% 2-year or some college, 10% high school). In these previous studies the samples about 55% of the participants were male.

Our first study of 20 participants set out to establish whether when given practice on a set of pyramid problems participants recruited from Mechanical Turk, who complete the experiment on computers in unknown environments, would display the same memorability differences as our sample collected in the fMRI scanner. The structure of the experiment was nearly identical to the one run in the fMRI scanner with some slight modifications. First we removed the n-back task separating the problem solving trials because we were not collecting brain response data. Second, we increased the number of practice blocks to 9, thus giving participants 42 opportunities to practice the repeated problems. These problems were interspaced with low-practice items seen 3 times over the course of 9 blocks. Although we could not perform the HSMM-MVPA analysis of this paper on our Mechanical Turk data set, a repeated-measures ANOVA showed significant differences in response latencies between problems of the three heights (f(2,38) = 6.95, MES = 118.1, p < .005) with height 4 (M 2.8, SD 2.7) problems taking on average longer to solve than height 3 (M 2.2, SD 2.1) or height 5 (M 2.7, SD 3.2). After completion of the study we had participants report if they used any strategies to remember the problems, or found any formulas to solve the problems. We were interested if the use of alternative problem solving strategies may explain the learning differences between the three heights. Five participants reported devising some type of formula for solving some of the problems. However examining the distribution of these strategies over height showed formulas were used on 3 height 3 problems, 5 height 4 problems, and 6 height 5 problems. These results provide little evidence that the faster learning of height 3 and 5 problems is due to participants developing alternative problem solving strategies.

Predictor	Coefficient	SE	Wald's Z	<i>p</i> -Value
Intercept (Height 3)	1.4	.18	7.5	<.001
Height 4	97	.13	-7.3	<.001
Height 5	07	.14	53	.59

 Table A.1

 Summary of fixed effects in mixed logit mode (N = 100, BIC = 2180.7).

Our next study of 100 participants investigated the hypothesis that the terms of height 3 and height 5 problems were more memorable than the terms of height 4 problems. The experiment consisted of 4 blocks, during the memorization period of each block the participant would be shown five pyramid problems three times each. These problems appeared as full expressions, such as "8\$5 = 30", and remain on the screen for 2 s. Participants in this study were not given the introduction explaining what the \$ operator meant, instead they were told that they would be memorizing numerical expressions. After the memorization period, each of the 5 expressions would be presented again, but with the solution missing (i.e. 8\$5) and participants would have to recall the value. To test whether across participants the height of the item impacted recall, we ran a mixed logit model using the lmer package in R (Bates, Maechler, Bolker, & Walker, 2014). We included main effects for height, and random intercept for participant (including random effects for height did not improve model fit). The main effects of height supported the significant distinction of height 4 from height 3, but no significant distinction between height 3 and 5 problems (Table A.1).

Appendix B. Upper bounds on the number of apparent regressions by chance

There are 27 trials where the context probability of being in the first learning phase is less than .1 and the signal probability is greater than .9. We can place a bound on how often this would occur by chance. These 27 trials would have to come from one of two categories:

- 1. They actually are in Phase 1. There are 3519 trials overall with a context probability of being in Phase 1 less than .1. These include the 27 trials in question. They have an average context probability of .0027. Thus, we would expect $3519 \times .0027 = 9.4$ trials actually in Phase 1 when the context probability is less than .1.
- 2. They are not in Phase 1. There are 246 trials overall with a signal probability of being in the first learning phase greater than .9 that include these 27 trials. These trials have an average signal probability of *not* being in Phase 1 of .0246, implying that no more than $246 \times .0246 = 6.0$ of the 27 trials are not in Phase 1.

Thus, 9.4 + 6.0 = 15.4 is the expected number of trials in the *disjunction* of being in learning phase 1 with a context probability less than .1 *or* not in Phase 1 with a signal probability greater .9. This number is an upper bound on the expected number of trials in the *conjunction* of these two constraints, which is observed to be 27.

There are 68 trials where the contextual probability of being in learning phase 3 is greater than .9 and the signal probability is less than .1. These 68 trials would have to come from one of two categories:

- 1. They are not in Phase 3. There are 2257 trials overall with a context probability greater than .9 of being in the Phase 3 that include the 68 trials. They have an average context probability of .0046 of *not* being in Phase 3, implying an expected number of $2257 \times .0046 = 10.4$ trials.
- 2. They are in Phase 3. There are 995 trials overall with a signal probability of being in the Phase 3 less than .1 that include these 68 trials. These trials have an average signal probability of being in Phase 3 of .0258, implying an expected number of $995 \times .0258 = 25.6$ trials.

Thus, the expected number of trials in the *disjunction* is 10.4 + 25.6 = 36.0 trials while there are observed 68 trials in the *conjunction*.

References

- Albouy, G., Sterpenich, V., Vandewalle, G., Darsaud, A., Gais, S., Rauchs, G., ... Maquet, P. (2012). Neural correlates of performance variability during motor sequence acquisition. *Neuroimage*, 60, 324–331.
- Anderson, J. R. (1982). Acquisition of cognitive skill. Psychological Review, 89, 369.
- Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem situations. Psychological Review, 94, 192.

Anderson, J. R. (2007). How can the human mind occur in the physical universe? 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, 1036.
- Anderson, J. R., & Fincham, J. M. (2014a). Discovering the sequential structure of thought. Cognitive Science, 38, 322-352.

Anderson, J. R., & Fincham, J. M. (2014b). Extending problem-solving procedures through reflection. *Cognitive Psychology*, 74, 1–34.

- Anderson, J. R., & Lebiere, C. (1998). Cognitive arithmetic. Erlbaum Mahwah.
- Anderson, J. R., Taatgen, N. A., & Byrne, M. D. (2005). Learning to achieve perfect timesharing: Architectural implications of Hazeltine, Teague, and Ivry (2002). Journal of Experimental Psychology: Human Perception and Performance, 31, 749–761.
- Averell, L., & Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55, 25–35.

Bajic, D., Kwak, J., & Rickard, T. C. (2011). Specificity of learning through memory retrieval practice: The case of addition and subtraction. Psychonomic Bulletin & Review, 18, 1148–1155.

- Bajic, D., & Rickard, T. C. (2009). The temporal dynamics of strategy execution in cognitive skill learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35*, 113.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). Ime4: Linear mixed-effects models using Eigen and S4. R package version 1.1-7. http://CRAN.R-project.org/package=lme4>.
- Bothell, D. et al. (2009). ACT-R Tutorial. Distributed with ACT-R 6.0 version 1.3 r766. http://act-r.psy.cmu.edu/software>.
- Cabeza, R., Locantore, J. K., & Anderson, N. D. (2003). Lateralization of prefrontal activity during episodic memory retrieval: Evidence for the production-monitoring hypothesis. *Journal of Cognitive Neuroscience*, *15*, 249–259.
- Chein, J. M., & Schneider, W. (2005). Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. *Cognitive Brain Research*, 25, 607–623.
- Compton, B. J., & Logan, G. D. (1991). The transition from algorithm to retrieval in memory-based theories of automaticity. *Memory & Cognition*, 19, 151–158.
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Computers and Biomedical Research, 29(3), 162–173.
- Curran, T. (2000). Brain potentials of recollection and familiarity. Memory & Cognition, 28, 923-938.
- Delaney, P. F., Reder, L. M., Staszewski, J. J., & Ritter, F. E. (1998). The strategy-specific nature of improvement: The power law applies by strategy within task. *Psychological Science*, 9, 1–7.
- Fitts, P. M., & Posner, M. I. (1967). Human performance.Oxford: Brooks/Cole.
- Friston, K. (2011). A short history of SPM. Statistical parametric mapping: The analysis of functional brain images: The analysis of functional brain images (p. 3).

Glover, G. H. (1999). Deconvolution of impulse response in event-related bold fMRI 1. Neuroimage, 9, 416-429.

- Haider, H., & Frensch, P. A. (2002). Why aggregated learning follows the power law of practice when individual learning does not: Comment on Rickard (1997, 1999), Delaney et al. (1998), and Palmeri (1999). Journal of Experimental Psychology: Learning, Memory, and Cognition, 28, 392–406.
- Heathcote, A., Brown, S., & Mewhort, D. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, 7, 185–207.
- Kelly, A. C., & Garavan, H. (2005). Human functional neuroimaging of brain changes associated with practice. Cerebral Cortex, 15, 1089–1102.
- Lee, H. S., Betts, S., & Anderson, J. R. (2015). Learning problem-solving rules as search through a hypothesis space. *Cognitive Science*.
- Lee, H., Voss, M. W., Prakash, R. S., Boot, W. R., Vo, L. T., Basak, C., ... Kramer, A. F. (2012). Videogame training strategy-induced change in brain function during a complex visuomotor task. *Behavioural Brain Research*, 232, 348–357.

Logan, G. D. (1988). Toward an instance theory of automatization. Psychological Review, 95, 492.

- Logan, G. D. (2002). An instance theory of attention and memory. Psychological Review, 109, 376.
- Lövdén, M., Wenger, E., Mårtensson, J., Lindenberger, U., & Bäckman, L. (2013). Structural brain plasticity in adult learning and development. Neuroscience & Biobehavioral Reviews, 37, 2296–2310.
- Luu, P., Tucker, D. M., & Stripling, R. (2007). Neural mechanisms for learning actions in context. Brain Research, 1179, 89–105.
- Myung, I. J., Kim, C., & Pitt, M. A. (2000). Toward an explanation of the power law artifact: Insights from response surface analysis. *Memory & Cognition*, 28, 832–840.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice (Vol. 1)Hillsdale, NJ: Erlbaum.
- Patel, R., Spreng, R. N., & Turner, G. R. (2013). Functional brain changes following cognitive and motor skills training a quantitative meta-analysis. *Neurorehabilitation and Neural Repair*, 27, 187–199.
- Petersen, S. E., Van Mier, H., Fiez, J. A., & Raichle, M. E. (1998). The effects of practice on the functional anatomy of task performance. *Proceedings of the National Academy of Sciences*, 95, 853–860.
- Poldrack, R. A. (2000). Imaging brain plasticity: Conceptual and methodological issues A theoretical review. *Neuroimage*, *12*, 1–13.
- Reas, E. T., & Brewer, J. B. (2013). Retrieval search and strength evoke dissociable brain activity during episodic memory recall. Journal of Cognitive Neuroscience, 25, 219–233.

- Reber, P. J. (2013). The neural basis of implicit learning and memory: A review of neuropsychological and neuroimaging research. *Neuropsychologia*, 51, 2026–2042.
- Rickard, T. C. (1997). Bending the power law: A CMPL theory of strategy shifts and the automatization of cognitive skills. *Journal of Experimental Psychology: General, 126, 288.*
- Rickard, T. C. (2004). Strategy execution in cognitive skill learning: An item-level test of candidate models. Journal of Experimental Psychology: Learning, Memory, and Cognition, 30, 65.
- Rubin, D. C., Hinton, S., & Wenzel, A. (1999). The precise time course of retention. Journal of Experimental Psychology: Learning, Memory, and Cognition, 25, 1161.
- Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: A bayesian hierarchical approach. *Psychological Review*, 120, 39.
- Schneider, W., & Chein, J. M. (2003). Controlled and automatic processing: Behavior, theory, and biological mechanisms. *Cognitive Science*, 27, 525–559.
- Taatgen, N. A., & Anderson, J. R. (2002). Why do children learn to say "broke" a model of learning the past tense without feedback. *Cognition*, 86, 123–155.
- Taatgen, N. A., & Lee, F. J. (2003). Production compilation: A simple mechanism to model complex skill acquisition. *Human* Factors: The Journal of the Human Factors and Ergonomics Society, 45, 61–76.
- Tang, Y.-Y., & Posner, M. I. (2014). Training brain networks and states. Trends in Cognitive Sciences, 18, 345-350.
- Tenison, C., & Anderson, J. (2015). Modeling the distinct phases of skill acquisition. Journal of Experimental Psychology: Learning, Memory, and Cognition.
- Zbrodoff, N. J. (1995). Why is 9+7 harder than 2+3? Strength and interference as explanations of the problem-size effect. Memory & Cognition, 23, 689–700.