

# Scaling Effective Learning Strategies: Retrieval Practice and Long-Term Knowledge Retention in MOOCs

Dan Davis<sup>1\*</sup>, René F. Kizilcec<sup>2</sup>, Claudia Hauff<sup>1</sup>, Geert-Jan Houben<sup>1</sup>

## Abstract

Scientific evidence for effective learning strategies is primarily derived from studies conducted in controlled laboratory or classroom settings with homogeneous populations. Large-scale online learning environments such as MOOCs provide an opportunity to evaluate the efficacy of these learning strategies in an informal learning context with a diverse learner population. *Retrieval practice*, a learning strategy focused on actively recalling information, is recognized as one of the most effective learning strategies in the literature. Here, we evaluate the extent to which retrieval practice facilitates long-term knowledge retention among MOOC learners using an instructional intervention. We conducted a pre-registered randomized encouragement experiment to test how retrieval practice affects learning outcomes in a MOOC. We observed no effect on learning outcomes and high levels of treatment non-compliance. This finding supports the view that evidence-based strategies do not necessarily work "out of the box" in new learning contexts and require context-aware research on effective implementation approaches. In addition, we conducted a series of exploratory studies into long-term recall of knowledge acquired in MOOCs. Surprisingly, both passing and non-passing learners scored similarly on a knowledge post-test, retaining approximately two-thirds of what they learned over the long term.

## Notes for Practice

- We develop and test a system that prompts students with previous questions to encourage retrieval practice in a MOOC.
- Retrieval practice showed no effect on learning outcomes or engagement.
- Learners who passed the course scored similarly on a knowledge post-test as do learners who didn't pass.
- Both passing and non-passing learners retain approximately two-thirds of the knowledge from the course over a two-week testing delay.

## Keywords

Retrieval Practice, Testing Effect, Experiment, Knowledge Retention, Learning Outcomes

**Submitted:** June 1, 2018 — **Accepted:** — — **Published:** —

<sup>1</sup> Delft University of Technology, Delft, the Netherlands. email: (d.j.davis, c.hauff, g.j.p.m.houben) @tudelft.nl

<sup>2</sup> Cornell University, Ithaca, NY, USA. email: kizilcec@cornell.edu

## 1. Introduction

A substantial body of research in the learning sciences has identified techniques and strategies that can improve learning outcomes, but translating these findings into practice continues to be challenging. An initial challenge is access to reliable information about which strategies work, which is the focus of scientific evidence aggregation hubs like the What Works Clearinghouse<sup>1</sup>. A subsequent challenge is knowing how to successfully implement and scale up an evidence-based strategy in potentially different learning contexts. This article addresses the latter challenge, which is becoming increasingly complex with the rise in popularity of online learning for diverse learner populations. We focus on a particular strategy that has been established as one of the most effective strategies to facilitate learning. *Retrieval practice*, also known as the testing effect, is the process of reinforcing prior knowledge by actively and repeatedly recalling relevant information. This strategy is more effective in facilitating robust learning—the committing of information to long-term memory (Koedinger et al., 2012)—than passively revisiting the same information, for example by going over notes or book chapters (Adesope et al., 2017; Clark & Mayer, 2016; Roediger & Butler, 2011; Henry L Roediger & Karpicke, 2016; Lindsey et al., 2014; Karpicke & Roediger, 2008; Karpicke & Blunt, 2011).

<sup>1</sup><https://ies.ed.gov/ncee/wwc/>

Given the wealth of scientific evidence on the benefits of retrieval practice (cf. Section 2), we test the extent to which the testing effect can be leveraged in one of today's most popular digital learning settings: Massive Open Online Courses (MOOCs). Research into both MOOC platforms and MOOC learners' behavior has found learners to take a distinctly linear trajectory (Davis, Chen, Hauff, & Houben, 2016; Wen & Rosé, 2014; Geigle & Zhai, 2017) through course content. Many learners take the path of least resistance towards earning a passing grade (Zhao et al., 2017) which does not involve any back-tracking or revisiting of previous course units—counter to a regularly-spaced retrieval practice routine.

Although MOOC platforms are not designed to encourage retrieval practice, prior work suggests that MOOC learners with high Self-Regulated Learning (SRL) skills tend to engage in retrieval practice of their own volition (Kizilcec, Pérez-Sanagustín, & Maldonado, 2017). These learners strategically seek out previous course materials to hone and maintain their new skills and knowledge. However, these learners are the exception, not the norm. The vast majority of MOOC learners are not self-directed autodidacts who engage in such effective learning behavior without additional support. This motivated us to create the Adaptive Retrieval Practice System (ARPS for short), a tool that encourages retrieval practice by automatically delivering quiz questions from previously studied course units, personalized to each learner. The system is automatic in that the questions appear without any required action from the learner and personalized in that questions are selected based on a learner's current progress in the course. We designed ARPS in accordance with the best practices in encouraging retrieval practice based on prior studies finding it to be effective (Adesope et al., 2017). We deployed ARPS in an edX MOOC (GeoscienceX) in a randomized controlled trial with more than 500 learners assigned to either a treatment (ARPS) or a control group (no ARPS but static retrieval practice recommendations).

Based on the data we collected in this experiment, we investigate the benefits of retrieval practice in MOOCs guided by the following research questions:

- RQ1** How does an adaptive retrieval practice intervention affect learners' course engagement, learning outcomes, and self-regulation compared to generic recommendations of effective study strategies in a MOOC?
- RQ2** How does a push-based retrieval practice intervention (which requires learners to act) change learners' retrieval practice behavior in a MOOC?

We also conduct exploratory analyses on the relationship between learner demographics (age, prior education levels, and country) and their engagement with ARPS:

- RQ3** Is there demographic variation in learners' engagement in retrieval practice?

In addition to collecting behavioral and performance data inside of the course, we invited learners to complete a survey two weeks after the course had ended. This self-report data enabled us to address the following two research questions:

- RQ4** To what extent is long-term knowledge retention facilitated in a MOOC?
- RQ5** What is learners' experience using ARPS to engage in retrieval practice?

The remainder of the article is organized as follows: In Section 2 we review the literature on retrieval practice, spacing, and knowledge retention. In Section 3 we introduce ARPS in detail. Section 4 discusses the methodology and our study design. Results are presented in Section 5, and Section 6 offers a discussion of the main findings, limitations to consider, and implications for future research.

## 2. Related Work

We now review prior research in the areas of retrieval practice, spaced vs. massed practice, and long-term knowledge retention, which will serve as a basis for the design of the current study.

### 2.1 Retrieval Practice

Decades of prior research on the effectiveness of different learning strategies has found retrieval practice to be effective at supporting long-term knowledge retention (Adesope et al., 2017; Clark & Mayer, 2016; Roediger & Butler, 2011; Henry L Roediger & Karpicke, 2016; Lindsey et al., 2014; Karpicke & Roediger, 2008; Karpicke & Blunt, 2011; Custers, 2010; Lindsey et al., 2014). However, how to effectively support retrieval practice in digital learning environments has not yet been thoroughly examined. The vast majority of prior work was conducted in “offline” learning environments, including university laboratory settings. Common methods of retrieval practice include (i) taking practice tests, (ii) making & using flashcards, (iii) making one's own questions about the topic, (iv) writing and/or drawing everything one can on a topic from memory, and (v) creating a concept map of a topic from memory<sup>2</sup>. Given the online medium and affordances of the edX platform, we here adopt

<sup>2</sup><http://www.learningscientists.org/blog/2016/6/23-1>

the practice test approach.

A key line of inquiry in the present study is the extent to which effective retrieval practice can be fostered in the MOOC setting. Retrieval practice has been evaluated in a variety of contexts in prior research. The vast majority (82%) of studies considered in the recent meta-review by Adesope et al. (2017) were conducted in laboratory settings, compared to 11% in physical class settings (7% did not reported the setting). The authors observed consistently large and significant effect sizes across both contexts. Furthermore, only 22% of those studies employed a randomized trial design, and, again, among those studies that did randomly assign participants to conditions, there was a consistently large and significant effect size. This trend of significant and large effect sizes holds for studies of post-secondary learners as well as studies specifically measuring learning outcomes (as opposed to engagement or retention). The design of our experiment is informed by these findings in order to maximize the similarity between our implementation and those which have found retrieval practice to be highly effective.

Critically evaluating the efficacy of retrieval practice in large-scale digital learning environments promises to advance theory by developing a deeper understanding of how retrieval practice can be effectively deployed/designed in a digital context as well as in a highly heterogeneous population that is embodied by MOOC learners. Prior research on Intelligent Tutoring Systems (ITS) has evaluated the benefits of teaching effective learning strategies and the effectiveness of various prompting strategies (Roll et al., 2011; VanLehn et al., 2007; Bouchet et al., 2016). The experimental results presented by Bouchet et al. (2016), for example, found that a scaffolding condition in which fewer SRL supports in the ITS were delivered each week was beneficial to learners even after the supports were removed completely. However, ITS studies are typically conducted in controlled lab settings with a considerably more homogeneous population than in MOOCs (Davis et al., 2018; Kizilcec & Brooks, 2017), so it is imperative that we scale such interventions up to the masses to gain a robust understanding of their heterogeneous effects.

Adesope et al. (2017) conducted the most recent meta-analysis of retrieval practice. They evaluated the efficacy of retrieval practice compared to other learning strategies such as re-reading or re-watching, the impact of different problem types in retrieval practice, the role of feedback, context, and students' education level. The effect of retrieval practice is strong enough overall for the authors to recommend that frequent, low-stakes quizzes be integrated into learning environments so that learners can assess knowledge gaps and seek improvement (Adesope et al., 2017). They also found that multiple choice problems not only require low levels of cognitive effort, they were the most effective type of retrieval practice problem in terms of learning outcomes compared to short answer questions. And while certainly a boon to learners (the majority of studies in the review endorse its effectiveness), feedback is actually not required or integral to effective retrieval practice. From studies that did incorporate feedback, the authors found that delayed feedback is more effective in lab studies, whereas immediate feedback is best in classroom settings.

Roediger & Butler (2011) also offer a synthesis of published findings on retrieval practice (from studies which were also carried out in controlled, homogeneous settings). From the studies reviewed, the authors offer five key points on retrieval practice for promoting long-term knowledge: (i) retrieval practice is superior to reading for long-term retention, (ii) repeated testing is more effective than a single test, (iii) providing feedback is ideal but not required, (iv) benefits are greatest when there is lag time between learning and practicing/retrieving, and (v) retrieval practice increases the likelihood of learning transfer—the application of learned knowledge in a new context (Roediger & Butler, 2011).


Consistent with the above findings, Johnson & Mayer (2009) evaluated the effectiveness of retrieval practice in a digital learning environment focused on lecture videos. In the study, learners who answered test questions after lecture videos—pertaining to topics covered in the videos—outperformed learners who merely re-watched the video lectures in terms of both long-term knowledge retention (ranging from one week to six months) and learning transfer.

A related body of research on spaced versus massed practice has found that a higher quantity of short, regularly-spaced study sessions is more effective than a few long, massed sessions (Clark & Mayer, 2016). There is overlap in the research on retrieval practice and that on spaced practice. Combining the insights from both literatures, an optimal study strategy in terms of achieving long-term knowledge retention is one of a regularly spaced retrieval practice routine (Miyamoto et al., 2015; Cepeda et al., 2008; Clark & Mayer, 2016). In fact, Miyamoto et al. (2015) provide evidence for this recommendation in the MOOC setting. They analyzed learners' log data and found that learners who tend to practice effective spacing without guidance or intervention are more likely to pass the course relative to those learners who do not engage in spacing.

## 2.2 Knowledge Retention

Scientific evaluation of the human long-term memory began at the end of the 19<sup>th</sup> century, leading to the earliest model of human memory loss/maintenance: the Ebbinghaus curve of forgetting (Ebbinghaus, 1885). The curve begins at time 0 with 100% knowledge uptake with a steep drop-off in the first 60 minutes to nine hours, followed by a small drop from nine hours to 31 days. Figure 1 is our own illustration/depiction of Ebbinghaus' memory curve based on his findings.

Custers (2010) conducted a review of long-term retention research and found considerable evidence in support of the Ebbinghaus curve in terms of shape—large losses in short-term retention (from days to weeks) which level off for longer intervals (months to years)—but not always in terms of scale. The result of their meta-analysis shows that university students




**Figure 1.** Rendition of Ebbinghaus' memory curve. Recall rate as a function of lag time in days.

typically lose one third of their knowledge after one year, even among the highest-achieving students. While the *shape* of Ebbinghaus' forgetting curve (cf. Figure 1) seems to hold across contexts, the actual proportion of knowledge remembered, varies.

Considering the effect on retrieval practice on long-term retention, Lindsey et al. (2014) conducted a similar study to the present research in a traditional classroom setting and found that their personalized, regularly spaced retrieval practice routine led to higher scores on a cumulative exam immediately after the course as well as a cumulative exam administered one month after the course. In their control condition (massed study practice), learners scored just over 50% on the exam, whereas those exposed to the retrieval practice system scored 60% on average. For the control group, this marked an 18.1% forgetting rate, compared to 15.7% for those with retrieval practice. They also found that the positive effect of retrieval practice was amplified with the passing of time.

It is worth noting that forgetting can be viewed as an adaptive behavior. Specifically, forgetting liberates the memory of outdated, unused information to create space for new, immediately relevant memories and knowledge (Richards & Frankland, 2017). Retrieval works against this natural, adaptive tendency to forget in that by regularly reactivating and revisiting knowledge, the brain recognizes its frequent use, labels it as such, and stores it in long-term memory accordingly so that it is readily accessible going forward.

Duolingo<sup>3</sup>, a popular language learning platform with hundreds of thousands of daily users, has developed their own forgetting curve to model the “half-life” of knowledge—their operates on a much smaller time scale, with a 0% probability of remembering after seven days. Based on the retrieval practice and spacing effect literature, they also developed a support system to improve learners' memory. Findings show that their support system, tuned to the “half-life regression model” of a learner's knowledge, significantly improves learners' memory (Streeter, 2015).



**Figure 2.** From Settles & Meeder (2016): Example student word-learning traces over 30 days: check marks indicate successful retrieval attempts, x's indicate unsuccessful attempts, and the dashed line shows the predicted rate of forgetting.

Researchers have used this model of a learner's knowledge at Duolingo to explore the effect of retrieval practice on learning and knowledge retention. Settles & Meeder (2016) in Figure 2 illustrate this effect by showing how retrieval interrupts the forgetting curve and how each successful retrieval event makes the ensuing forgetting curve less steep. The design of the

<sup>3</sup>[www.duolingo.com](http://www.duolingo.com)




Duolingo platform is centered around retrieval events as most learning activities constitute knowledge assessments. This is not the case in most MOOCs, where lecture videos constitute the primary instructional content. Thus, understanding how to foster effective retrieval practice in MOOCs warrants further investigation.

### 3. Adaptive Retrieval Practice System

ARPS is a client-server application (written in JavaScript/node.js) that provides *automated*, *scalable* and *personalized* retrieval practice questions to MOOC learners. We developed ARPS for use within the edX platform, taking advantage of the RAW HTML input affordance. This allows course creators to build custom interfaces within the platform that render along with the standard edX content (such as videos, quizzes, etc.). We designed the user interface in accordance with previous approaches to encouraging retrieval practice (Adesope et al., 2017; Clark & Mayer, 2016; Roediger & Butler, 2011; Henry L Roediger & Karpicke, 2016; Lindsey et al., 2014; Karpicke & Roediger, 2008; Karpicke & Blunt, 2011; Custers, 2010; Lindsey et al., 2014).

As illustrated in Figure 3, the ARPS back-end keeps track of the content a MOOC learner has already been exposed to through client-side sensor code that logs a learner's progress through the course and transmits it to the back-end. Once the back-end receives a request from the ARPS front-end (a piece of JavaScript running in a learner's edX environment on pages designated to show retrieval practice questions), it determines which question to deliver to a learner at a given time based on that learner's previous behavior in the course by randomly selecting from a pool of questions only pertaining to content the learner has already been exposed to<sup>4</sup>. Each question is pushed to the learner in the form of a qCard, an example of which is shown in Figure 4. These qCards appear to the learner as a pop-up within the browser tab. We log all qCard interactions: whether it was ignored or attempted, the correctness of the attempt, and the duration of the interaction.



**Figure 3.** ARPS System Architecture. Solid outlines indicate back-end components; dashed lines indicate front-end components.

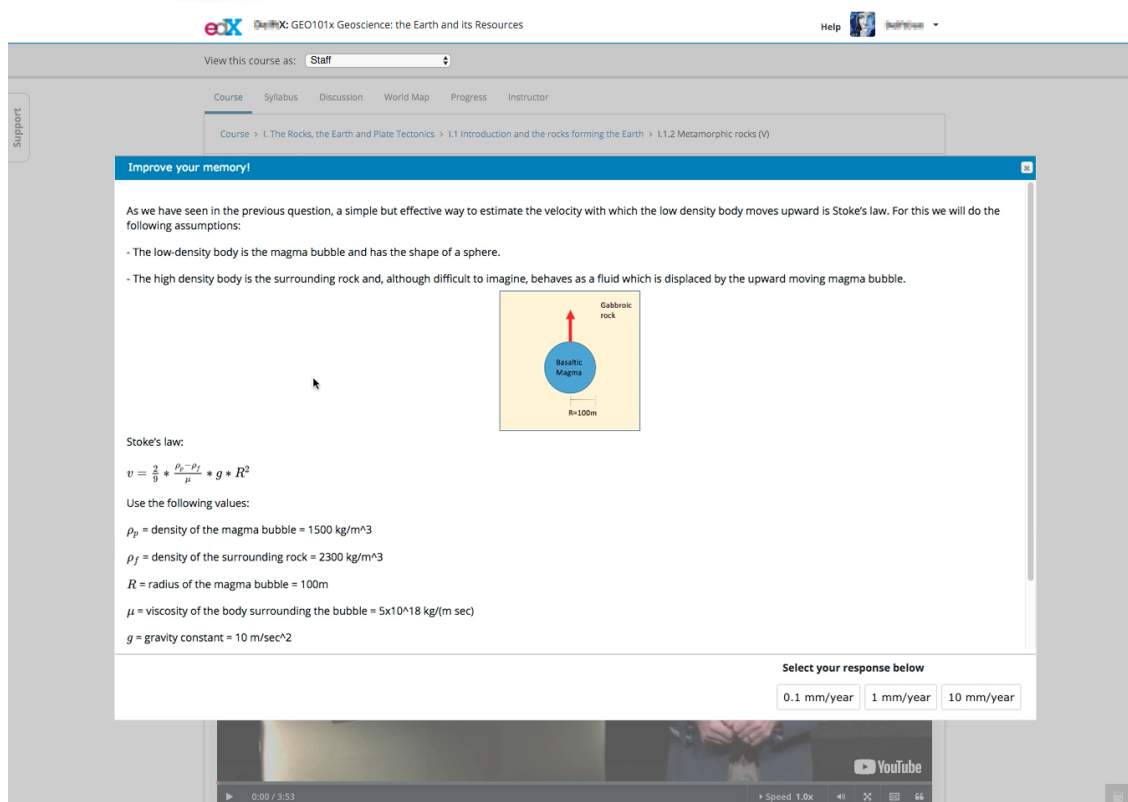
In every course page (see Figure 5 for the standard appearance and organization of a course page in edX) we inserted our code (a combination of JavaScript, HTML, and CSS) in a RAW HTML component. This code was either **sensor code** (included in every page) or a **qCard generator** (included in some pages). The sensor code did not result in any visible/rendered component; it tracked behavior and logged learners' progress through the course. We inserted the qCard generator code on each lecture video page.

The learners' course progress and all of their interactions with the qCards were stored in our server, to enable real-time querying of each learner's quiz history. By default, edX provides user logs in 24 hour intervals only, making it impossible to employ edX's default logs for ARPS. Our system design thus effectively acts as a RESTful API for the edX platform in allowing the real-time logging of information and delivery of personalized content.

In contrast to previous interventions in MOOCs (Rosen et al., 2017; van der Zee et al., 2018; Davis, Chen, Van der Zee, et al., 2016; Davis et al., 2017; Kizilcec & Cohen, 2017; Kizilcec, Saltarelli, et al., 2017; Yeomans & Reich, 2017; Gamage et al., 2017), we push questions to learners instead of requiring the learner to seek the questions out (e.g. by pressing a button). We adopted this design in order to allow learners to readily engage with the intervention with minimal interruption to the course experience. This design also addresses the issue of treatment noncompliance that has arisen in past research (Davis, Chen, Van der Zee, et al., 2016; Davis et al., 2017). In the case of Multiple Choice (MC) questions (example problem text in Figure 4), the entire interaction requires just a single click: the learner selects their chosen response and, if correct, receives positive feedback (a ✓ mark accompanied by encouraging text), and the qCard disappears. For incorrect responses, a learner receives a negative feedback symbol (an x alongside text encouraging them to make another attempt) which disappears after 4 seconds and redirects the learner to the original question to reattempt the problem<sup>5</sup>. Learners are not informed of the correct response after incorrect attempts in order to give them a chance to actively reflect and reevaluate the problem. The qCard can also be closed without attempting the question.

<sup>4</sup>Note that learners could have engaged with a course week's content without having engaged with all of its quiz questions.

<sup>5</sup>A video demonstration of the system in action is available at <https://youtu.be/25ckrrMzCr4>.



edX GEO101x Geoscience: the Earth and its Resources

View this course as: Staff


Course Syllabus Discussion World Map Progress Instructor

Course > 1. The Rocks, the Earth and Plate Tectonics > 1.1 Introduction and the rocks forming the Earth > 1.1.2 Metamorphic rocks (V)

**Improve your memory!**

As we have seen in the previous question, a simple but effective way to estimate the velocity with which the low density body moves upward is Stoke's law. For this we will do the following assumptions:

- The low-density body is the magma bubble and has the shape of a sphere.
- The high density body is the surrounding rock and, although difficult to imagine, behaves as a fluid which is displaced by the upward moving magma bubble.



Stoke's law:

$$v = \frac{2}{9} * \frac{\rho_b - \rho_f}{\mu} * g * R^2$$

Use the following values:

$\rho_b$  = density of the magma bubble = 1500 kg/m<sup>3</sup>

$\rho_f$  = density of the surrounding rock = 2300 kg/m<sup>3</sup>

$R$  = radius of the magma bubble = 100m

$\mu$  = viscosity of the body surrounding the bubble = 5x10<sup>-18</sup> kg/(m sec)

$g$  = gravity constant = 10 m/sec<sup>2</sup>

Select your response below

☐ 0.1 mm/year ☐ 1 mm/year ☐ 10 mm/year

**Figure 4.** Example qCard in the GeoscienceX course. The main body of the qCard contains the question text, and the bar at the bottom contains the MC answer buttons. The grey “x” at the top right corner closes the window and dismisses the problem.

Apart from MC questions, we also enabled one other question type<sup>6</sup> to appear in qCards: the Numeric Input (NI) type. These questions require the learner to calculate a solution and enter the answer. While requiring more effort than a single click response, we included this question type to allow for a comparison between the two. Even though MC questions were most common in GeoscienceX, this is not the case across all MOOCs.

The following is an example of an easy (less than 5% of incorrect responses) multiple choice question in GeoscienceX:

*A body with a low density, surrounded by material with a higher density, will move upwards due to buoyancy (negative density difference). We analyze the situation of a basaltic magma generated at a depth of 10 km and surrounded by gabbroic rocks. Will the magma move downward, remain where it is or move upward?*

The following is an example of a difficult (5% correct response rate) numerical input question in GeoscienceX:

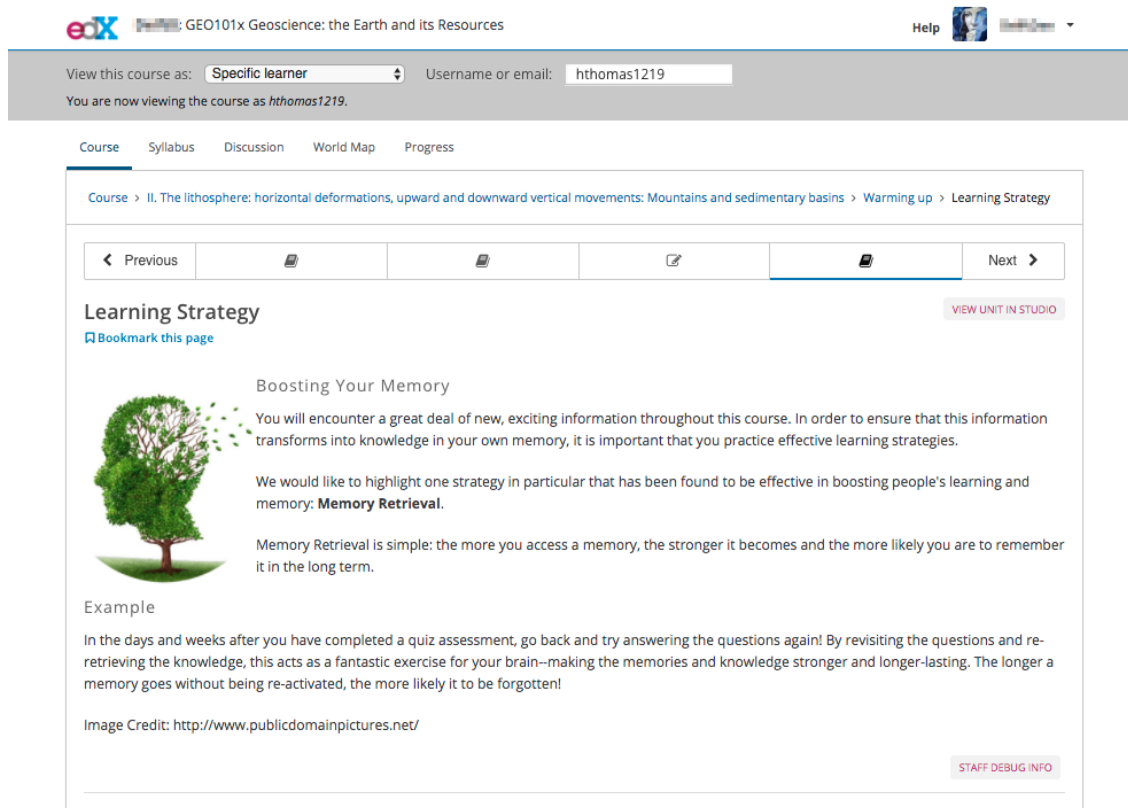
*Suppose an earthquake occurred at a depth of 10 kilometers from the surface that released enough energy for a P-wave to travel through the center of the Earth to the other side. This is for the sake of the exercise, because in reality sound waves tend to travel along the boundaries and not directly through the Earth as depicted. Assume the indicated pathway and the given thicknesses and velocities. How many seconds does it take for the seismic P-wave to reach the observatory on the other side of the Earth?*

These two sample problems are representative of the set of questions delivered in ARPS in that they prompt the learner to mathematically simulate an outcome of a situation based on given circumstances.

## 4. Study Design

In this section, we describe the design and context of our empirical study.

<sup>6</sup> Additional question types that are supported by the edX platform can easily be added to ARPS; in this paper we focus exclusively on MC and NI questions as those are the most common question types in the MOOC we deployed ARPS in.



edX GEO101x Geoscience: the Earth and its Resources

View this course as: **Specific learner** Username or email: **hthomas1219**

You are now viewing the course as **hthomas1219**.

Course Syllabus Discussion World Map Progress

Course > II. The lithosphere: horizontal deformations, upward and downward vertical movements: Mountains and sedimentary basins > Warming up > Learning Strategy

< Previous [icon] [icon] [icon] [icon] Next >

### Learning Strategy

[Bookmark this page](#) [VIEW UNIT IN STUDIO](#)

#### Boosting Your Memory

You will encounter a great deal of new, exciting information throughout this course. In order to ensure that this information transforms into knowledge in your own memory, it is important that you practice effective learning strategies.

We would like to highlight one strategy in particular that has been found to be effective in boosting people's learning and memory: **Memory Retrieval**.

Memory Retrieval is simple: the more you access a memory, the stronger it becomes and the more likely you are to remember it in the long term.

#### Example

In the days and weeks after you have completed a quiz assessment, go back and try answering the questions again! By revisiting the questions and re-retrieving the knowledge, this acts as a fantastic exercise for your brain—making the memories and knowledge stronger and longer-lasting. The longer a memory goes without being re-activated, the more likely it is to be forgotten!

Image Credit: <http://www.publicdomainpictures.net/>

[STAFF DEBUG INFO](#)

**Figure 5.** Page shown to learners in the control condition at the beginning of each course week describing how to practice an effective memory retrieval routine.

## 4.1 Participants

A total of 2,324 learners enrolled in the MOOC *Geoscience: the Earth and its Resources* (or GeoscienceX), which was offered on the edX platform between May 23, 2017 and July 26, 2017<sup>7</sup>. The course consists of 56 lecture videos and 217 graded quiz questions. We manually determined 132 of those 217 questions to be suitable for use with qCards (multi-step problems were excluded so that each qCard could be answered independently), 112 were MC and 20 were NI questions.

Based on self-reported demographic information (available for 1,962 learners), 35% of participants were women and the median age was 27. This course drew learners from a wide range of educational backgrounds: 24% held at least a high school diploma, 7% an Associate's degree, 42% a Bachelor's degree, 24% a Master's degree, and 3% a PhD. Learners were not provided any incentive beyond earning a course certificate for participating in the study.

We define the study sample as the 1,047 enrolled learners who entered the course *at least once*: 524 assigned to the control condition and 523 to the treatment condition. The random assignment was based on each learner's numeric edX ID (modulo operation).

We collected a stratified sample for the post-course survey & knowledge retention quiz. We sent an invitation email to all 102 learners who attempted a problem within the ARPS system—9 completed the survey, a 8.8% response rate—and the 150 learners with the highest final grades in the control condition—11 completed the survey, a 7.3% response rate.

## 4.2 Procedure

This study was designed as a randomized controlled trial. Upon enrolling in GeoscienceX, learners were randomly assigned to one of two conditions for the duration of the course:

- **Control condition:** A lesson on effective study habits was added to the weekly introduction section. The lesson explained the benefits of retrieval practice and offered an example of how to apply it (Figure 5).
- **Treatment condition:** ARPS was added to the course to deliver quiz questions from past weeks. The same weekly lesson on study habits as in the control condition was provided to help learners understand the value of the tool. In addition, information on how ARPS works and that responses to the qCard do not count towards learners' final grade was provided. The qCards were delivered before each of the 49 course lecture videos (from Weeks 2–6) across the six

<sup>7</sup>The study was pre-registered at <https://osf.io/4py2h/>.

course weeks. A button at the bottom of each lecture video page enabled learners to receive a new qCard on demand after the initial one to keep practicing.

To assess how well learners retained their knowledge from the course, we sent a post-course survey to a stratified sample of learners in the course (as defined before) two weeks after the course had ended. The survey contained a random selection of ten assessment questions from GeoscienceX. Learners in the treatment condition additionally received eight questions about their experience with ARPS. We evaluated the results of this post-course assessment with respect to differences between the two cohorts in long-term knowledge retention.

### 4.3 Measures

To compare the behavior of learners in the control and treatment conditions, we consider the following in-course events:

- **Final grade** (a score between 0 and 100);
- **Course completion** (binary indicator: pass, no-pass);
- Course activities:
  - Video interactions (play, pause, fast-forward, rewind, scrub);
  - Quiz submissions (number of submissions, correctness);
  - Discussion forum posts;
  - Duration of time in course;
- ARPS interactions:
  - Duration of total qCard appearance;
  - Response submissions (with correctness);
  - qCard interactions (respond, close window).

We collected the following data in the post-course survey:

- Course survey data
  - **Post-exam quiz score** (between 0-100%);
  - Learner intentions (e.g., to complete or just audit);
  - Prior education level (highest degree achieved).

The three variables printed in bold serve as primary outcome variables in our study for the following reasons: (i) a learner's final grade is the best available indicator of their performance in the course in terms of their short-term mastery of the materials, and (ii) the post-exam quiz score measures how well learners retained the knowledge two weeks after the end of the course.

## 5. Results

This section presents results from five sets of analyses we conducted: (i) estimating the causal effect of the retrieval practice intervention (RQ1), (ii) examining how learners interacted with ARPS (RQ2), (iii) exploring demographic variation in the use of ARPS (RQ3), (iv) modeling how learners' knowledge evolves and decays over time (RQ4), and, (v) understanding learners' experience with ARPS from a qualitative angle using survey responses (RQ5).

**Table 1.** Summary statistics for the mean value of the measures listed in Section 4.3 for analyses including all learners in both conditions who logged at least one session in the platform. The differences observed between Control and Treatment groups are not statistically significant.

| Group     | N=  | <b>Final Grade</b> | <b>Passing Rate</b> | <b># Video Interactions</b> | <b># Quiz Submissions</b> | <b># Forum Posts</b> | <b>Time in Course</b> | <b>Time with qCards</b> | <b>qCards Seen</b> |
|-----------|-----|--------------------|---------------------|-----------------------------|---------------------------|----------------------|-----------------------|-------------------------|--------------------|
| Control   | 524 | 9%                 | 8%                  | 6.52                        | 34.77                     | 0.26                 | 4h47m                 | –                       | –                  |
| Treatment | 523 | 8%                 | 7%                  | 5.83                        | 30.88                     | 0.29                 | 3h40m                 | 23m9s                   | 7.71               |

### 5.1 Effect of Encouraging Retrieval Practice

The goal of the randomized experiment is to estimate the *causal effect* of retrieval practice (RQ1). By comparing learners in the control and treatment group, we can estimate the effect of the encouragement to engage in retrieval practice with ARPS with the following pre-registered confirmatory analyses. Table 1 provides summary statistics for learners in the two experimental conditions, showing no significant differences between conditions, according to tests of equal or given proportions (all  $p > 0.05$ ). However, many learners who were encouraged did not engage in retrieval practice, which is a form of treatment noncompliance. Specifically, of the 523 learners assigned to the treatment, only 102 interacted at least once with a qCard (i.e. complied with the treatment). For this reason, in order to estimate the effect of retrieval practice itself, we also analyze the experiment as an encouragement design. Due to the small sample size and low compliance rate, however, we adjusted our analytic approach.



Specifically, we analyze the experiment as an encouragement design beyond estimating average treatment effects, and we do not apply the pre-registered sample exclusion criteria because they could inadvertently bias the causal inference.

The primary outcome measure is the final course grade, which determines certificate eligibility (the passing threshold is 60%). Table 2 contains summary statistics for grade and certification outcomes in the control group and the treatment group, overall and separately for treatment compliers and noncompliers. First, we estimate the Intent-to-treat Effect (ITT), which is the difference in average outcomes between the treatment and control groups. We find that the ITT is not significant for certification (log odds ratio =  $-0.215, z = -0.920, p = 0.357$ ), getting a non-zero grade (logOR =  $0.143, z = 1.08, p = 0.280$ ), and the continuous grade itself (Kruskal-Wallis  $\chi^2_{df=1} = 0.592, p = 0.442$ ).

Next, we use an instrumental variable approach (two-stage least squares) to estimate the effect of retrieval practice for those who used it (i.e. a Local Average Treatment Effect, or LATE) (Angrist et al., 1996). For a binary instrument  $Z$ , outcome  $Y$ , and compliance indicator  $G$ , we can compute the Wald estimator:

$$\beta^{IV} = \frac{E(Y|Z=1) - E(Y|Z=0)}{E(G|Z=1) - E(G|Z=0)} \quad (1)$$

The LATE is not significant for certification ( $\beta^{IV} = -0.078, z = -0.893, p = 0.371$ ), getting a non-zero grade ( $\beta^{IV} = 0.160, z = 1.11, p = 0.267$ ), and the continuous grade itself ( $\beta^{IV} = -0.066, z = -0.889, p = 0.374$ ).

Finally, we conduct a per-protocol analysis, which considers only those learners who completed the protocol for their allocated treatment. We estimate the per-protocol effect as the difference in average outcomes between treatment compliers and control compliers, which is the entire control group by design. We find large differences in terms of certification (logOR =  $1.74, z = 6.66, p < 0.001$ ), getting a non-zero grade (logOR =  $2.00, z = 7.94, p < 0.001$ ), and the continuous grade itself (Kruskal-Wallis  $\chi^2_{df=1} = 99, p < 0.001$ ). However, the per-protocol analysis does not yield causal estimates due to self-selection bias; in particular, we compare everyone in the control group with those highly motivated learners who comply in the treatment group. In fact, treatment compliance is strongly correlated with receiving a higher grade (Spearman's  $r = 0.56, p < 0.001$ ).

**Table 2.** Course outcomes in the control and treatment group. Compliers are learners who complied with the intervention; non-compliers are those who ignored the intervention.

| Condition | Subset      | N   | Non-Zero Grade | Passing Rate | Grade Quantiles |
|-----------|-------------|-----|----------------|--------------|-----------------|
| Control   | All         | 524 | 31%            | 8%           | [0, 0, 2]       |
| Treatment | All         | 523 | 34%            | 7%           | [0, 0, 2]       |
| Treatment | Complier    | 102 | 76%            | 34%          | [2, 19, 74]     |
| Treatment | Noncomplier | 421 | 23%            | 0.2%         | [0, 0, 0]       |

In addition to estimating effects based on the final course grade, the pre-registration also specifies a number of process-level analyses (RQ2). In particular, we hypothesized that learners who receive the treatment would exhibit increased self-regulatory behavior in terms of (i) revisiting previous course content such as lecture videos, (ii) self-monitoring by checking their personal progress page, and (iii) generally persisting longer in the course. No evidence in support of the hypothesized behavior was found, neither in terms of the ITT (Kruskal-Wallis  $\chi^2_{df=1}$ -values  $< 0.68$ ,  $p$ -values  $> 0.41$ ) nor in terms of the LATE ( $|z|$ -values  $< 0.98$ ,  $p$ -values  $> 0.32$ ). Focusing on learners in the treatment group, we also hypothesized that learners who attempt qCards at a higher rate would learn more and score higher on regular course assessments, which is supported by the data (Spearman's  $r = 0.42, p < 0.001$ ). In summary and in contrast to previous studies on the topic (Adesope et al., 2017; Clark & Mayer, 2016; Roediger & Butler, 2011; Henry L Roediger & Karpicke, 2016; Lindsey et al., 2014; Karpicke & Roediger, 2008; Karpicke & Blunt, 2011) we find that:

A causal analysis yields no evidence that encouraging retrieval practice raised learning, performance, or self-regulatory outcomes in this course.

We also observe a selection effect into using ARPS among highly motivated learners in the treatment group. Among those learners, increased engagement with qCards was associated with higher grades, though this pattern could be due to self-selection (e.g., more committed learners both attempt more qCards and put more effort into assessments). To better understand how different groups of learners used ARPS and performed on subsequent learning assessments, we conduct a series of exploratory analyses.

## 5.2 Navigational Behavior Following Retrieval Cues

With the following exploratory analyses we now aim to understand the extent to which engagement with the qCards (i.e. compliance with the experimental manipulation) led to changes in learners' navigational behavior through the course (RQ2). Specifically, we examine whether or not there were marked changes in behavior immediately following the appearance of a qCard—as the qCard is from a prior unit in the course, we expected this to stimulate learners to skip backwards in the course to revisit old content to refresh their memory.

Once again, we divide learners in the treatment condition into compliers and noncompliers. For the two groups, we then isolate every page load immediately following the appearance of a qCard (i.e., where did learners navigate to after being exposed to a qCard?), measured by the proportion of times where learners either continue forwards or track backwards to an earlier point in the course. There is no significant difference between compliers and noncompliers in terms of their forward or back-seeking behavior immediately following exposure to a qCard ( $\chi^2 = 1.5$ ,  $p = 0.22$ ). Both groups navigated backwards after qCards 8–10% of the time. This indicates that qCards did not have a significant effect on learners' navigation and did not motivate them to revisit previous course topics to retrieve relevant information.

## 5.3 Engagement with Retrieval Cues

As prior work has found significant demographic variation in learner behavior and achievement in MOOCs (Kizilcec, Saltarelli, et al., 2017; Kizilcec & Halawa, 2015; Davis et al., 2017), we explore such variation in the use of ARPS (RQ3).

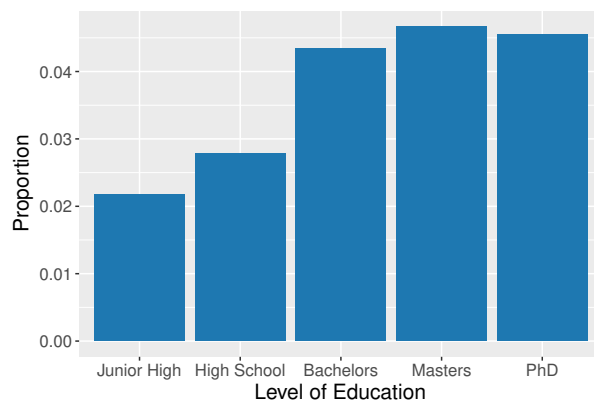
### 5.3.1 Demographic Analysis

We here consider learner characteristics that have been identified as important sources of variation: age, prior education level, and the human development index of the country in which a learner is located (HDI: the United Nations' combined measure of a country's economic, health, and educational standing (Sen, 1994)). The following analyses explore whether the retrieval practice system was similarly accessible and helpful to all groups of learners, as we intended it to be in the design phase.

We find that learners with a higher level of education are more likely to interact with ARPS (i.e. intervention compliers), as illustrated in Figure 6. This suggests that the design of the system may be less accessible or usable for learners with lower levels of education. Next, we test for age differences between compliers and noncompliers. The average complier is 38 while the average noncomplier is only 32 years old, a statistically significant difference ( $p = 0.003$ ,  $t = -2.99$ ). Finally, we examine whether learners from more developed countries would be more or less likely to engage with qCards. We find that compliers came from significantly more developed countries (higher HDI) than noncompliers, even when adjusting for prior education levels (ANCOVA  $F = 4.93$ ,  $p = 0.027$ ).

One goal of open education (and ARPS) is to make education more accessible to historically disadvantaged populations. We find that ARPS is not used equally by different demographic groups. Younger and less educated learners from developing countries are less likely to use the system and thus miss an opportunity to benefit from the testing effect facilitated by ARPS.


Learners' interaction with retrieval prompts varies systematically with demographic characteristics, namely age, country, and prior education level.



**Figure 6.** Compliance vs. prior education level: proportion of all learners of each level of prior education to have complied with the intervention (engaged with at least one qCard).

### 5.3.2 Question-level Analysis


Figure 7 illustrates learners' responses for every question delivered by ARPS, which indicates which questions learners tended to struggle with (or ignore). The figure reveals that the choice to attempt or ignore a qCard is strongly associated with a learner's eventual passing or failing of the course. Moreover, it shows a steady decrease in learner engagement over time, not only among non-passing learners, but also among those who earned a certificate. Thus, attrition in MOOCs is not limited to those who do not pass the course; even the highest-achieving learners show a tendency of slowing down after the first week or two (also observed in Zhao et al. (2017); Kizilcec et al. (2013)).



**Figure 7.** A question-by-question breakdown of every learner interaction with the qCards. The top two figures represent the behavior of passing learners—the upper image shows the number of learners being served that question, and the lower shows the proportion of ignored, incorrect, and correct responses—and the bottom two show that of non-passing learners. Questions are shown in order of appearance in the course (left to right), and the solid vertical line indicates a change in course week (from 1–5). Best viewed in color.

From Figure 8, we observe that passing and non-passing learners do not differ in their rate of giving incorrect responses, which would indicate misconceptions or a lack of understanding the materials. Instead, they differ in their choice to ignore the problems altogether. When removing the instances of ignored qCards and focusing only on attempted problems (right-hand side of Table 3), we observe a significant albeit small difference (6 percentage points,  $\chi^2 = 9.63$ ,  $p = 0.002$ ) between the proportion of correct or incorrect responses between passing and non-passing learners (cf. Table 3). In other words, passing and non-passing learners both perform about the same on these quiz problems—and yet, with no discernible difference in their assessed knowledge, only some go on to earn a passing grade and course certificate. Figure 8a shows that there is a single learner who answered more than 300 qCard questions. This is the only learner who took advantage of the ARPS functionality that allowed learners to continue answering qCard questions after engaging with the original (automatically appearing) one, which enabled a Duolingo-like approach where a learner can engage with an extended sequence of assessments (or retrieval cues) from past course weeks<sup>8</sup>. This functionality was previously not available in edX, as quiz questions can only be attempted a finite number of times—most commonly once or twice.

<sup>8</sup>We found that this learner followed a pattern of regular weekly engagement in the course over the span of 1.5 months and that their high level of ARPS usage was not a result of gaming the system or a single massed practice session, but rather an indication of a highly disciplined pattern of spaced learning.



**Figure 8.** Each bar corresponds to one learner. Only one learner took advantage of the “infinite quizzing” capability by frequently using the “Generate new qCard” button. Best viewed in color.

**Table 3.** ARPS problem response (left) and correctness (right). Significant differences ( $p < 0.001$ ) between passing and non-passing learners in each column are indicated with †.

|             | Attempted <sup>†</sup> | Ignored <sup>†</sup> | Correct | Incorrect |
|-------------|------------------------|----------------------|---------|-----------|
| Non-passing | 47%                    | 53%                  | 76%     | 24%       |
| Passing     | 73%                    | 27%                  | 82%     | 18%       |

### 5.3.3 First Question Response

To explore the predictive power of a learner’s choice to either attempt or ignore the qCards, we next analyze each learner’s *first* interaction with a qCard. Figure 9a shows the passing rate of learners segmented according to their first interaction with a qCard. Learners who attempted the first qCard instead of ignoring it had a 47% chance of passing the course. In contrast, learners who ignored their first qCard only had a 14% chance of passing. Figure 9b additionally illustrates the relationship between the result of the first qCard attempt and course completion. There were notably few learners who responded incorrectly, but their chance of passing the course was still relatively high at 33% compared to those who simply ignored the qCard.

While Figure 9a shows passing rates conditional on *attempting* a problem, Figure 9b shows passing rates conditional on the result of the first attempt. The latter adds no predictive power; getting the first question correct is still associated with about a 50% chance of passing the course.

To evaluate whether the response of a learner’s second qCard problem adds any predictive value, we replicate the analysis shown in Figure 9 for the responses to the first two qCards delivered to each learner. No difference in predictive power was observed by considering the second consecutive response—learners who answered their first two consecutive qCards correctly had a 53% chance of earning a passing grade.

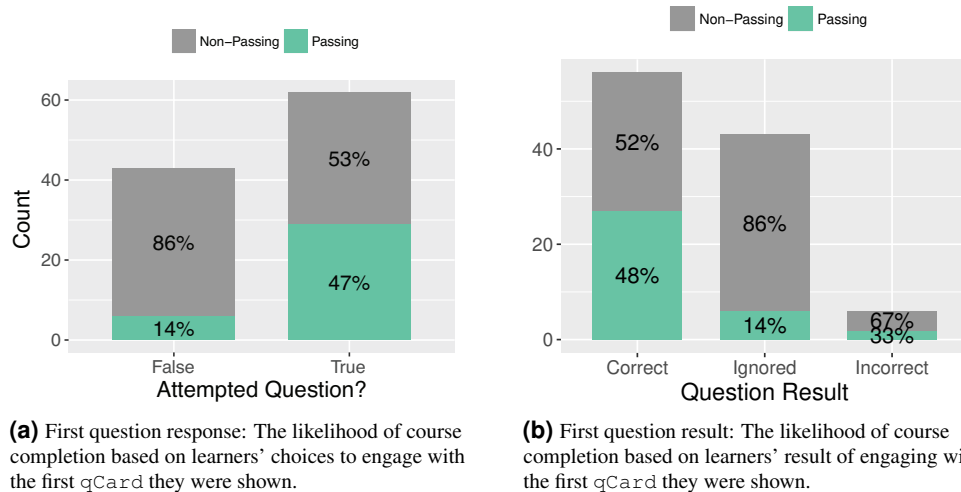
We conclude that initial adoption of ARPS appears to depend partly on learners’ motivation to complete the course.

### 5.3.4 Response Duration


We next explore how much time learners spent interacting with qCards and how the amount of time they spent predicts the outcome of the interaction. Figure 10 shows the proportion of correct, incorrect, and ignored responses as a function of time elapsed with a qCard. We find that the decision to ignore the qCard happened quickly, with a median duration of 7 seconds (from the time the qCard appeared to the time the learner clicked the “x” button to close it). When learners did attempt to answer the question, the amount of time they spent was not predictive of the correctness of their response, with a median duration for correct and incorrect responses of 18 and 16 seconds, respectively.

From the question-level, first question response, and response duration analyses, we conclude:

There is no significant difference in assessed knowledge between passing and non-passing learners; the key difference lies in a learner’s willingness to engage with the retrieval practice questions.



**Figure 9.** The predictive value of learners' engagement with the first qCard to which they were exposed. "True" indicates both correct or incorrect responses, and "False" indicates the qCard was ignored. Best viewed in color.



**Figure 10.** Result likelihood vs. time: Kernel Density Estimation plot showing the relationship between time with the qCard visible and the result of a learner's response. The median time for each result is indicated with the dashed vertical line. Best viewed in color.

## 5.4 Modeling Knowledge Over Time

One of the contributions of ARPS is the data set that it generates: by tracking learners' responses to these periodic, formative, and ungraded questions throughout the entire course, we have a longitudinal account of learners' evolving knowledge state throughout the entire process of instruction. In this section, we explore how learners' knowledge (as measured by performance with the qCards) deteriorates over time (RQ4).

Figure 11 shows the cumulative week-by-week performance of both passing and non-passing learners. As qCards could only be delivered with questions coming from *prior* course weeks, the x-axis begins with Week 2, where only questions from Week 1 were delivered. This continues up to Week 6 where questions from Weeks 1–5 could be delivered.

Figure 11a illustrates the forgetting curve for passing learners in GeoscienceX. We observe a statistically significant decrease in performance between Weeks 2 and 6 (the correct response rate drops from 67% to 49%;  $\chi^2 = 32.8$ ,  $p < 0.001$ ). While the proportion of ignored responses remains steadily low, the proportion of correct responses drops by 18 percentage points (nearly identical to the forgetting rate found in Lindsey et al. (2014)). The rate of incorrect responses increased from 4% to 25% ( $\chi^2 = 87.8$ ,  $p < 0.001$ ).

Figure 11b illustrates the forgetting curve for non-passing learners. We observe that the choice to ignore qCards was common throughout the duration of the course, with a slight increase in the later weeks. We also observe a significant decrease in correct response rates for non-passing learners ( $\chi^2 = 15.7$ ,  $p < 0.001$ ). However, unlike passing learners who exhibited a significant increase in incorrect responses, there is no significant change for non-passing learners. The change is in the rate of ignored responses, which increases from 47% in Week 2 to 61% in Week 6.

We identify two main contributing factors to this decline in performance over time. First, the amount of assessed content increases each week; in Week 6 there are five course weeks worth of content to be assessed, whereas in Week 2 there is only content from Week 1 being assessed. Second, people simply forget more with the passing of time (Richards & Frankland,



2017); each passing week moves the learner temporally farther away from when the content was initially learned and therefore, without regular retrieval, increases the likelihood of this knowledge being marked as non-essential and thus forgotten.

We next explore the relationship between testing delay and learners' memory and performance on qCards. In Figure 12, the x-axis represents the difference between a learner's current week and the week from which the qCard came. For example, if a learner was watching a lecture video in Week 5 and the qCard delivered was a question from Week 2, that would be a difference of three. While Figure 11 shows how the amount of content covered/assessed is related to performance, Figure 12 illustrates how the delay in testing relates to performance.

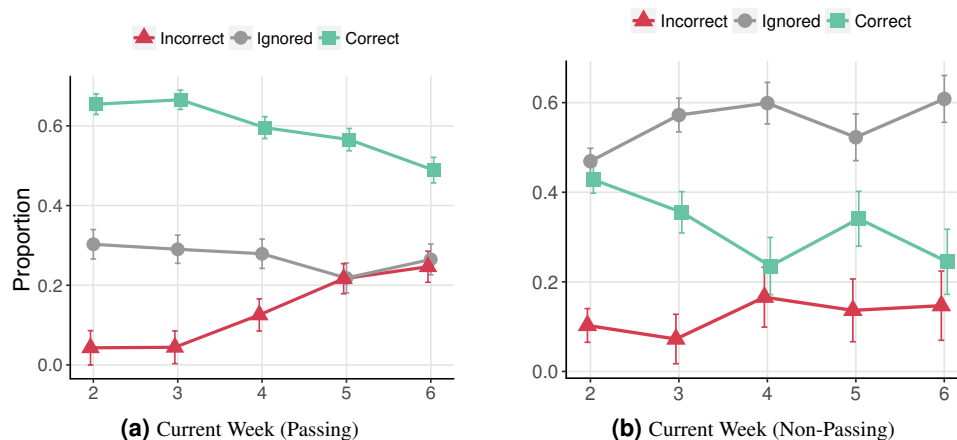
We observe very similar trends for both passing and non-passing learners. For passing learners there is a drop in the correct response rate from 1 Week Elapsed to 5 Weeks Elapsed (65% to 42%,  $\chi^2 = 23.6$ ,  $p < 0.001$ ). Similarly, there is an increase in the incorrect response rate (8% to 21%,  $\chi^2 = 17.5$ ,  $p < 0.001$ ). The increase in ignored question frequency is not significant for passing learners, though it is large and significant for non-passing learners: between 1 Week Elapsed and 5 Weeks Elapsed, ignored questions increased from 50% to 72% ( $\chi^2 = 4.9$ ,  $p = 0.025$ ). Overall, for non-passing learners, we observe increased ignoring, decreased correct problem attempt rates, and steady incorrect problem attempt rates.

This pattern shows that non-passing learners are able to recognize, attempt, and correctly answer qCard problems that are more proximate to their current stage in the course. This suggests a high level of self-efficacy, especially among non-passing learners; they are able to identify questions that they likely do not know the answer to and choose to ignore them.

Another promising finding from this analysis concerns learners' short-term knowledge retention. As illustrated by Figure 12, passing learners correctly answer 88% of problems that were attempted with 1 Week Elapsed. Non-passing learners also exhibit high recall levels with 79% correct (note that the required passing grade for the course was 60%). This, in tandem with the rapid rates of forgetting, suggests that there is room for a learning intervention to help learners commit new knowledge to long-term memory. We expected ARPS to accomplish this, but it was ultimately ineffective at promoting long-term knowledge retention.

From the above findings for assessed knowledge as a function of both time and course advancement, we conclude:

Knowledge recall deteriorates with the introduction of more course concepts/materials and the passing of time.

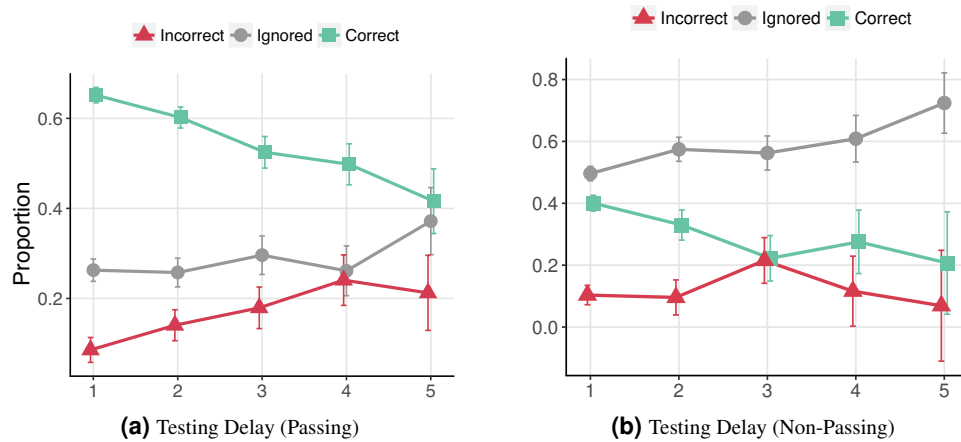


**Figure 11.** Week-by-week results of learners' interaction with the qCards. The x-axis represents the course week  $w$ , and the y-axis represents the proportion (%) of correct, incorrect, or ignored responses (with qCards showing queries from course weeks 1 to  $w - 1$ ). Error bars show standard error. Best viewed in color.

#### 5.4.1 Considering the Initial Knowledge State

The questions shown on qCards are sourced directly from the questions that appear earlier in the course as created by the course instructor. The results illustrated in Figures 11 & 12 show learner performance on qCards without taking into consideration learners' performance on the original graded question which appeared earlier in the course. We may not expect a learner to correctly answer a question later if they never got it right in the first place. With Figures 13–15, we explore the learners engagement and performance with the qCards conditional on their performance the first time they answered the question.

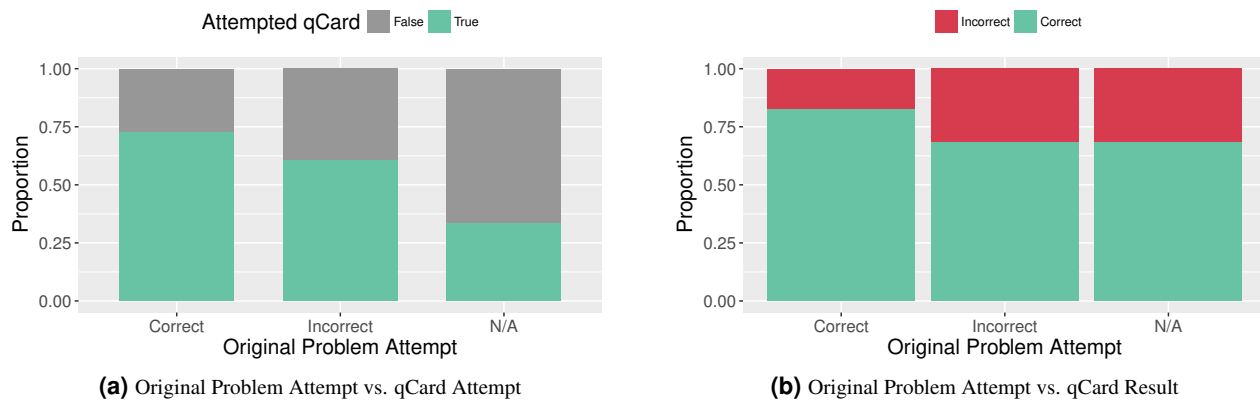
Figure 13a shows the likelihood of a learner attempting a qCard based on their performance on the original question in the course. When a qCard is shown to a learner who initially answered the question correctly, the learner will attempt the qCard



**Figure 12.** The x-axis represents the number of elapsed weeks between the course week where the topic was introduced and the course week in which the qCard was delivered (testing delay), and the y-axis represents the proportion of each response. Error bars show standard error. Best viewed in color.

73% of the time. However, if the question was initially answered incorrectly, only 61% reattempt it. And for questions that were ignored or never answered, those qCards were only attempted 34% of the time. These percentage values are all significantly different from each other ( $\chi^2 > 22$ ,  $p < 0.0001$ ). This indicates that a learner's performance on the original problem in the course is significant for their likelihood of engaging with a qCard showing that same question.

Figure 13b shows the likelihood of a learner either answering the qCard correctly, incorrectly, or ignoring it based on the learner's performance on the original quiz question in the course. When faced with qCards showing questions that the learner initially answered correctly, they get it correct again 61% of the time. In contrast, if the initial question was answered incorrectly or ignored, the chance of getting it right in a qCard later on drops to 42% and 23%, respectively. These percentages are significantly different from each other ( $\chi^2 > 40.0$ ,  $p < 0.0001$ ). This indicates that a learner's initial performance on the question significantly predicts their likelihood of answering the corresponding qCard correctly. Although learners do not perform well on the qCards in general, they demonstrate consistent knowledge of mastered concepts over time.

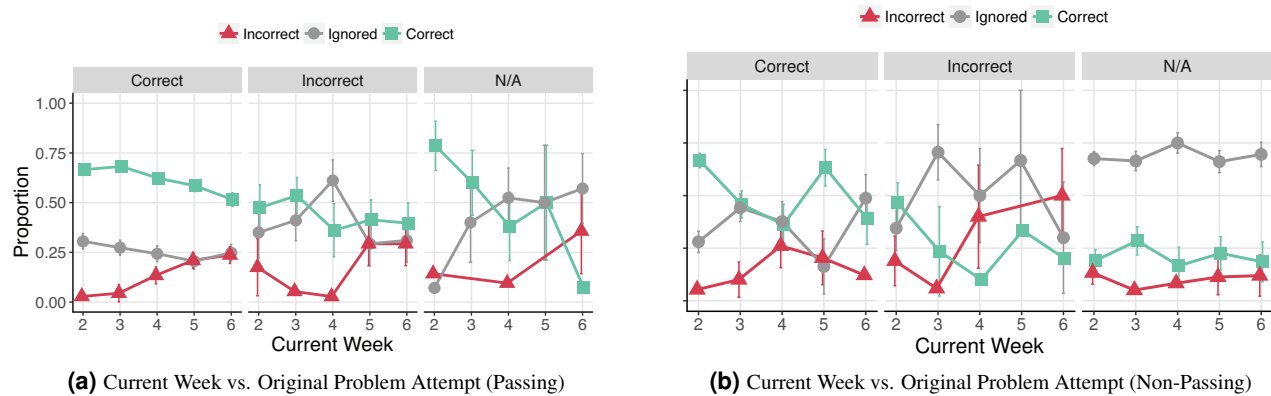


**Figure 13.** qCard performance vs. original problem attempt. N/A means not attempted. Best viewed in color.

We next replicate the analyses done in Figures 11 & 12, except this time we again take into consideration the learners' performance on the original problem in the course. These results are shown in Figures 14 & 15.

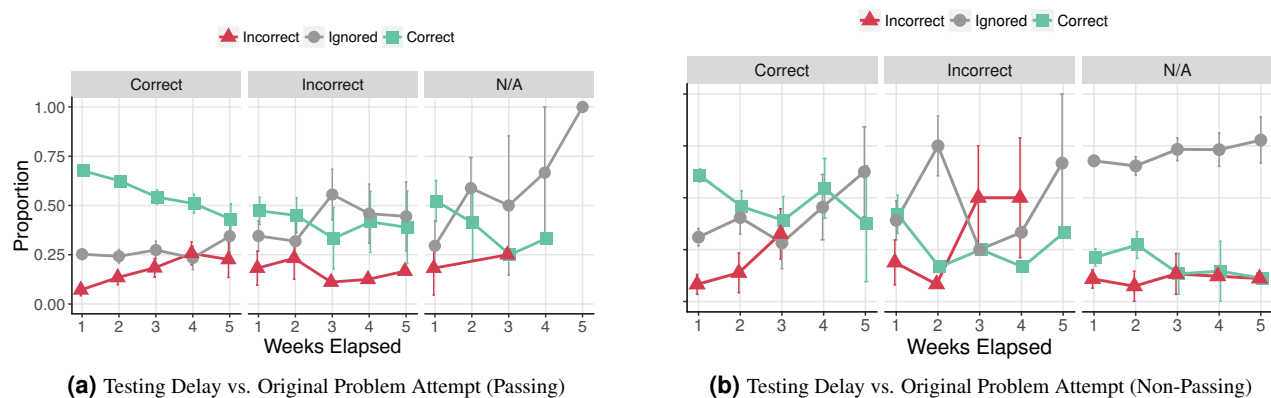
We find from Figure 14a that for passing learners, with each successive course week there is a steady decrease in the likelihood of learners answering qCards correctly, even when they had answered the original question correctly. We likewise observe a weekly increase in the proportion of incorrect answers and a slight decline in the proportion of ignored qCards. This trend is amplified for questions that were originally ignored; in Week 2, qCard questions which were originally ignored were answered correctly 79% of the time and this proportion drops to 7% in Week 6. This indicates the importance of merely attempting problems in the course, as the Week 6 rate of correct qCard responses to questions that were originally answered incorrectly is 40%—in other words, learners are 5.7 times more likely to answer a qCard correctly when they attempted the

original question and got it wrong than when they ignored/did not attempt the original question.



**Figure 14.** qCard performance vs. current course week and original problem attempt. Error bars show standard error. Best viewed in color.

Figure 15 shows the effect of testing delay and original problem performance on qCard performance. For passing learners, the most prominent trend to consider is the steady decline in performance over time for questions that were originally answered correctly. In conducting this follow-up analysis, we expected to find that taking learners' original performance on individual questions into consideration would reduce the apparent rate of forgetting, however the trend does hold. For qCard questions with a 1-week testing delay, passing learners answer them correctly 68% of the time; compare this to a 5-week testing delay where qCards are answered correctly only 39% of the time (a chi-Squared test indicates that this difference is statistically significant:  $p < 0.0001$ ,  $\chi^2 = 21.8$ ). With this result we identify more evidence of the trend of MOOC learners' knowledge (as ascertained by quiz question performance) to steadily deteriorate over time, even if they had initially exhibited some level of mastery.



**Figure 15.** qCard performance vs. testing delay and original problem attempt. Error bars show standard error. Best viewed in color.

#### 5.4.2 Long-Term Knowledge Retention

Well-spaced retrieval practice is known to improve long-term knowledge retention, which is typically evaluated by a post-test with some lag time between exposure to the learning materials and the assessment (Lindsey et al., 2014; Custers, 2010). As the GeoscienceX course used weekly quizzes for assessment, we took a random selection of ten quiz questions from throughout the six end-of-week quizzes and created a post-course knowledge assessment that was delivered to learners in a survey format two weeks after the course had ended.

We compare performance on the post-test between the two experimental conditions and found no significant difference (RQ4). The mean score for learners in the control and treatment conditions was 6.2 ( $SD = 1.9$ ) and 6.6 ( $SD = 1.8$ ), respectively, out of a possible 10 points ( $N = 20$ ,  $t(17.6) = -0.45$ ,  $p = 0.66$ ). Notably, our findings for long-term knowledge retention are consistent with prior literature (Lindsey et al., 2014; Custers, 2010) which finds that, regardless of experimental condition and whether or not a learner passed the course:

Approximately two thirds of course knowledge is retained over the long-term.

## 5.5 Learner Experience

To evaluate learners' experience with ARPS, we provided a stratified sample of learners a mix of rating and open-ended questions in a post-course survey (RQ5). Using an adapted version of the System Usability Scale (Brooke et al., 1996), we find that the average usability score of ARPS is 73.9 ( $\sigma = 12.2$ ) out of 100. According to Bangor et al. (2009), this is categorized as "acceptable usability" and the system's usability falls into the third quartile of SUS scores overall (Bangor et al., 2009). For a research prototype that was not developed as a production system, this is a notably positive rating. We do acknowledge, however, that there is room for improvement in the system's usability, and that this improvement could cause more learners to engage and benefit. We find evidence for survey response bias in terms of course performance. Learners who responded to the survey achieved higher grades in the course (median = 82%) than learners who did not respond (median = 4%; Kruskal-Wallis  $\chi^2 = 26.7, p < 0.0001$ ).

To gain deeper insight into learners' experience and find out which specific aspects of the system could be improved, we also offered learners the opportunity to describe their experience with ARPS in two open response questions. One prompted them to share which aspects of ARPS they found to be the most enjoyable and another asked about less desirable aspects of ARPS.

One learner explained how the type of problem delivered was a key factor in their use of ARPS, and how, for qCard questions they would prefer not to have to do any calculations:


*"It [would] be better if only conceptual questions [were] asked for [the] pop quiz, it's troublesome if calculation is required. If calculation is required, I would prefer that the options are equations so that we can choose the right equation without evaluating them."*

Other learners reported similar sentiments and also shared insights that indicate a heightened level of self-awareness induced by the qCards. Learners shared their perspectives talking about how the system helped "...remind me [of] things that I missed in the course" and how it gave them "the chance to see what I remembered and what I had learned." These anecdotes are encouraging as for these learners the system was able to encourage a deliberate activation of previously-learned concepts which may have otherwise been forgotten. One learner also reported that ARPS helped to "...reinforce my learning from previous chapters," and another learner reported: "It helps me to review important information in previous sections." So while the challenge of compliance and engagement remains the priority, it is encouraging that those who did choose to engage with ARPS did indeed experience and understand its benefits.

Upon seeing the learner feedback about how the problem type affected the learner's experience, we conducted a follow-up analysis to see if there was any indication that other learners felt the same way (as expressed through their interaction with ARPS). Figure 16 reveals that, indeed, this learner was not alone in their sentiment; we found that there was a 69% likelihood of learners attempting a MC qCard problem type compared to 41% attempting NI problems ( $p < 0.001$ ). Given that the question type (mostly evaluations of mathematical equations) is consistent across both problem types (MC and NI), we can conclude that these differences are indeed an effect of the problem type. This finding supports our initial design decision for a highly efficient interaction process—learners are far more likely to attempt a problem which only requires a single click selecting from a list of answers than one that requires two extra steps: generating an answer from scratch and then typing it out. Nevertheless, we cannot identify from the data which of these two steps contributes more to questions being ignored.

## 6. Discussion

In this study we have evaluated the extent to which the effectiveness of retrieval practice directly translates from traditional classroom and laboratory settings to MOOCs. We replicated the implementation of retrieval practice from past studies as closely as possible in this new context. The design therefore aligns with the best practices in the retrieval practice literature as outlined by Adesope et al. (2017). While Adesope et al. (2017) showed consistent positive effects on learning outcomes in randomized experiments in post-secondary physical classroom settings, no significant effects were found in the present study in the context of a MOOC. This indicates that traditional instructional design practices for encouraging retrieval practice in a traditional classroom or laboratory setting do not directly translate to this online context. We therefore highlight the need for further research to empirically investigate how evidence-based instructional/learning strategies can be successfully implemented in online learning contexts. In the case of retrieval practice, we find that directly translating what works in traditional contexts to MOOCs is especially ineffective for younger and less educated learners who joined the MOOC from developing countries. More research is needed to understand how to adapt learning strategies for novel learning contexts in a way that elicits the same cognitive benefits as have been observed in classroom and lab settings.



**Figure 16.** Response by problem type: Breakdown of qCard interaction results across the two problem types. Best viewed in color.

We evaluated an Adaptive Retrieval Practice System to address the emerging issue of supporting learning strategies at large scale (Davis et al., 2018) and to bridge retrieval practice theory into the digital learning space. As was the case with most MOOC studies in the past (Rosen et al., 2017; van der Zee et al., 2018; Davis, Chen, Van der Zee, et al., 2016; Davis et al., 2017; Kizilcec & Cohen, 2017; Kizilcec, Saltarelli, et al., 2017; Yeomans & Reich, 2017; Gamage et al., 2017), we found non-compliance to be a major limitation in our evaluation of the system and its effectiveness. Many learners did not engage with the intervention, which limits our ability to draw causal inferences about the effect of retrieval practice on learners' achievement and engagement in the course.

Despite the lack of causal findings, the data collected from ARPS allowed us to gain multiple insights into the online learning process as it pertains to the persistence and transience of knowledge gains. By examining learner performance and engagement with the retrieval prompts, we were able to track changes in knowledge retention for the same concept over time and with the introduction of new course materials. We summarize our key findings:

- We find no evidence that encouraging retrieval practice raised learning, performance, or self-regulatory outcomes.
- Demographic factors play a significant role in learner's compliance with retrieval cues, with older & more highly educated learners being most likely to engage.
- There is no significant difference in assessed knowledge between passing and non-passing learners; the key difference lies in a learner's willingness to engage with the retrieval practice questions.
- Learner quiz performance deteriorates with the introduction of more course concepts/materials and the passing of time.
- Approximately two thirds of course knowledge is retained over the long-term.

We observed an encouraging trend of learners' high levels of short- and medium-term knowledge retention, which is indicative of the early stages of learning. To what extent this newly gained knowledge is integrated into long-term memory warrants further research in the context of large online courses. Despite the null results from our causal analysis (Section 5.1), the wealth of evidence showing that retrieval practice is one of the most effective strategies to support knowledge retention makes this approach ripe for further investigation in online learning settings. To this end, researchers will need to find better ways of encouraging learners to use activities designed to support the learning process by designing interfaces that foster high levels of engagement.

## References

- Adesope, O. O., Trevisan, D. A., & Sundararajan, N. (2017). Rethinking the Use of Tests: A Meta-Analysis of Practice Testing. *Review of Educational Research*.
- Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434), 444–455.
- Bangor, A., Kortum, P., & Miller, J. (2009). Determining what individual sus scores mean: Adding an adjective rating scale. *Journal of usability studies*, 4(3), 114–123.
- Bouchet, F., Harley, J. M., & Azevedo, R. (2016). Can adaptive pedagogical agents' prompting strategies improve students' learning and self-regulation? In *International conference on intelligent tutoring systems* (pp. 368–374).



- Brooke, J., et al. (1996). Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194), 4–7.
- Cepeda, N. J., Vul, E., Rohrer, D., Wixted, J. T., & Pashler, H. (2008, November). Spacing Effects in Learning: A Temporal Ridgeline of Optimal Retention. *Psychological Science*, 19(11), 1095–1102.
- Clark, R. C., & Mayer, R. E. (2016). *E-learning and the science of instruction: Proven guidelines for consumers and designers of multimedia learning*. John Wiley & Sons.
- Custers, E. J. (2010). Long-term retention of basic science knowledge: a review study. *Advances in Health Sciences Education*, 15(1), 109–128.
- Davis, D., Chen, G., Hauff, C., & Houben, G.-J. (2016). Gauging mooc learners' adherence to the designed learning path. In *Proceedings of the 9th international conference on educational data mining* (pp. 54–61).
- Davis, D., Chen, G., Hauff, C., & Houben, G.-J. (2018). Activating learning at scale: A review of innovations in online learning strategies. *Computers & Education*.
- Davis, D., Chen, G., Van der Zee, T., Hauff, C., & Houben, G.-J. (2016). Retrieval practice and study planning in moocs: Exploring classroom-based self-regulated learning strategies at scale. In *European conference on technology enhanced learning* (pp. 57–71).
- Davis, D., Jivet, I., Kizilcec, R. F., Chen, G., Hauff, C., & Houben, G.-J. (2017). Follow the successful crowd: raising mooc completion rates through social comparison at scale. In *Proceedings of the seventh international learning analytics & knowledge conference* (pp. 454–463).
- Ebbinghaus, H. (1885). *Über das gedächtnis: untersuchungen zur experimentellen psychologie*. Duncker & Humblot.
- Gamage, D., Whiting, M., Rajapakshe, T., Thilakarathne, H., Perera, I., & Fernando, S. (2017). Improving Assessment on MOOCs Through Peer Identification and Aligned Incentives. In *the fourth (2017) acm conference* (pp. 315–318). New York, New York, USA: ACM Press.
- Geigle, C., & Zhai, C. (2017). Modeling student behavior with two-layer hidden markov models. *JEDM-Journal of Educational Data Mining*, 9(1), 1–24.
- Henry L Roediger, I., & Karpicke, J. D. (2016, June). The Power of Testing Memory: Basic Research and Implications for Educational Practice. *Perspectives on Psychological Science*, 1(3), 181–210.
- Johnson, C. I., & Mayer, R. E. (2009). A testing effect with multimedia learning. *Journal of Educational Psychology*, 101(3), 621.
- Karpicke, J. D., & Blunt, J. R. (2011). Retrieval practice produces more learning than elaborative studying with concept mapping. *Science*, 331(6018), 772–775.
- Karpicke, J. D., & Roediger, H. L. (2008). The critical importance of retrieval for learning. *science*, 319(5865), 966–968.
- Kizilcec, R. F., & Brooks, C. (2017). Diverse big data and randomized field experiments in moocs. *Handbook of Learning Analytics*, 211–222.
- Kizilcec, R. F., & Cohen, G. L. (2017). Eight-minute self-regulation intervention raises educational attainment at scale in individualist but not collectivist cultures. *Proceedings of the National Academy of Sciences*, 201611898.
- Kizilcec, R. F., & Halawa, S. (2015). Attrition and achievement gaps in online learning. In *Proceedings of the second (2015) acm conference on learning@ scale* (pp. 57–66).
- Kizilcec, R. F., Pérez-Sanagustín, M., & Maldonado, J. J. (2017). Self-regulated learning strategies predict learner behavior and goal attainment in massive open online courses. *Computers & education*, 104, 18–33.
- Kizilcec, R. F., Piech, C., & Schneider, E. (2013). Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *Proceedings of the third international conference on learning analytics and knowledge* (pp. 170–179).
- Kizilcec, R. F., Saltarelli, A. J., Reich, J., & Cohen, G. L. (2017). Closing global achievement gaps in moocs. *Science*, 355(6322), 251–252.

- Koedinger, K. R., Corbett, A. T., & Perfetti, C. (2012). The knowledge-learning-instruction framework: Bridging the science-practice chasm to enhance robust student learning. *Cognitive science*, 36(5), 757–798.
- Lindsey, R. V., Shroyer, J. D., Pashler, H., & Mozer, M. C. (2014). Improving students' long-term knowledge retention through personalized review. *Psychological science*, 25(3), 639–647.
- Miyamoto, Y. R., Coleman, C. A., Williams, J. J., Whitehill, J., Nesterko, S., & Reich, J. (2015, December). Beyond time-on-task: The relationship between spaced study and certification in MOOCs. *Journal of Learning Analytics*, 2(2), 47–69.
- Richards, B. A., & Frankland, P. W. (2017). The Persistence and Transience of Memory. *Neuron*, 94(6), 1071–1084.
- Roediger, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in cognitive sciences*, 15(1), 20–27.
- Roll, I., Aleven, V., McLaren, B. M., & Koedinger, K. R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learning and Instruction*, 21(2), 267–280.
- Rosen, Y., Rushkin, I., Ang, A., Federicks, C., Tingley, D., & Blink, M. J. (2017). Designing Adaptive Assessments in MOOCs. In *the fourth (2017) acm conference* (pp. 233–236). New York, New York, USA: ACM Press.
- Sen, A. (1994). Human development index: Methodology and measurement.
- Settles, B., & Meeder, B. (2016). A trainable spaced repetition model for language learning. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: Long papers)* (Vol. 1, pp. 1848–1858).
- Streeter, M. (2015). Mixture modeling of individual learning curves. *International Educational Data Mining Society*.
- van der Zee, T., Davis, D., Saab, N., Giesbers, B., Ginn, J., van der Sluis, F., ... Admiraal, W. (2018). Evaluating retrieval practice in a mooc: how writing and reading summaries of videos affects student learning. In *Proceedings of the 8th international conference on learning analytics and knowledge* (pp. 216–225).
- VanLehn, K., Graesser, A. C., Jackson, G. T., Jordan, P., Olney, A., & Rosé, C. P. (2007). When are tutorial dialogues more effective than reading? *Cognitive science*, 31(1), 3–62.
- Wen, M., & Rosé, C. P. (2014). Identifying latent study habits by mining learner behavior patterns in massive open online courses. In *Proceedings of the 23rd acm international conference on conference on information and knowledge management* (pp. 1983–1986).
- Yeomans, M., & Reich, J. (2017). Planning prompts increase and forecast course completion in massive open online courses. In *Proceedings of the seventh international learning analytics & knowledge conference* (pp. 464–473).
- Zhao, Y., Davis, D., Chen, G., Lofi, C., Hauff, C., & Houben, G.-J. (2017). Certificate achievement unlocked: How does mooc learners' behaviour change? In *Adjunct publication of the 25th conference on user modeling, adaptation and personalization* (pp. 83–88).