Thinking about doing: Representations of skill learning

Xiuyuan Zhang (flora.zhang@yale.edu)
Samuel D. McDougle (samuel.mcdougle@yale.edu)
Julia A. Leonard (julia.leonard@yale.edu)
Department of Psychology, Yale University

Abstract

Skill learning usually unfolds exponentially — we improve rapidly early on, and then performance levels off. However, we do not know whether people's *representations* of skill learning accurately reflect this fact. Here, we asked people to predict the learning trajectory for a novel visuomotor task, "Lollitoss." First, we established that skill learning unfolds exponentially on Lollitoss (Exp. 1). Across two experiments probing people's trial-by-trial predictions of learning in Lollitoss using direct performance (Exp. 2a) and likelihood estimates (Exp. 2b), we found that people accurately represent the learning curve as exponential. However, we also found systematic errors - people think individuals start out better, make less errors, and learn slower in the task than in reality. Taken together, we find that people are surprisingly accurate at representing the overall shape of learning, but misestimate certain features, like the rate of learning, which may potentially have downstream effects on self-directed learning.

Keywords: intuitive theories of learning; visuomotor learning; learning curve

How does a child go from never riding a bike before to confidently riding one all the way to school? Even if you have never ridden a bike, you probably have some intuition about how learning this skill might unfold over time. The child probably didn't hop on the bike and master it right away. Instead, she most likely made gradual improvements, from biking with training wheels, to riding just a bit by herself, to eventually riding all the way to school. Although tracking major skill learning milestones is intuitive, it is not clear if people have fine-grained, accurate representations of how learning typically unfolds *over time* (e.g., from minute to minute, hour to hour). Critically, our beliefs about how learning proceeds over time can influence our actions: If a child thinks their progress should be swift the first day that they try learning to ride a bike, but it turns out to be slow, then they may prematurely quit. Here, we ask whether people have granular and accurate representations of skill learning.

Decades of research have tracked how skill learning unfolds across a wide range of task domains, from simple motor skills, to complex routines, to perception and memory tasks(Thorndike, 1913; Newell & Rosenbloom, 1981; Heathcote, Brown, & Mewhort, 2000; Gallistel, Fairhurst, & Balsam, 2004). These studies reveal that performance gains, especially during the acquisition of motor skills, proceed according to a decelerating exponential (or power) function, just like riding a bike. That is, a naïve learner will experience a lot of initial improvement over a short period of time, then the amount of improvement per unit time will decrease as their performance approaches an asymptote (Heathcote et al., 2000; Krakauer, Hadjiosif, Xu, Wong, & Haith, 2019).

Although individuals can approach these asymptotes at different rates, the general shape of learning across individuals remains the same.

However, reasoning about exponential functions is notoriously difficult. In seminal work, Wagenaar and Sagaria (1975) showed that people tend to "linearize" observations that actually are generated by exponential functions. This is true not only when people have to extrapolate from a few data points presented numerically, but also when data is presented graphically (Wagenaar & Sagaria, 1975), or in contexts of visual storytelling (e.g., watching a pond fill up with duckweed and predicting the amount of time left before the whole pond is covered; Wagenaar & Timmers, 1978). Even in more consequential, real-world situations, like planning for long-term financial investments, people believe that the growth of their savings is linear when it is exponential, and they misestimate how much to invest (known as the "exponential growth bias", see Mckenzie & Liersch, 2011; Stango & Zinman, 2009). Although it has been well documented that people have trouble reasoning about exponential functions in a wide variety of domains (e.g., financial decisions in Mckenzie & Liersch, 2011, and Covid-19 growth forecasting in Lammers, Crusius, & Gast, 2020, Hutzler et al., 2021), we know surprisingly little about people's exponential reasoning in the domain of learning.

When it comes to intuitions about our own and others' learning, often known as "metacognition," past work has primarily focused on single item or single time point predictions, and not the shape of learning over time. One line of work focuses on people's judgements of learning specific items before a test (see Finn & Metcalfe, 2008, 2014; Richardson, Sheskin, & Keil, 2021). For instance, Finn and Metcalfe (2008) found that college students were more accurate at judging which specific cue-target word-pairs they would remember after repeated exposure. Importantly, this literature has shown that metacognition impacts actual learning — accurate representations of learning individual items leads people to hone in on areas that need the most attention (Thiede, Anderson, & Therriault, 2003; Son & Metcalfe, 2005; Metcalfe & Finn, 2008; Son & Sethi, 2010). Another line of research in metacognition concerns people's broad intuitions about learning particular subjects (see Keil, Lockhart, & Schlegel, 2010; Letourneau & Sobel, 2020; Lockhart, Goddu, & Keil, 2021). For example, to understand what children think learning means, Letourneau and Sobel (2020) asked four- to eight-year-old children open-ended questions, such as "Can you think of something that you have learned?"

and "How did you learn that?". However, prior work leaves open the question of whether people can accurately simulate the time course of learning over longer time scales. Understanding how people intuitively represent learning progress is important not only to informing theory related to representations of the mind, but also practices related to effectively guiding learners' efforts on the path to mastery.

Here, we ask whether people intuitively represent skill learning as it actually progresses - that is, exponentially. To compare people's conceptions of skill learning to actual learning, we first collected learning data from naïve participants playing a novel visuomotor learning task, "Lollitoss" (Exp. 1). These data established "ground-truth" exponential learning trajectories in Lollitoss. Next, we examined a separate group of participants' predictions of learning in Lollitoss on a trial-by-trial basis, with no prior practice on the task. In one experiment, participants simulated a precise prediction of others' behavior across trials (Exp. 2a) and in another experiment, participants specified a set of numerical likelihood estimations for possible outcomes across trials (Exp. 2b). We probed representations of learning in two disparate ways to validate findings across methods and reduce the chance that the way we queried learning biased participant's answers. Across all three experiments, we fit the actual and simulated learning curves with either exponential or linear functions, and compared the fit of these models. To more precisely examine how people's predictions of learning differed from actual learning, we compared the parameters (i.e., the starting point, slope, and asymptote) obtained from the best model fits of actual learning data to people's simulated learning data. Together, these studies are designed to probe people's intuitions (Exp. 2a & b) about the time course of skill learning.

Experiment 1

The goal of Experiment 1 is to establish the ground-truth learning trajectory of individuals on our novel visuomotor task. Based on nearly a century of work on the shape of skill learning curves (see Heathcote et al., 2000 for a summary), we hypothesized that learning would unfold following an exponential decay function - fast at first, then slowly thereafter. Exp. 1 was preregistered here.

Methods

Participants   We recruited 55 adult U.S. participants online through Prolific. We expected to observe learning in the task (i.e., minimization of errors over time), so we preregistered fitting a simplified linear model (error ~ trial_number) for each participant and excluding any participant who did not show a negative slope (i.e., no evidence for a trend of learning). Based on this criteria, we excluded five participants (final $n = 50$).

Procedure   Participants were introduced to a novel game - Lollitoss. The goal of Lollitoss is to try to get as many points as possible by 'tossing' lollipops that move back and forth
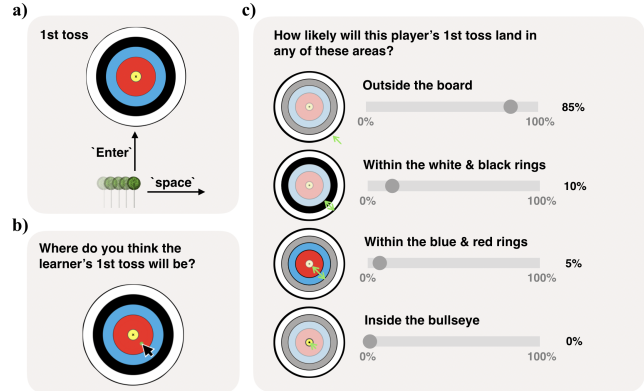

Figure 1: Schematics for (a) Exp.1 - actual learning, (b) Exp.2a - direct prediction, and (c) Exp.2b - likelihood estimation.

along the bottom of the screen towards the middle of a target (see Figure 1). To 'toss' the lollipops, participants first needed to stop the lollipop using the 'space bar' key, which moved at speed 970pixel/1000ms across the screen horizontally. Then, participants 'launched' the lollipop by pressing the 'Enter' key, with the amount of time they held down the key dictating the vertical distance that the lollipop traverses. We set the optimal interval to 1412 ms, with hold intervals less than 1079 ms or more than 1746 ms fully missing the target. If a player got the lollipop in the bullseye of the target (yellow circle, see Figure 1), they got 50 points. For each concentric ring outside of the bullseye (starting from red, ending at white) they got 30, 20, 10, and 5 points respectively, and 0 points for landing the lollipop outside of the target. After reading the instructions, participants had to pass two comprehension questions before proceeding to game play: "How many points will you get if you hit the lollipop on the bullseye?" and, "Which key will be used to toss the lollipop toward the board?".

Participants played Lollitoss on a 640pixel × 1,000pixel window. Each participant completed a total of 50 trials during 10 rounds of the game, with 5 trials per round. The lollipops always appeared first on an extreme side of the window before starting to move horizontally (randomized left and right start location). After each toss, participants saw their toss score and their total score. For a given trial, the Euclidean distance from the center of the tossed lollipop head to the center of the target was calculated and recorded to mark participant's deviation from the goal (i.e., their error for a given toss). By calculating participant's errors over time, we constructed individual learning curves.

Results and discussion

As expected, learning in Lollitoss was best characterized by an exponential decay function (Figure 2a). We fit each participant's learning data with an exponential decay function (error ~ a * exp(-b * trial_number) + c) and a linear function (error ~ b * trial_number + c), as a comparison. For

Figure 2: The left panel shows model estimations for individual learning rates fitted with exponential decay for each of the three experiments, and the right panel includes exponential vs. linear model comparisons for individual subject using AIC for the corresponding experiment. (a) & (b) Exp.1 - actual learning, (c) & (d) Exp.2a - direct prediction, and (e) & (f) Exp.2c - likelihood estimation. The average actual learning (a) and average predicted learning (c) & (e) are all presented with 95% bootstrapped CIs. Here, the y-axes show different scales to present clearly the learning trajectories from the three experiments.

all but one participant, learning data were fit by the exponential decay function. Focusing on these 49 participants, we evaluated model performance using the Akaike information criterion (AIC). As shown in Figure 2b, after computing the ΔAIC between the exponential and the linear models for each participant, we found that the exponential model outperformed the linear model in 44 out of 49 participants ($\Delta AIC < 0$). A paired Wilcoxon signed-rank test across the 49 participants revealed that the exponential model had significantly lower AIC scores (i.e., a better model-fit) than the linear model ($V = 41$, $p < .001$). Thus, in line with previous research on skill learning, our results show that the learning curve for Lollitoss is best described by an exponential decay function.

Experiment 2a

In Experiment 2a, we asked whether people can accurately simulate the trajectory of learning in Lollitoss. We probed people's intuitions by asking them to predict the precise locations of an imagined naïve learner's tosses at a few time points across 50 trials. Exp. 2a was preregistered here.

Methods

Participants   We recruited 54 adult U.S. participants online through Prolific. We expected participants to predict the occurrence of learning during the game, so we pre-registered fitting a linear model (predicted_error ~ trial_number) for each participant and excluding any participant data if their predicted errors over the total trials do not show a negative slope (i.e., no trend of expectation of learning). Based on this criteria, four participants were excluded (final $n = 50$).

Procedure   As in Exp. 1, participants were introduced to the instructions for Lollitoss. Additionally, to get a feel for the task, participants were given one opportunity to experience how the 'space' and 'Enter' keys work to move the lollipop without the presence of the target board. However, instead of performing any actual trials of Lollitoss (which would confound our results), participants were told that they would predict a naïve player's progress on the game. To proceed to the prediction phase, participants had to pass three comprehension questions, including the two comprehension questions in Exp. 1 and a third question "whose performance will you be predicting in this game?".

In the prediction phase, participants were asked to predict the landing location of the lollipop tosses for a first-time learner. To make a prediction, participants simply clicked on the screen (including inside and outside the target board) where they thought the lollipop would land on that trial. Participants completed seven rounds of predicting bins of trials, instead of the full 50 trials that participants played in Exp. 1, to reduce task redundancy and fatigue. In the first six rounds, participants were asked to predict three consecutive trials per round with a nine-trial interval between rounds (e.g., 1st, 2nd, 3rd trials, then 10th, 11th, and 12th trials, etc.). To ensure a matched ending trial between the prediction responses and the learning responses from Exp.1, participants also predicted the landing location of the last toss (i.e., the 50th trial) in the seventh round. The Euclidean distance between the reported center of the lollipop head and the center of the board was recorded, as in Exp.1.

Results & discussion

People intuitively represented the trajectory of learning over time as being exponential (Figure 2c). As in Exp. 1, we fit individual participant's predictions with an exponential decay function and a linear function for comparison. The exponential model provided the fits for 49 out of 50 participants. We computed the AIC score (see Figure. 2e for individual ΔAIC) for these 49 participants, where 37 out of 49 participants had a negative ΔAIC (Figure 2d). A paired Wilcoxon signed-rank test revealed that the exponential models had significantly lower AIC scores than the linear models ($V = 158$, $p < .001$).

These data suggest that people intuit that the shape of learning in Lollitoss as exponential rather than linear, accurately reflecting the general shape of the ground-truth learning trajectories on the task. However, it is possible that the way we queried participants could influence people's predictions. To find converging evidence about people's intuitions across methods, we ran a second experiment in which we asked participants to reason about learning in this task probabilistically.

Experiment 2b

In Experiment 2b, we tested people's trial-by-trial representation of the learning trajectory using a distinctly different measure: people predicted the likelihood of lollipops landing at specific regions on and outside the target board at different time points of learning. Exp. 2b was preregistered here.

Methods

Participants We recruited 63 adult U.S. participants online through Prolific. As in Exp. 2a, we expected participants to predict learning during the game, so we preregistered fitting a linear model (predicted_probability_for_bullseye ~ trial_number) to each participant's data and excluding any participant if their predicted probability for bullseye over 19 trials did not show a positive slope (i.e., no trend of expectation of learning). Thirteen participants were excluded based

on preregistered inclusion criteria (final $n = 50$).

Procedure To ensure that participants understood the goal of the task, we gave them two practice scenarios in which they predicted the likelihood of a ball landing in a certain color area on the floor below after being dropped. In one example, the area below the ball was equal colors yellow and green. In the other, the green to yellow ratio was 4:1. For each scenario, participants were asked two questions 'How likely will it (the ball) land inside the green/yellow region?' and recorded their responses using sliders (0%-100%). Participants who responded correctly to the two training scenarios continued to the instruction phase for Lollitoss. Participants were then introduced to Lollitoss and provided the same instruction and comprehension questions as in Exp. 2a.

In the prediction phase, participants were asked to predict where the lollipop tosses would land for a first-time learner, just as in Exp. 2a. However, instead of having participants directly place the lollipop on the screen on a given trial, participants were asked to estimate the likelihood of the toss landing on each of the four possible regions: outside the board, on the white and the black rings, on the blue and the red rings, and inside the yellow circle. Four sliders (0%-100%) were provided, one for each region (see Fig. 1c). Participants were able to submit their response and proceed to the next trial if the sum of all four sliders was 100%. They completed the same seven rounds of predictions with a total of 19 trials as in Exp. 2a.

Results & Discussion

As in Exp. 2a, we found that people's predictions of learning over time were again best fit by an exponential decay function when probed using likelihood estimations (Figure 2e). To convert the reported probabilities to errors in the task space (e.g., distances from the center), we computed weighted distances for participants using their likelihood estimations on each trial (i.e., P(outside) * $D_{to\ center}$(outside) + P(white & black) * $D_{to\ center}$(white & black) + P(blue & red) * $D_{to\ center}$(blue & red) + P(bullseye) * $D_{to\ center}$(bullseye)). 'Distance to center' measures were the distance from the midpoint of each set of concentric rings to the center of the board, except the 'outside the board' area. Since this area was not bounded by an outer circumference, we used the mean distance observed in 'outside' area trial predictions from Exp. 2a. We again fit two possible functions, exponential decay and linear, to participant's predictions of learning data over the 19 trial bins. As in the last two experiments, 49 out of 50 participants' predictions converged for the exponential model, so we focused on these 49 participants for model comparisons. 40 out of 49 participant data were better described by an exponential model rather than a linear model and had negative ΔAIC. Furthermore, we found that the exponential models had significantly lower AIC scores than the linear models using a paired Wilcoxon test (see Figure 2f for ΔAIC; $V = 95$, $p < .001$).

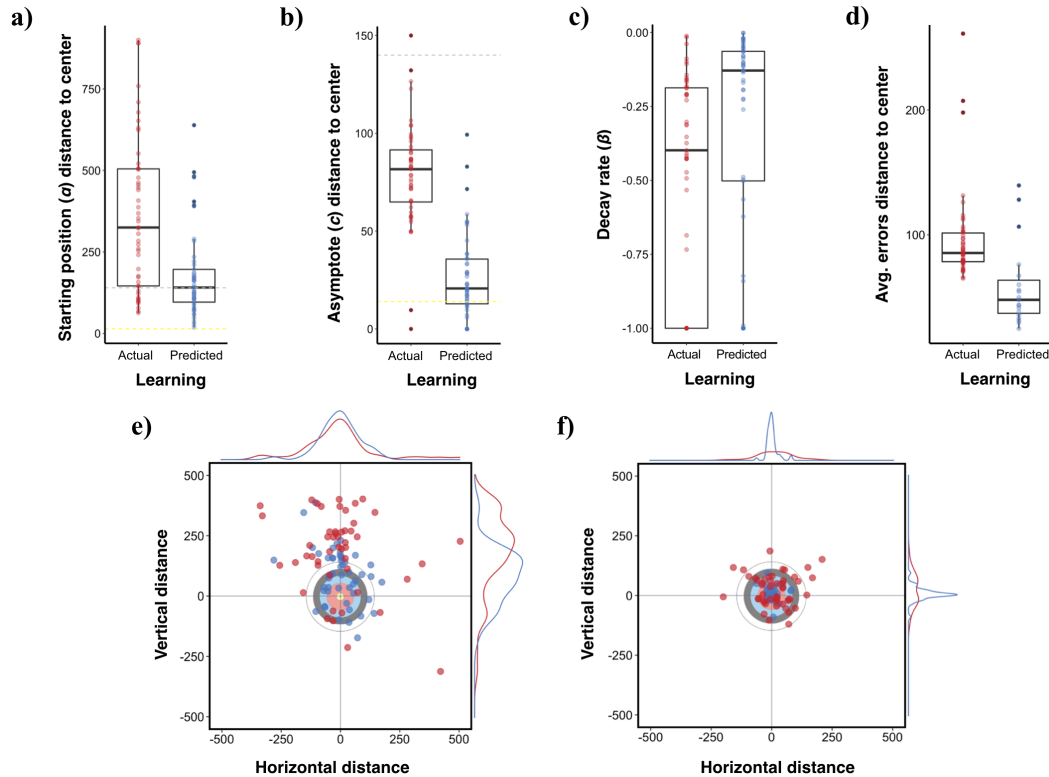Together with the results from Exp. 2a, we found con-

285

Figure 3: Model parameters from the best exponential decay model fits in Exp. 1 - actual learning and Exp. 2a - predicted learning for individual participants: (a) estimates for the starting points (α), (b) estimates for the asymptote (c), and (c) estimates for the learning rate (β). The grey dashed line indicates the outer border of the target board and the yellow dashed line the border of the bullseye. (d) average errors across 50 trials for Exp. 1 and Exp. 2a. Participant's 1st (e) and 50th (f) actual tosses (Exp. 1, dark red) and predicted tosses (Exp. 2a, blue) are overlaid for visualization. Data and estimates shown here are from 49 participants (Exp. 1) and 49 participants (Exp. 2a) where the models converged.

verging evidence that people represent learning trajectories as exponential when asked to reason about learning by making either point estimates (Exp. 2a) or likelihood estimates (Exp. 2b) of performance. Although people accurately represent the overall shape of learning as exponential, it is possible that they are inaccurate in representing specific features of the learning curve, like the starting point and the learning rate.

Actual vs. Predicted Learning

To probe people's intuitions about specific features of the learning curve, we compared parameters from predicted learning curves (Exp. 2a) to parameters from actual learning curves (Exp. 1). Exp. 2b was not included in this analysis because it produced less precise estimates (weighted distance estimates from participants' likelihood responses) than Exp. 2a. Comparing Exp. 1 and 2a, we found that people assumed learners would have better starting performance ($M = 179.27$, $CI = [141.34, 218.43]$) when providing trial-by-trial point estimates than they actually do during learning ($M = 363.09$, $CI = [301.40, 432.03]$; Wilcoxon test, $W = 1774$, $p < .001$). When estimating the errors a naïve learner would make on their first toss, people's prediction were near the outer border of the target board, while

actual learners' starting positions are further away from the target board (Figures 3a & e). People also believed that learners would get closer to the bullseye ($M = 26.38$, $CI = [20.42, 32.84]$) by the end of the learning process than they actually do ($M = 79.22$, $CI = [72.02, 86.50]$; Wilcoxon test, $W = 2218.5$, $p < .001$; Figures 3b & f). When we uncovered the learning rates from people's point estimates, we found that people predicted that learners would have a slower learning rate ($M = -.32$, $CI = [-.42, -.22]$) than their actual learning rates ($M = -.48$, $CI = [-.58, -.38]$; Wilcoxon test, $W = 768.5$, $p < .005$; Figure 3c). Similarly, people believed that learners would make significantly fewer errors on average ($M = 58.29$, $CI = [44.89, 73.89]$) than they actually made ($M = 96.26$, $CI = [87.48, 106.94]$) across 50 trials (Wilcoxon test, $W = 805$, $p < .001$; Figure 3d).

Thus, although the general shape of people's predictions about learning matched the actual learning data, the exact features (e.g., the range of people's errors, the starting performance, and the learning rate, etc.) did not. People assumed that learners start closer and end closer to the bullseye than they actually do across 50 trials. In turn, people assume that learners make fewer errors than in reality. Miscalibrating the learner's starting point may have led people to further

underestimate naive player's learning rate. This suggests that people's assumptions about task difficulty, indexed by starting point, in Exp. 2a may be key for "parameterizing" the intuited exponential functions.

General Discussion

Our work shows that people accurately represent the exponential trajectory of skill learning over time. Replicating prior work (Heathcote et al., 2000), we first established that learning on a novel visuomotor task, Lollitoss, was best fit by an exponential decay function (Exp. 1). Across two studies probing trial-by-trial predictions of learning in Lollitoss with point estimates (Exp. 2a) and with likelihood estimates (Exp. 2b), we found converging evidence that people accurately represent learning as unfolding exponentially. However, we also found that people misestimate specific features of the learning curve: people tend to overestimate learners starting and ending performance and underestimate their rate of learning.

In contrast to prior work, we found that people are surprisingly good at exponential thinking. Participants predicted exponential learning when providing explicit trial-by-trial point estimates (Exp. 2a) and more abstract likelihood estimates (Exp. 2b). Why do our results differ from prior work on the Exponential Growth Bias (Wagenaar & Sagaria, 1975)? One reason may be due to *how* we asked people to make predictions. For instance, in the classic duckweed and pond paradigm, Wagenaar and Timmers (1978) asked participants to indicate "the proportion of elapsed time" in comparison to when the pond will be fully filled. This task question is very cognitively demanding, requiring participants to reason and connect abstract properties like time and growth. In contrast, participants in our work were asked to do something less cognitively demanding - they directly predicted a series of learning outcomes at specific moments in time (Exp. 2a & 2b). Importantly, our findings are in line with recent work showing that exponential reasoning is more accurate when individuals make predictions on shorter time scales (e.g., how much will Covid-19 cases increase in 3 days vs. 15 days, Lammers et al., 2020).

Although people's predictions of skill learning matched the exponential decay shape of actual skill learning, we found interesting points of misalignment. People systematically misrepresent the precise learning curve parameters (i.e., the intercept, rate, and asymptote). For example, individuals think that naïve learners will start out and end closer to the bulls-eye than in reality. People also underestimate the amount of improvement over time (i.e., the rate). This result is in line with prior work showing that people tend to underestimate the amount of exponential growth (Wagenaar & Sagaria, 1975; Wagenaar & Timmers, 1978). Many factors could lead to discrepancies between predicted and actual learning curve parameters, including the perceived difficulty of the task (note that people who predicted learning had never played Lollitoss), specific task features (e.g., how long to hold the 'enter'

key), and the perceived skill level of individuals playing the game. However, it is exactly these same 'free parameters' that make the alignment between people's predicted and actual exponential learning functions even more impressive.

Past work has demonstrated that the way people think about moment to moment learning influences their actual learning. For example, before a test, people tend to prioritize studying vocabulary words that are of medium difficulty and avoid studying the words they already know or the words that are very difficult (Thiede et al., 2003; Finn & Metcalfe, 2008). Similarly, adults are sensitive to learning trajectories and choose to spend their energy on tasks with steeper learning curves, where they can experience more improvement within a shorter period of time (Ten, Kaushik, Oudeyer, & Gottlieb, 2021). Thus, a natural extension of the current work is to ask whether representations of the learning process also guide actual learning. There are two ways to approach this question. The first is to see whether individual differences in people's prediction of the learning process (the overall shape, as well as the intercept and rate) impacts the way they approach learning (e.g., their motivation, learning preferences, etc). A second way is to prime individuals to think of learning more or less accurately and see how that impacts their approach to learning. For example, if we prime people to incorrectly think that their rate of learning will increase over time, this may motivate learners initially (they experience more improvement than expected) but demotivate them over time (when their actual rate of learning is much lower than expected).

One open question is when in development people possess this capacity to intuit the shape of learning. On one hand, work in developmental psychology and cognitive science points to the remarkable sophistication of human metacognitive reasoning. Preschoolers accurately track their past performance (Hembacher & Ghetti, 2014) and selectively choose to switch to easier tasks when their performance hasn't improved over time (Leonard, Duckworth, Schulz, & Mackey, 2021). Furthermore, children allocate attention towards displays of intermediate complexity (Kidd, Piantadosi, & Aslin, 2010) and possess an intuitive sense of difficulty, at least on simple tasks (Gweon, Asaba, & Bennett-Pierre, 2017). However, simulating the time course of learning may be a cognitive challenging task that emerges later in development. Ongoing work is exploring whether young children intuit that learning unfolds non-linearly using a simplified paradigm.

In sum, we found that people are remarkably accurate at representing the shape of learning over time in the domain of skill learning. However, we also identified specific misconceptions of parts of the learning process - people overestimate learner's starting and ending performance and underestimate their rate of learning. This works sets the stage for a series of follow-up studies probing representations of learning across ages, question types, and domains, as well as work exploring the functional consequences of misrepresentations of learning trajectories.

Acknowledgements

References

Finn, B., & Metcalfe, J. (2008). Judgments of learning are influenced by memory for past test. *Journal of memory and language*, *58*(1), 19–34.

Finn, B., & Metcalfe, J. (2014). Overconfidence in children's multi-trial judgments of learning. *Learning and Instruction*, *32*, 1–9.

Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: implications of a quantitative analysis. *Proceedings of the National Academy of Sciences*, *101*(36), 13124–13131.

Gweon, H., Asaba, M., & Bennett-Pierre, G. (2017). Reverse-engineering the process: Adults' and preschoolers' ability to infer the difficulty of novel tasks. In *Cogsci*.

Heathcote, A., Brown, S., & Mewhort, D. J. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic bulletin & review*, *7*(2), 185–207.

Hembacher, E., & Ghetti, S. (2014). Don't look at my answer: Subjective uncertainty underlies preschoolers' exclusion of their least accurate memories. *Psychological science*, *25*(9), 1768–1776.

Hutzler, F., Richlan, F., Leitner, M. C., Schuster, S., Braun, M., & Hawelka, S. (2021). Anticipating trajectories of exponential growth. *Royal Society open science*, *8*(4), 201574.

Keil, F. C., Lockhart, K. L., & Schlegel, E. (2010). A bump on a bump? emerging intuitions concerning the relative difficulty of the sciences. *Journal of Experimental Psychology: General*, *139*(1), 1.

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2010). The goldilocks effect: Infants' preference for stimuli that are neither too predictable nor too surprising. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 32).

Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor learning. *Compr Physiol*, *9*(2), 613–663.

Lammers, J., Crusius, J., & Gast, A. (2020). Correcting misperceptions of exponential coronavirus growth increases support for social distancing. *Proceedings of the National Academy of Sciences*, *117*(28), 16264–16266.

Leonard, J. A., Duckworth, A. L., Schulz, L. E., & Mackey, A. P. (2021). Leveraging cognitive science to foster children's persistence. *Trends in Cognitive Sciences*.

Letourneau, S. M., & Sobel, D. M. (2020). Children's descriptions of playing and learning as related processes. *PloS one*, *15*(4), e0230588.

Lockhart, K. L., Goddu, M. K., & Keil, F. C. (2021). How much can you learn in one year? how content, pedagogical resources, and learner's age influence beliefs about knowledge acquisition. *Cognitive Development*, *60*, 101115.

Mckenzie, C. R., & Liersch, M. J. (2011). Misunderstanding savings growth: Implications for retirement savings behavior. *Journal of Marketing Research*, *48*(SPL), S1-S13. doi: 10.1509/jmkr.48.SPL.S1

Metcalfe, J., & Finn, B. (2008). Evidence that judgments of learning are causally related to study choice. *Psychonomic Bulletin & Review*, *15*(1), 174–179.

Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. *Cognitive skills and their acquisition*, *1*(1981), 1–55.

Richardson, E., Sheskin, M., & Keil, F. C. (2021). An illusion of self-sufficiency for learning about artifacts in scaffolded learners, but not observers. *Child Development*.

Son, L. K., & Metcalfe, J. (2005). Judgments of learning: Evidence for a two-stage process. *Memory & Cognition*, *33*(6), 1116–1129.

Son, L. K., & Sethi, R. (2010). Adaptive learning and the allocation of time. *Adaptive Behavior*, *18*(2), 132–140.

Stango, V., & Zinman, J. (2009). Exponential growth bias and household finance. *The Journal of Finance*, *64*(6), 2807–2849.

Ten, A., Kaushik, P., Oudeyer, P.-Y., & Gottlieb, J. (2021). Humans monitor learning progress in curiosity-driven exploration. *Nature communications*, *12*(1), 1–10.

Thiede, K. W., Anderson, M., & Therriault, D. (2003). Accuracy of metacognitive monitoring affects learning of texts. *Journal of educational psychology*, *95*(1), 66.

Thorndike, E. L. (1913). *The psychology of learning* (Vol. 2). Teachers College, Columbia University.

Wagenaar, W. A., & Sagaria, S. D. (1975). Misperception of exponential growth. *Perception & Psychophysics*, *18*(6), 416–422.

Wagenaar, W. A., & Timmers, H. (1978). Extrapolation of exponential time series is not enhanced by having more data points. *Perception & Psychophysics*, *24*(2), 182–184.