




# Reinforcement learning tutor better supported lower performers in a math task

Sherry Ruan<sup>2</sup> · Allen Nie<sup>1</sup> · William Steenbergen<sup>2</sup> · Jiayu He<sup>2</sup> · J. Q. Zhang<sup>2</sup> · Meng Guo<sup>1</sup> · Yao Liu<sup>2</sup> · Kyle Dang Nguyen<sup>1</sup> · Catherine Y. Wang<sup>1</sup> · Rui Ying<sup>2</sup> · James A. Landay<sup>1</sup> · Emma Brunskill<sup>1</sup> 

Received: 10 April 2023 / Revised: 27 September 2023 / Accepted: 3 October 2023 /  
Published online: 9 February 2024  
© The Author(s) 2023

## Abstract

Resource limitations make it challenging to provide all students with one of the most effective educational interventions: personalized instruction. Reinforcement learning could be a pivotal tool to decrease the development costs and enhance the effectiveness of intelligent tutoring software, that aims to provide the right support, at the right time, to a student. Here we illustrate that deep reinforcement learning can be used to provide adaptive pedagogical support to students learning about the concept of volume in a narrative storyline software. Using explainable artificial intelligence tools, we extracted interpretable insights about the pedagogical policy learned and demonstrated that the resulting policy had similar performance in a different student population. Most importantly, in both studies, the reinforcement-learning narrative system had the largest benefit for those students with the lowest initial pretest scores, suggesting the opportunity for AI to adapt and provide support for those most in need.

**Keywords** Reinforcement learning · Education · Children · Artificial intelligence

## 1 Introduction

Many children fail basic reading and math standards,<sup>1</sup> and the number of such students has greatly increased during the covid-19 pandemic. One-on-one human tutoring can be highly effective (Nickow et al., 2020), in part because it enables students to receive personalized,

---

<sup>1</sup> See for example the 2019 US's National Assessment of Educational Progress report: <https://www.nationsreportcard.gov/highlights/reading/2019/>.

---

Editors: Minmin Chen, Omer Gottesman, Lihong Li, Yuxi Li, Zonging Lu, Niranjani Prasad, Zhiwei Qin, Csaba Szepesvari, Matthew Taylor

---

Sherry Ruan and Allen Nie have contributed equally to this work.

---

Due to a conflict of interest, EB did not review submitted papers for this special issue.

---

Extended author information available on the last page of the article

differentiated instruction, but it is often prohibitively expensive. Educational software aims to provide some of this personalized instruction at scale, but can still be costly and slow to build.

Reinforcement learning (RL) could reduce the cost of developing effective learning technology by automating the process of specifying how best to support a student through their learning journey. RL algorithms learn from data to choose an intervention (such as a hint), given the current context (such as an estimate of a student's knowledge) to maximize the expected value of some desirable outcome, such as test scores. Preliminary work on using RL for improving educational software has enabled encouraging gains on learning outcomes (Mandel et al., 2014; Chi et al., 2011; Park et al., 2019; Bassen et al., 2020) or student persistence (Mandel et al., 2014; Bassen et al., 2020). Such systems have been limited to selecting among practice items, and not all experiments with using RL for enhancing educational technology have yielded positive outcomes (see review (Doroudi et al., 2019)). It is unknown if reinforcement learning could be used to automatically tune and optimize broader types of learning systems, such as the pedagogical feedback provided in a narrative environment. We also seek to do so in a way that is interpretable and robust—two important aspects of AI for societally impactful applications, that are receiving increasing attention broadly, but have not yet been considered as much in the AI for education space.

To address this, we created a narrative-based adaptive pedagogical-supported educational software to support math concept learning for students roughly ages 9–12 and used reinforcement learning to adaptively (machine) learn the responses to provide support for student learning. Recent advances in explainability methods for deep neural networks (e.g. Lundberg and Lee (2017); Sundararajan et al. (2017)) have made it possible to use advanced tools for modeling without sacrificing interpretability. We used these methods to help understand if and how the system is learning to differentiate in order to optimize desired outcomes. An additional key consideration is whether the learned pedagogical support would generalize to a different student community, as all schools may not be able to support online adaptive RL systems. We tested if a distilled version of the decision policies learned in the first study could be used in a different population of students that was a more geographically diverse population with a lower household income distribution. In both studies, students with the lowest pretest scores improved using our RL-powered narrative AI system, and more than compared to students using a baseline system. This highlights the potential for reinforcement learning to tune educational software parameters to enhance effectiveness, in a way that is interpretable, transfers to other populations, and can help those most in need of support.

## 2 Related work: reinforcement learning for student learning

Reinforcement learning has seen impressive successes in areas like robotics (Levine et al., 2016) and game playing (Silver et al., 2018). The goal of a reinforcement learning algorithm is to compute a strategy (referred to as a “policy”) that specifies the intervention (such as a pedagogical activity) to choose in a particular context (e.g., a learner's knowledge state and frustration level), in a way that is expected to maximize desired outcomes (e.g., test scores, engagement, retention). A key challenge is that the algorithm does not have prior knowledge of the statistical parameters governing the process by which contexts evolve, and outcomes occur. Instead, an algorithm must learn from experience by analyzing actual decisions made and their outcomes, a strategy with high expected outcomes.

In the context of education, there have been some promising results that reinforcement learning can improve word acquisition of preschoolers interacting with a social robot (Park et al., 2019), the persistence of learners during a fractions game (Mandel et al., 2014), the performance of college students learning introductory physics (Chi et al., 2011), undergraduates learning discrete mathematics (Zhou et al., 2019), and the outcomes and efficiency of working adults learning linear algebra (Bassen et al., 2020). However, in other settings, there has been little benefit over a reasonable control condition (Rowe & Lester, 2015; Doroudi et al., 2019). More broadly, work on intelligent tutoring systems and computer-assisted learning suggests that personalized feedback and support in educational software can be an effective way to support student learning (Corbett, 2001; Beal et al., 2010; VanLehn, 2011), but most prior work has focused on software designed to be used in the classroom where there are additional mechanisms to keep students' attention.

We hypothesize that reinforcement learning may be particularly beneficial when learning is happening out of the classroom, or motivation and engagement are particularly critical, or in less traditional curricula that move towards different forms of instruction rather than lecture and practice. Learning sciences offer less guidance about how to best support students in these settings. Yet, such educational settings are likely to be increasingly important in the future, both due to immediate challenges due to the covid-19 pandemic and aftermath, as well as due to the types of skills needed for success in the 21st century. Reinforcement learning may inform data-driven instruction for such settings, and we focus our attention on learners outside the classroom in this work.

As another contrast between our focus and prior related work, in the context of education, it is both important and of interest to understand what the algorithm learns to do: what personalized decisions are made for different contexts and individuals, and who is most helped by the algorithm. Such issues have been historically largely unstudied in the reinforcement learning research community, with some notable exceptions (e.g. Shen et al. (2016); Zhou et al. (2022)), but are an important part of our current work.

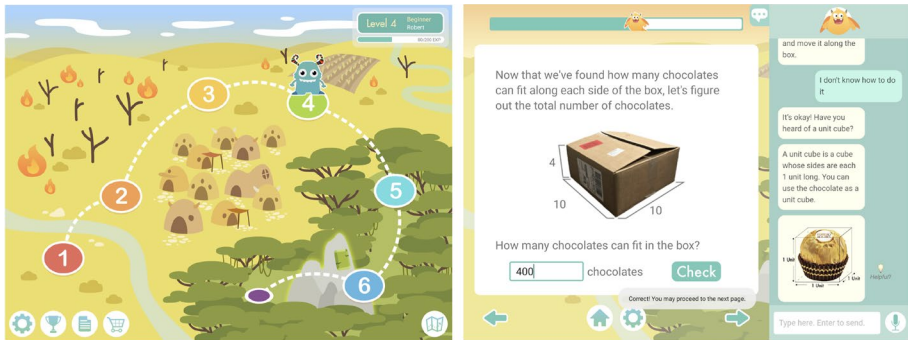
### 3 Interface design

Learning science principles can often be too broad to inform the specific design decisions needed to create engaging, effective educational software. For example, a narrative-based, basic chatbot-supported<sup>2</sup> educational interface can lead to significant learning and engagement gains over a no-narrative, no-chatbot variant (Ruan et al., 2020), but doing so well is subtle. Here the effective chat tutoring system actually used humans to act as chatbots, in a wizard-of-oz style study. In contrast, a different narrative-based system with standard step-by-step hints (which are common in intelligent tutoring systems) provided no benefit over the no-narrative, no-hint control condition (Ruan et al., 2020).

RL has the potential to be particularly helpful in such situations where personalization may be key. In this work, we used an informal online learning environment to teach students about the concept of volume that was previously developed (Ruan et al., 2020).<sup>3</sup> Learning tasks in this system are embedded in a narrative storyline. In response to student input, a companion AI tutor selects among four common pedagogical strategies: providing

<sup>2</sup> Note that our work was conducted before the public launch of ChatGPT in November 2022.

<sup>3</sup> Code for the interface is available here: <https://github.com/StanfordAI4HI/smart-primer-website-public>



**Fig. 1** Tutoring AI Guide Interface: A child solves a math problem while interacting with the AI-driven tutoring guide. The child can click on the “helpful?” button if they consider the AI tutor’s response to be helpful. The child can also click on “I want to stop playing” to quit the activity at any time

direct hints, generic encouragement, and guided prompts that scaffold the student (e.g., “Have you heard of a unit cube?”), or passive positive acknowledgment (emoticon smiley face). Figure 1 shows a screenshot of the software used.

## 4 Approach

### 4.1 Feature space

Due to past success in RL systems for adult learning (Chi et al., 2011; Bassen et al., 2020), we use a small set of features, specifically an eight-dimensional state space, described in detail below:

- *Grade*: The elementary school grade a child is in, ranging from 3–5.
- *Pre-score*: The score a child receives for the pre-test, ranging from 0–8.
- *Step*: The step of the task a child is in, ranging from 1–6. (This is automatically defined by the task interface).
- *Failed attempts*: The number of failed attempts made by the child in the current step. It is a non-negative integer. There is a single correct answer to each step.
- *NLP positive score*: A score that reflects the positive sentiment in the last phrase typed in by the child. It is a float ranging from 0–1. An automatic sentiment analysis tool from NLTK (Bird et al., 2009) is used to calculate this.
- *NLP negative score*: A score that reflects the negative sentiment in the last sentence mentioned by the child. It is a float ranging from 0–1. An automatic sentiment analysis tool from NLTK (Bird et al., 2009) is used to calculate this.
- *NLP help score*: A score that reflects the extent to which the child asks for help in the message sent. It is a float ranging from 0–1 and calculated as the semantic similarity between the child’s message and “help”.
- *Anxiety score*: The score of the math anxiety test (Carey et al., 2017) that the child takes prior to beginning the activity.

The observation vectors were normalized element-wise before being used for training and prediction. Grade, pre-score, and anxiety score are static variables. Other variables are affected by the actions the policy takes and change as the child is solving each step of the task.

## 4.2 RL policy learning

### 4.2.1 The simulation phase

RL algorithms were run on a simulator before any real-world experiments were done to get an initial estimate of the performance and test the algorithm's potential. The simulator models children with various characteristics and their interactions with the math problem and interactive teaching support (the actions selected by the RL policy). Note that this simulator used simple, hand designed models of student learning and was not intended as a high fidelity replica of student learning: rather we used it to help explore how quickly a reinforcement learning agent might be able to learn an effective policy in such an environment (aka we expected our later experiment to be performed with hundreds, not millions, of students), and to tune the hyperparameters of our setup.

These early simulations informed our choice of a small function model for use in our later experiments. For example, we explored various multiple policy architectures and converged on 2 hidden layers, since in our simulations the parameters for a small instructional model could be learned within a couple of hundred simulated students. In this way, these experiments can serve as a very rough sanity check prior to the first experiment. The code for these simulations is available here: [https://github.com/StanfordAI4HI/SmartPrimer\\_Gym](https://github.com/StanfordAI4HI/SmartPrimer_Gym).

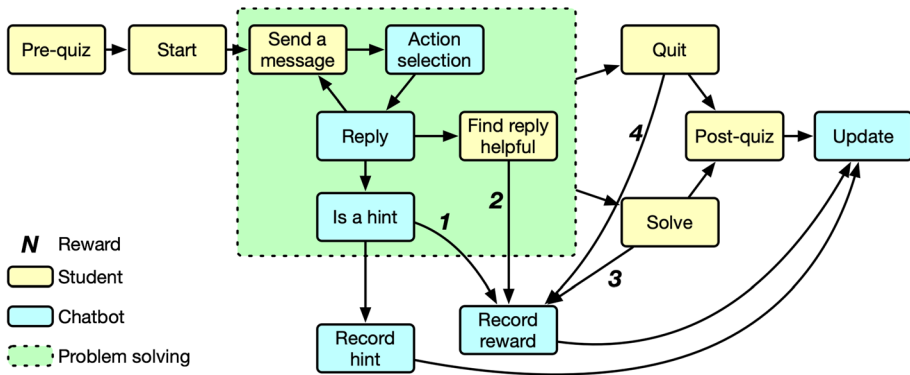
### 4.2.2 Online learning Phase

Throughout the math-learning activity, children have access to an AI guide on a side panel that provides encouragement, hints, and companionship. The goal is for the AI guide to provide additional engagement with the math activity and provide adaptive support that facilitates learning gains. The AI guide takes on the persona of the monster that children select in the fantasy-based narrative. Before entering the math learning activity, children are brought through a short tutorial in which they communicate with the AI guide, which introduces itself and asks about the children. This tutorial serves to familiarize the children with the AI guide interface and build social rapport between the AI guide and the children. We provide a workflow in Fig. 2.

The RL decision policy takes in a vector describing features of the learner state and outputs a particular support type (of the 4 options) to provide. The RL algorithm aims to learn an automated decision policy to maximize the expected reward function, which should capture the key desired outcomes. We specify the reward when a student  $j$  finishes as:

$$R_j = \sum_{i=1}^8 [\max(0, post_{ij} - pre_{ij})] - \lambda * n_{hj} + \beta n_{uj} + \mathbb{1}(quit_j),$$

( $\lambda = 0.013$ ,  $\beta = 0.1$ ), where the first term is the sum over items of the  $j$ -th student's clipped learning gain from pretest to post-test on item  $i$  of the assessment, the second term is a tiny penalty on the number of hints  $n_{hj}$  given by the system to the student (since too many hints may reduce learning), the third term provides a small bonus for the number of times  $n_{uj}$



**Fig. 2** The interaction between user and RL AI guide. The RL AI guide selects one of four actions and replies to the user once a message is received. The reward function is updated both during the interaction and after the child completes the post-quiz. Rewards 1–4 correspond to the reward functions described in Sect. 4.2 (Reward Function). The RL AI guide performs an update after every five children

child  $j$  marked an AI guide reply as helpful<sup>4</sup>, and the last term  $\mathbb{1}(quit_j) = -8$  is a penalty if the learner quits before completing the task. Note that we choose to use a clipped learning gain (pair-wise increases between pretest to post test per problem). Problems were matched on the two tests to be similar with different specific numerical quantities (aka problem 3 on the test 1 was similar to problem 3 on test 2). We expected it is highly unlikely for the policy and practice with the math software to cause negative learning gains, and clipping the signal at 0 means that if a student did a problem correctly initially but not on the post test (which could occur for many reasons, including a student not focusing on the post test), did not impact the resulting reward signal. Bassen et al. (2020) previously stated that using such a clipped signal improved the stability and efficiency of reinforcement learning in their learning task. We set the hyperparameters ( $\lambda = 0.013$  and  $\beta = 0.1$ ),<sup>5</sup>

The proximal policy optimization (PPO) algorithm (Schulman et al., 2017) was used to learn the decision policy to optimize the expected reward.<sup>6</sup> The policy architecture is stochastic. The hyper-parameter used in the online study was  $\epsilon = 0.2$ . Both the policy neural network and value function neural network had two hidden layers with 16 nodes and a tanh activation function. We used an Adam optimizer with a learning rate of 0.0025 for both. The RL policy is implemented with the RLGraph package (Schaarschmidt et al., 2019). This optimization method was chosen as it has shown potential in similar situations, for example, in Bassen et al. (2020).

<sup>4</sup> Post study we analyzed if the number of times a student marked an AI guide as helpful was correlated with post test - pretest score gains. We used both a Pearson correlation coefficient to estimate a linear relationship, and also performed a Kendall correlation test. The estimated relationship was small in both cases, suggesting that in the future, omitting this term may yield a more stable reward function.

<sup>5</sup> As has been done in prior work (e.g. Bassen et al. (2020)) we created a simple simulator of student learning that we have made available [https://github.com/StanfordAI4HI/SmartPrimer\\_Gym](https://github.com/StanfordAI4HI/SmartPrimer_Gym) and explored a small set of hyperparameter values for  $\lambda$  and  $\beta$  in that simulator, centered around those proposed in Bassen et al. (2020). As the simulator was extremely coarse, it is possible that other hyperparameters would do better in the real experiment.

<sup>6</sup> Our code is available at [https://github.com/StanfordAI4HI/smart\\_primer\\_bot\\_public](https://github.com/StanfordAI4HI/smart_primer_bot_public).

### 4.2.3 Offline reinforcement learning

We also performed offline reinforcement learning to extract another policy for use in a subsequent experiment. We did this for multiple reasons. First, as described later, during online reinforcement learning, the policy had not yet converged by the end of study 1, and we wanted to compare a static learned policy to a control, where the differences might be clearer. Second, we were curious whether we might extract a higher-performing decision policy using offline learning. Third, in most experimental sciences, research is hoped to provide findings that generalize beyond the specific research setting. Such generalizability is also of key interest in machine learning. Therefore an important open issue is whether automated pedagogical strategies obtained using reinforcement learning in one setting will transfer to similar settings.

We used offline reinforcement learning policy evaluation to select among potential new automated instructional policies using the data gathered from online reinforcement learning (in our study 1, as we will shortly describe). We considered two sets of algorithms for training potential instructional policies. The first is behavior cloning (Pomerleau, 1990; Sammut et al., 1992), a popular method for leveraging offline data to train an automated policy. Behavior cloning trains the model to imitate the probability distribution of actions that are outputted by the online policy.

Recall that during our online RL experiment, PPO was used to update the RL policy deployed at regular intervals. This meant that only a few students got the same identical policy in the RL condition. Therefore behavior cloning can be used to output a *single* RL policy that essentially distills an aggregate policy over the entire online RL experiment: in a sense, intuitively, though PPO does not have cumulative regret guarantees in our setting, our procedure at a high level is similar to theoretical proofs that show how an algorithm that achieves a particular cumulative regret can be used to output a single decision policy with a small simple regret by constructing a new decision policy that is an average over all the policies deployed by the algorithm up to a certain point. More precisely, behavior cloning minimizes the following loss:

$$\mathcal{L}_{\text{BC}}(\theta, \mathcal{D}) = \mathbb{E}_{(s,a,s') \sim \mathcal{D}} [D_{\text{KL}}(\pi_{\theta}(s) \| p(a|s))]$$

which in our setting, will create a single stochastic policy. Note that this policy may be different than any of the decision policies deployed during online RL.

The second style of algorithms we explored was offline policy gradient on the estimated performance of the trained instructional policy. This method has been used in several other offline RL optimization papers (see e.g. Metelli et al. (2018); Liu et al. (2020)). Here we used a weighted importance sampling (WIS) estimator to estimate the value of the policy,

$$\begin{aligned} \mathcal{L}_{\text{WIS}}(\theta, \mathcal{D}) = & \frac{1}{\sum_{i=1}^{|\mathcal{D}|} \left( \prod_{t=1}^L \frac{\pi_{\theta}(a_t|s_t)}{p(a_t|s_t)} \right)} \sum_{i=1}^{|\mathcal{D}|} \left( \prod_{t=1}^L \frac{\pi_{\theta}(a_t|s_t)}{p(a_t|s_t)} \right) R_i \\ & + \eta \cdot \frac{1}{\sum_{i=1}^{|\mathcal{D}|} \left( \prod_{t=1}^L \frac{\pi_{\theta}(a_t|s_t)}{p(a_t|s_t)} \right)^2} \end{aligned} \quad (1)$$

where  $R_i$  is the total reward for student  $i$ . This is called policy gradient via importance sampling (POIS). We also explored whether adding an effective sample size (ESS) penalty with hyperparameter  $\eta$  would help – ESS regularizes the difference between the learned policy  $\pi_{\theta}$  and the behavior policy  $p$ .

**Table 1** Hyperparameters considered during offline batch reinforcement learning. The policy network dimension describes the network structure: e.g. [4] is one layer of 4 hidden nodes, [16,16] is two hidden layers, each with 16 nodes

	Hyperparameter range
Training algorithm	BC, POIS
Policy network dimension	[4], [8], [16], [4,4], [8,8], [16,16]
Training epochs	[1, 5, 10]
ESS penalty $\eta$	[0, 0.01, 0.05]

We considered multiple hyperparameters for each of the two algorithm procedures (see Table 1). There are 108 hyperparameter combinations to learn our policy. We use an algorithm evaluation procedure where we partition the collected dataset into a train and validation set by randomly allocating 50% of students into one group and the rest into another. We repeat this strategy 10 times. We use this split dataset to choose the best model architecture, hyperparameters, and learning objectives, similar to what has been proposed in Nie et al. (2022). We trained our model on the training split and use weighted importance sampling (WIS) to evaluate the performance of this policy on the validation set. We apply the same learning procedure across all 10 splits and compute the average of the performances. We choose the best algorithm from the highest average performance on the validation set. We then apply this algorithm to train a policy that learns from the entire dataset.

In our evaluation, the behavior cloned policy was estimated to outperform the online policy in the majority of splits. Also, a small 1-layer fully connected neural network with 4-dimensional hidden state and Gaussian error linear unit (Hendrycks & Gimpel, 2016) activation function outperformed other model architectures.

Therefore we used the distilled, behavior cloned policy in our second experiment.

## 5 Experimental setups

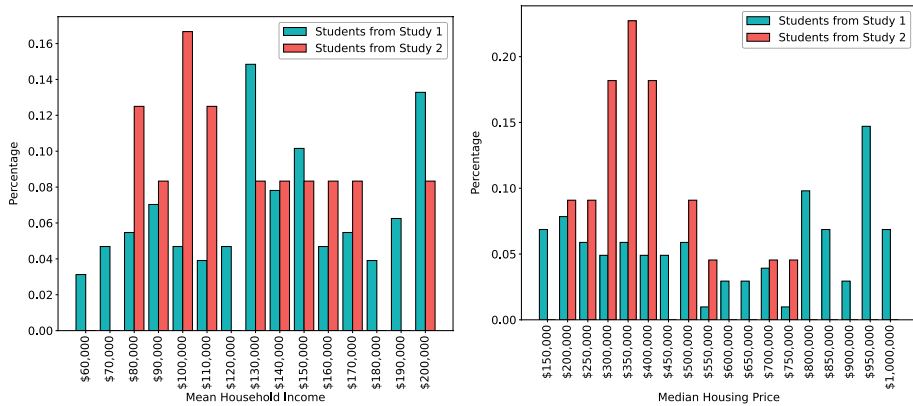
As a control condition, the interface included the mathematics task but had no narration and no adaptive support; similar to a mastery-style approach, students had to successfully complete one subpart before advancing.

While this may initially seem like a weak control condition, a past study (Ruan et al., 2020) on teaching an elementary school mathematics task had found that a similar control condition had performed similarly to a control condition with a narrative storyline, and slightly better than a control condition with a narrative storyline and step-wise hints (which are common in tutoring software).

In study 1, we examined the speed and effectiveness of using reinforcement learning to adapt the type of AI guide feedback given to learners. Due to covid-19 pandemic restrictions, all experiments were completed online. Subjects were randomly assigned to each condition, but with an unequal allocation—more students were assigned to the RL condition than the control condition. In total 269 elementary school students used the reinforcement learning-narrative educational software (RL). 70 students were in the control condition.

Subjects completed an 8 item assessment and a math anxiety survey (Carey et al., 2017), then used the volume education software, and then completed another assessment (identical up to numerical values, and cross-randomized across students), and Giggle Gauge, an engagement measure designed for studies with children (Dietz et al., 2020).





**Fig. 3** Distribution of household income (left) or median housing price (right) in the zip codes provided by subjects in study 1 and study 2. There were significant difference in the subject pools between the two studies

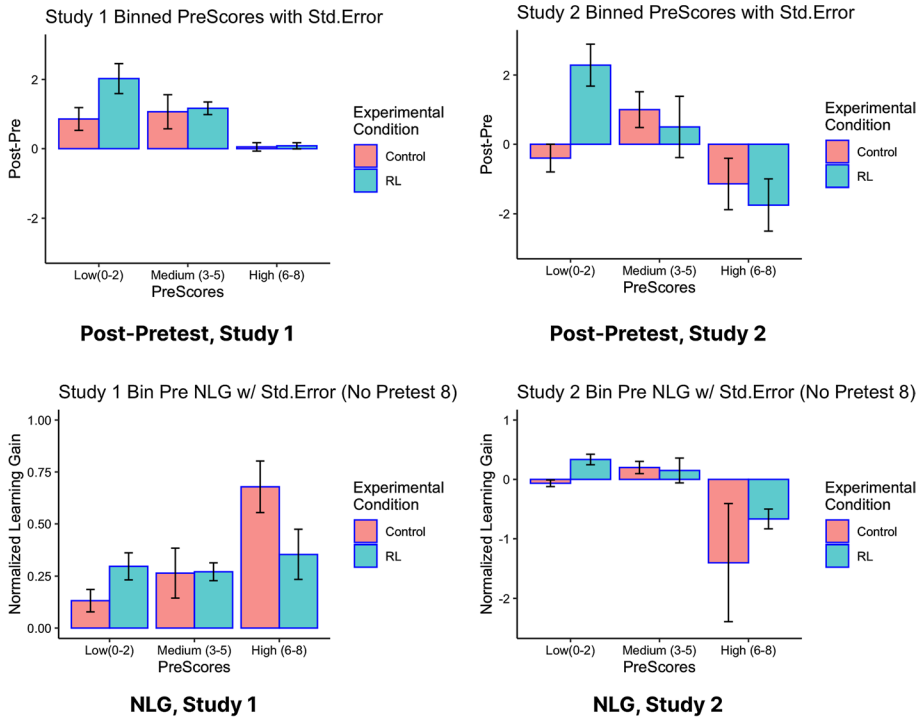
**Table 2** Mean (std. dev) results of children in both studies

	Online RL study 1		Distilled policy study 2	
	Control	Narrative AI	Control	Narrative AI
Number	68	258	18	17
Pretest	5.46 (2.73)	4.87 (2.38)	4.06 (2.58)	3.41 (2.31)
Posttest	5.89 (2.57)	5.72 (2.31)	3.83 (2.53)	4.12 (1.87)
Improvement	0.44 (1.26)	0.84 (1.89)	-0.22 (1.7)	0.7 (2.33)
Engagement	3.16 (0.62)	3.4 (0.52)	3.17 (0.52)	3.28 (0.57)
Completion	92.6%	94.5%	100%	100%

More specifically, Giggle gauge is a seven item self-report measure of engagement that was designed to be appropriate given children’s development.

In study 2 we were interested to see if the distilled behavior cloned policy learned from the online RL process (Sect. 4.2), would transfer to a new population of subjects. We then conducted study 2 with a new set of subjects (37 participants used for analysis): subjects were randomized into the same control condition as study 1, or using the single distilled RL policy.

In study 2, we recruited a broader population more similar to that of the U.S.A. For the original study, 113 participants out of 203 provided home zip codes. For the follow-up study, 16 participants out of 30 provided home zip codes. For those that did not provide their home zip code, we use their school zip code. Using these zip codes, we obtained the median housing price and mean annual household income from the fifth American Community Survey (in 2020), accessible through an API provided by the United States Census Bureau. Figure 3 shows the difference between the student groups in study 1 and study 2. We conduct the Kolmogorov-Smirnov 2-sided test between student populations of two studies on these variables. For both mean annual household income ( $Pr(F(x) = G(x)) = 0.02 < 0.05$ ) and median housing price



**Fig. 4** Top row, Post-test - Pretest (y-axis), Bottom row, Normalized learning gain (NLG)  $\frac{Posttest-Pretest}{MaxScore-Pretest}$  (y-axis). Scores are clustered by subjects with low (0-2), medium (3-5), and high (6-8) initial pretest scores. Error bars show standard errors. Note the NLG (bottom row) calculations exclude students who scored 100% on the Pretest since the NLG is not well defined

( $Pr(F(x) = G(x)) = 0.0005 < 0.01$ ), we found a significant difference between two populations. In addition, subjects were more geographically and racially diverse (see Appendix). In addition, study 1 was done when many more U.S.A. children attended school remotely. Thus, study 2 offers a chance to examine the generalizability of learned RL policies.

## 6 Results

Aggregate summaries are shown in Table 2. Some subjects completed the pretest or post-test twice due to a limitation in the system. We excluded these subjects from the results presented. There was no significant difference in the amount of improvement (post-test - pretest score) between the RL narrative condition and control condition (study 1: Wilcoxon rank test  $W = 9632.5, p = 0.2$ , study 2: Wilcoxon rank test  $W = 185.5, p = 0.281$ ).

However, encouragingly, in both studies, there was a trend for subjects with a low initial pretest score (0-2) to have a much larger improvement between the pretest and post-test in the RL narrative condition (Fig. 4, top row). The average improvement for these students was 2.02 in study 1 (N=41), and 2.29 in study 2 (N=7), out of a total score range was (0-8). There was a significant difference in the change in scores between the RL condition and control condition in study 2 for those with low pretest scores (0-2) (Wilcoxon rank

test  $W = 2, p = 0.013$ ), though this difference does not persist after correcting for multiple-hypothesis testing, and all other differences for studies and pretest groups were not statistically significant under the same test.

Engagement scores range from 1 to 4 and subjects with low initial pretest scores (0-2) also trended to having much higher engagement in the RL AI guide condition (study 1 mean engagement score 3.29 (N=40), study 2, mean engagement score 3.28 (N=7)) than in the control condition (study 1 mean engagement score 2.7 (N=14), study 2, mean engagement score 2.7 (N=5)). Prior work suggests interpreting scores below 3.0 as low engagement and 3.0–3.6 as moderate engagement (Dietz et al., 2020).

The assessment used may be subject to ceiling effects, as a number of students did receive the maximum score (8) on either the pretest or the post-test. Though the pretest scores did not significantly differ between the two conditions, in either study, since the control pretest scores were slightly higher, ceiling effects may have impacted the control condition more.

To address this, we also repeated our analysis using normalized learning gains (NLG),  $\frac{\text{Posttest} - \text{Pretest}}{\text{Maximumscore} - \text{Pretest}}$ , which represent the fraction of improvement made by subjects, relative to the possible improvement. Note this excludes any subjects who scored the maximum score on the pretest since the NLG is not well-defined for such students. There was no significant difference between the RL narrative condition and control condition for NLG in either case (study 1,  $W = 4394.5$ ,  $p\text{-value} = 0.6978$ ; study 2,  $W = 104.5$ ,  $p\text{-value} = 0.3819$ ).

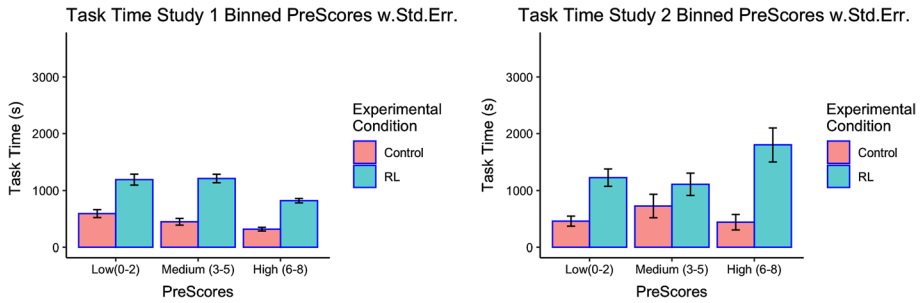
Like for posttest - pretest, we observe larger normalized learning gains for the RL narrative condition than the control condition for initially lower performing students, in both studies (Fig. 4, bottom row). The NLG performance for students with medium pretest scores is similar in both conditions, as was also seen for such subjects' posttest minus pretest scores. The pattern for the highest performing students is slightly different than for the post-test - pretest scores but should be taken lightly: as stated, the NLG analysis ignores all students with maximum pretest scores. Note that an NLG of 75% for the initially high-performing student group would be at most a  $2 * 0.75 = 1.5$  post-test - pretest improvement (since 2 is the largest possible gain, if the student scored 6 on the pretest, and it is lower if the student scored 7), whereas a 30% improvement for the initially low performing student group is at least a gain of  $6 * 0.3 = 1.8$  on the post-test - pretest (since  $\text{MaxScore} - \text{Pretest} \geq 6$  for such subjects).

Together these analyses encouragingly suggest that the RL narrative condition trends to provide a bigger benefit to initially lower-performing students than the control condition. We now provide some additional analyses into the RL process and the potential mechanisms underlying this difference.

## 6.1 RL online learning

In study 1, the RL agent updated the AI guide pedagogical policy over subjects, but during the 28 policy updates (after 10 subjects each), we observed significant variability, and the performance had not converged.

We hypothesize this may be due to several factors. Likely most importantly, we saw a significant variation in the pretest scores of subjects over time. This may be in part because we performed rolling recruitment, adding additional recruitment sources during the study, which likely caused some shift in the distribution of the underlying students. In addition, the natural variation across third to fifth-graders and student background skills means that



**Fig. 5** Time on task (sec) (y-axis) by low (0-2), medium (3-5), and high (6-8) initial pretest scores. Error bars show standard errors. Students whose time on task exceeded 90 min (8 students) were excluded from the analysis since it was likely such students might have taken significant breaks

across small sets (such as the 10 trajectories used each round for PPO), it is quite possible to have a substantial difference in the pretest scores of those subjects. If any of the students are already at or near the ceiling of the pretest scores, there will be almost no potential room for improvement for the RL policy. Indeed there may be some natural regression to the mean, which means that an RL policy that looked promising in prior rounds for related states, may now look worse (depending on the particular generalization). Even without this potentially shifting population, ten trajectories (subjects) is a small size to average over when performing policy updates, so the gradient may be quite noisy. This suggests that performing stratification and trying to ensure a stable distribution of initial start states over participants might lead to faster convergence and better results.

However, despite this, through training, subjects in the AI guide condition consistently match or exceed the average performance of those in the control condition.

## 6.2 Investigating other explanations for the benefit to low pretest subjects

A natural question is what is the mechanism behind the improved performance of subjects in the RL narrative condition over those in the control condition, for subjects with initially low pretest scores, and whether this could be due to factors beyond the RL-narration itself.

One potential hypothesis is that there were additional differences between the two conditions. Indeed, on average, subjects spend longer on the RL narrative condition task than in the control condition. As Fig. 5 shows,<sup>7</sup> this was consistent for students across all three groups of pretest performance, and the difference in time spent between the two conditions was largely similar for all three groups. However, only the students in the low pretest group seemed to have a significant benefit from the RL condition. It seems unlikely that time on task is the primary reason for improved performance in the RL narrative condition.

The study was conducted remotely, and a prescreening call was done with a guardian of each child participating to discuss the study, emphasize the child should do the task without assistance, and verify the child would be participating. However, it is still possible that guardians helped the children in some cases. It seems unlikely that for children with low

<sup>7</sup> We excluded individuals who took longer than 90 min on the task in this figure, since such subjects are likely to have taken breaks. All individuals who took at least 90 min took over 2 h, and there were 8 such individuals excluded using this restriction.

pretest scores, guardians helped them more if the child was in the RL condition than if they were in the control condition. Indeed the control condition offered less support and hints than the RL narrative condition, so the opposite seems more likely to be true. One potential exception is that the RL narrative condition involved a storyline, and while unlikely, depending on the subject's reading skills, it is possible that the guardian would have helped the subject to understand the text.

An interesting piece of evidence that it was the combination of the narrative and the RL text interaction that lead to student gains, is that we find students with low and medium pretest scores interacted more (sent more messages to the AI tutor) than students with high pretest scores in study 1. In particular, in study 1, the maximum number of messages sent by students was 20, with a long tail. The median number of messages sent in both the low and medium pretest score groups in study 1 was 4, and the high pretest score group in study 1 had a median of 2 messages sent. We conduct three two-sided two-sample Wilcoxon rank test on the number of messages sent by students in the RL condition (between the low and medium pretest group, low and high pretest group, and medium and high pretest group). There was no significant difference between the low and medium groups ( $W=1441.5$ ,  $p=0.92$ ), and there was a significant difference between the low and high ( $W=1348.5$ ,  $p=0.0015 < 0.0167=0.05/3$ , correcting for the 3 tests done here) and the medium and high groups ( $W=3167.5$ ,  $p=8.16 * 10^{-6}$ ,  $< 0.0167=0.05/3$ , correcting for the 3 tests done here). This helps to explain why the high pretest score students may not have benefited as much from the system – they did not (likely) need as much support, and did not interact much with the RL text based agent. However, this analysis only provides part of the insight for a potential mechanism, since the patterns of messages sent for the low and medium pretest groups was similar, and yet the performance gains (over the control condition) were larger for the low pretest students.

### 6.3 Integrated gradient analysis of policy on feature space

A natural question is whether benefits to subjects with low pretest scores may derive from the personalization capacity of the RL instructional policy. Indeed a key benefit of using RL to select activities is its potential to differentiate instruction if doing so is estimated to improve outcomes. Therefore it is of interest to evaluate what differentiation, if any, is done by the RL AI guide policy. However, most popular RL algorithms, including PPO, which we use here, use complex function approximators that are hard to interpret. Therefore we use a method in explainable machine learning, integrated gradient (Sundararajan et al., 2017), to decompose the multi-decision output of the RL policy used in study 2 into a linear additive sum of attribution for each input context feature.

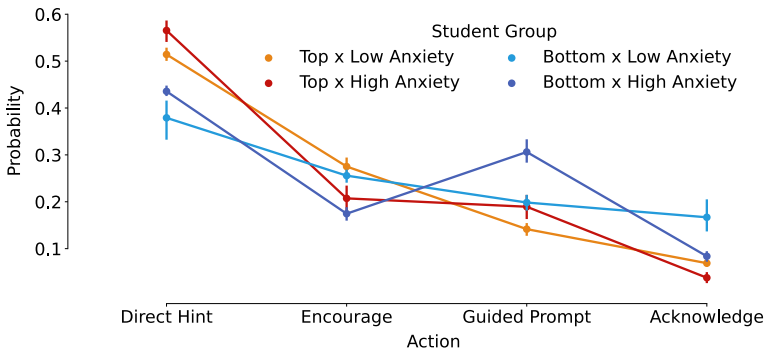
Table 3 shows that the feature importances computed for the policy selected from offline RL and deployed in the RL condition. Recall there are three primary categories of features used to select pedagogical strategies: static features of the learner, features about the stage of the learning activity, and features about the learner's interaction and performance during learning.

This analysis selected student's pretest score and their math anxiety score as the most influential contextual features on the AI guide's chosen response. Other student features had little to no effect. Figure 6 shows the probability of assigning actions for students from our distilled policy.

Students with higher pretest scores were more likely to receive direct hints: such students may require less of the productive struggle needed to learn new mathematics.

**Table 3** Feature importance calculated by the integrated gradient method. Numbers represent how on average, the feature (with its original value) will positively or negatively contribute to how our RL policy decides to increase or decrease the probability of choosing an action for the current student

RL policy action	Pre-test score (%)	Math anxiety (%)	Other features (%)
Pr (Direct hint)	+7.5	+8.8	−0.4
Pr (Acknowledgment)	−7.5	−7.4	+0.6
Pr (Encouragement)	+4.1	−8.2	+0.4
Pr (Guided prompt)	−4.2	+6.8	−0.7



**Fig. 6** The y-axis shows the probability of choosing the first action for each group of subjects, based on their pretest scores (Bottom (0-2), Top (6-8)), and math anxiety level (Low (9-13, corresponding to the bottom 25% percentile), and High (22-45, corresponding to the top 25% percentile)). Error bar shows 95% CI

Students with lower pretest scores may need more engaged practice, but those with high math anxiety may also perceive math as more effortful (Choe et al., 2019). Increasing the use of guided prompts may help support such students, as we observe in the policy instructional selections for low-performing, higher math anxiety students. These observed interactions between the multiple features describing student and context, and pedagogy choices, could inform expert analysis and support future hypothesis generation for learning sciences.

## 7 Discussion

Our work offers cautionary optimism on the potential role of reinforcement learning in optimizing pedagogical instructional policies. The personalized narrative AI guide may benefit students with the lowest pretest performance, without harming the performance of other learners. Indeed the average gain in scores for subjects with low (0-2) pretest scores was over 2 in both studies in the RL condition, which means the mean scores for such students at least doubled, in an assessment with 8 total points. Our results do not provide a definitive mechanism for this result, though the engagement scores suggest that the control condition was not engaging for subjects with low pretest scores. For such students, the RL

narrative AI guide condition yielded higher engagement, similar to those with higher pre-test scores. This is likely due to the RL AI guide, not the narrative, since prior work found narrative alone, with hints, yielded no benefit over no narrative and no AI guide in a volume learning task (Ruan et al., 2020).

Our encouraging result is consistent with limited prior work that personalized computer-assisted learning software may sometimes be similarly or only slightly more effective on average but may particularly benefit students who start with lower scores or take longer to complete problems (e.g. Shen et al. (2016); de Barros and Ganimian (2021)). Since the RL algorithm we used aims to maximize expected (test) outcomes, if differentiation within the available pedagogical supports can increase the outcome of any subgroups (without harming the outcomes of other subgroups), the algorithm should learn from data to provide such personalization. Our analysis did not find a significant benefit of RL over the control condition at the population level, though it is possible an effect would be observed with a larger sample size, or with different state feature representations, network architectures, or RL algorithms.

Across study 1 and 2 the comparison between the narrative RL condition and control conditions appear largely stable (Fig. 4), with a trend for the RL condition benefiting those with low pretest scores. This suggests an RL decision policy learned on one population can sometimes benefit other populations.

## 8 Conclusion

Our work was conducted on around 400 students, which is typically less than the number of third to fifth-graders in a school district, suggesting the feasibility of using this approach to quickly optimize digital learning environments. By combining reinforcement learning with explainable AI, this approach can provide new insights into the interaction of context and student learning that may prompt new research in learning sciences and has a high potential to help quickly identify and scale effective learning practices.

## Appendix A

### A real conversation example

To provide an indication of how well the RL AI guide performs after about 200 rounds of iterations, we present below a real conversation snippet between a user and the RL AI guide.

**Child:** Let's keep thinking.

**AI guide:** [Direct Hint] You can fit 3 chocolates along the height, 5 chocolates along the width, and 6 chocolates along the length of the box.

**Child:** Ok.

**AI guide:** [Guided Prompt] Hm... I suggest thinking about the concept of a unit cube and the volume of a unit cube. It may remind you of what you need in order to calculate the volume of the box!

*Child inputs 90 (correct answer) for the volume question.*

**Child:** We got another one right!

**AI guide:** [Acknowledgement]:)

**Child:** Do you think we multiply or divide.

**AI guide:** [Encouragement] I think I have some ideas! But I won't spoil it. Just let me know if you need a hint.

*Child inputs 180 (correct answer) for the weight question.*

**Child:** Also we got another one right!

**AI guide:** [Acknowledgement]:)

As can be seen from this example, the AI guide adopted all four types of teaching strategies: direct hint, guided prompt, encouragement, and acknowledgment. In particular, when the child asked "Do you think we multiply or divide.", instead of providing any meaningful help, the AI guide simply encouraged the child. This turns out to be an effective teaching strategy since the child successfully solved the problem without needing specific math-related help from the AI guide.

### The math learning activity and narrative

The central task of the learning platform is the completion of an online math learning activity focused on teaching 2 concepts related to volume and its knowledge components (measurement and multiplication). Similar to the task in Ruan et al. (2020), our math learning activity presented children with 6 tasks in order of increasing difficulty. All of these tasks were subparts of one overall problem (calculating if a box of chocolates weighed > 320 oz.) and results from prior tasks were used in later ones.

The first three tasks ask children to measure the length, width, and height of an on-screen box by dragging a bar of single-unit square chocolate along its edges. For the fourth task, children are asked to use these measurements to calculate the total number of chocolates that can fit in the box. In the fifth task, children are told that each chocolate weighs 0.5oz., the information they are required to use to help them calculate the total weight of the box. Finally, the sixth task asks children to determine if the box can be safely transported by a boat with a weight limit of 320oz.

Our AI guide support component replaces the remote human feedback support component used by Ruan et al. (2020). In addition, due to the constraints of the covid-19 pandemic situation at the time, children completed our math learning activity remotely through an online web app as opposed to in a physical lab setting. This means children complete the online activity asynchronously without the observation or interference of a researcher. We conducted a 10-minute video call with each guardian-child to confirm there was a child learner who intended to complete the task. During this video call, we emphasized to the subjects that they should complete the activity without the help of any outside resources, and guardians were asked to ensure their children completed the task without outside resources.

### AI guide support

During the math learning activity, each time the AI guide is sent a message, it can take one of several actions. 1) Provide an instructional hint. Hints are specific to the task the child is currently working on and are provided in a fixed order. Each time this action is taken, the next hint is provided, and when no hints are left for the current task, the AI guide sends an appropriate message. 2) Send acknowledgment. In this case, the agent decides that no action is appropriate; the AI guide acknowledges the child's message but otherwise



provides no assistance or encouragement (“:”). 3) Send encouragement. A random encouraging message from a predetermined list is sent to the user (for example, “You’re doing a great job. If we keep working like this, we’ll be done in no time!”). These messages were written to promote a growth mindset and excitement about the challenge of the problem without giving help to the problem itself. 4) Guided Prompt. As with normal hints, guided prompts are specific to the current task and are provided in a fixed order. In contrast to normal hints, the goal of guided prompts is to provide some assistance to children who do not need as much help as a standard hint provides (for example, “Try thinking about the concept of volume to solve this problem.”).

The AI guide only responds when spoken to with the exception of periodic “reminder” messages which remind the children that the AI guide is there (for example, “I think you’ve got this. But if you need help, just let me know!”). These messages are chosen randomly from a predetermined list. The goal is to provide children with social support as well as remind children to use the AI guide as a helpful resource if they become stuck. These reminders are sent every 120 s after user inactivity (including both speaking to the AI guide and interacting with the software). Additionally, the AI guide has a list of predicted responses that it ignores (such as “Okay”) or acknowledges with “You’re welcome!” (such as “Thanks”) to reduce noise from natural language responses that do not require one of the above actions.

In contrast to the experimental condition, there was no AI guide and no hint system present in the control condition.

The AI guide responds to input from the learner. There was an automated reminder for the child to engage if no prior interactions had happened during the 120 s. The automated instructional policy was trained using reinforcement learning. For the first phase, the reward model uses  $\alpha = 0.01$ ,  $\beta = 0.1$ ,  $\gamma = 0.3$ , based on the hyperparameter choices of prior work (Bassen et al., 2020) and our earlier simulations simulation.

All the hints and message templates were written and uploaded through an easy-to-use teacher-facing dashboard (see Fig. 7) by educators and designers without prior background in machine learning.

## Distribution of grade and pre-test scores in treatment and control

**In study 1**, 339 participants in grades 3–5 were recruited through Twitter, NextDoor, userinterview.com, school mailing lists, and word of mouth. Children came from 263 different schools. Of 339 participating children, 172 were boys and 167 were girls. 114 were in grade 3, 114 were in grade 4, and 111 were in grade 5. Children were randomly assigned to one of the two systems based on a predetermined ratio: 70 children used the control system and 269 children used the system with RL AI guide-mediated guidance. Gender and grade were balanced across the two conditions.

There is no significant difference between the treatment and control group in study 1 on pre-test score (Cohen’s  $d = -0.235$ , two-sample Wilcoxon rank test  $W = 10782$ ,  $p$ -value = 0.058) as well as grade (Cohen’s  $d = 0.008$ , two-sample Wilcoxon rank test  $W = 9371.5$ ,  $p$ -value = 0.9502).

**In study 2**, 35 participants were recruited using userinterview.com and childrenhelpingscience.com. There is no significant difference between the treatment and control group in study 2 on pre-test score (Cohen’s  $d = -0.262$ , two-sample Wilcoxon rank test  $W = 175.5$ ,  $p$ -value = 0.4634), as well as grade (Cohen’s  $d = 0.282$ , two-sample Wilcoxon rank test  $W = 136.5$ ,  $p$ -value = 0.5665).

**Fig. 7** Interface for teachers to write hints and prompts

## Details on repeated post-test taking in logged data

The software did not explicitly check for students repeating the pretest or post-test, and in our post-analysis, we found a few students took either the pretest or post-test multiple times. The logging software only recorded the score of the final time the student took the test. For this reason, we only analyzed students who took the pretest and post-test once. In study 1, this resulted in 68 (out of 70) students in the control condition being kept in the analysis (only 2 students took either the pretest or posttest twice) and 258 (out of 269) students in the RL condition. In study 2, 18 (out of 19) students in the control condition, and 17 (out of 18) students in the RL condition were included in the analysis. We computed our results after removing these duplicate entries.

## Report on time spent between control and RL condition

On average students do often spend longer on the RL narrative condition task than in the control condition: Fig. 5. This was consistent for students across all three groups of pretest performance, and the difference in time spent between the two conditions was largely similar for all three groups. As it was only students in the low pretest group that seem to have a

**Table 4** Average time (secs) spent on different parts of the tutoring session

	Online RL study 1		Distilled policy study 2	
	Control	Narrative AI	Control	Narrative AI
Pre-test	410.5	424.3	478.2	560.5
Task	592.3	2334.8	542.1	1978.9
Post-test	511.1	290.7	246.5	300.9

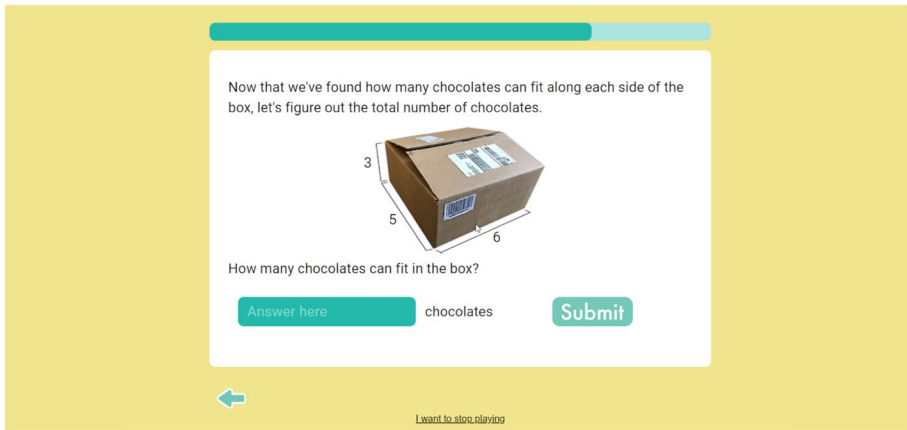
significant benefit from the RL condition, it seems unlikely that time on task is the primary reason for improved performance in the RL narrative condition.

We report the time spent on the pretest, task, and post-test (assessment), in each control and experiment, in both study 1 and study 2 (see Table 4). We conduct a two-sample Wilcoxon rank test on all pairs (between study 1 and study 2). We find no significance between the two studies.

## Engagement

In study 1, students with low pretest scores (scores 0-2) had an average engagement score of 2.67 (N=14, standard error = 0.23) in the control condition and an average engagement score of 3.29 (N=40, standard error = 0.11) in the RL narrative AI guide condition. In study 1, students with medium pretest scores (scores 3-5) had an average engagement score of 3.43 (N=15, standard error = 0.12) in the control condition, and an average engagement score of 3.36 (N=105, standard error = 0.05) in the RL narrative AI guide condition. In study 1, students with high pretest scores (scores 6-8) had an average engagement score of 3.24 (N=38, standard error = 0.07) in the control condition, and an average engagement score of 3.48 (N=108, standard error = 0.04) in the RL narrative AI guide condition. Three subjects in study 1 did not complete the engagement survey.

In study 2, students with low pretest scores (scores 0-2) had an average engagement score of 2.71 (N=5, standard error = 0.12) in the control condition and an average engagement score of 3.28 (N=7, standard error = 0.23) in the RL narrative AI guide condition. In study 2, students with medium pretest scores (scores 3-5) had an average engagement score of 3.21 (N=6, standard error = 0.23) in the control condition, and an average engagement score of 3.28 (N=6, standard error = 0.28) in the RL narrative AI guide condition. In study 2, students with high pretest scores (scores 6-8) had an average engagement score of 3.45 (N=7, standard error = 0.17) in the control condition, and an average engagement score of 3.29 (N=4, standard error = 0.23) in the RL narrative AI guide condition.



**Fig. 8** The control interface where the child attempts to solve the math problem without any narrative or help from a chatbot. The child can click on “I want to stop playing” to quit the activity at any time

## Implementation details

The platform consists of three major parts: a user-facing interactive website (Fig. 8 for control and Fig. 1 for AI guide), an admin dashboard (Fig. 7), and an AI guide server. Both the website and the dashboard were created using Web technologies, including ReactJS (Facebook, 2019) and TypeScript (Microsoft, 2019). The Python-based AI guide was hosted on an AWS server and used Flask (Projects, 2010) as its API gateway to expose essential functions. The interactive website communicated with a GraphQL (Hasura, 2019) API endpoint backed by Hasura Engine (Hasura, 2019) and PostgreSQL (Postgres, 2019). The stored user conversation data was reflected in real-time on the admin console, where researchers could view the chat history and modify message templates. All user data was uploaded to the backend by Google App Script upon the completion of the user’s session.

Questionnaires and quizzes were created using Google Forms, and we used HTML iframe to embed Google Forms into the website to automatically process the form responses so as to enable real-time RL. When users interact with the AI guide, the observation space is calculated in real-time, and the AI guide performs action selection to reply to users. When users completed the post-quiz, their answers were converted to vector inputs and fed into the RL AI guide in real-time, which triggered a webhook to request the AI guide server to update its model accordingly. A complete diagram showing the interaction between the user and the RL agent is displayed in Fig. 2.

**CONTACT INFORMATION:**

Questions:  
If you have any questions, concerns or complaints about this research, its procedures, risks, and benefits, contact a member from the research team: Smart Primer team ([smartest.primer@gmail.com](mailto:smartest.primer@gmail.com)), Sherry Ruan ([ssruan@stanford.edu](mailto:ssruan@stanford.edu)), or the protocol Director James Landay ([landay@stanford.edu](mailto:landay@stanford.edu)).

Independent Contact:  
If you are not satisfied with how this study is being conducted, or if you have any concerns, complaints, or general questions about the research or your rights as a participant, please contact the Stanford Institutional Review Board (IRB) to speak to someone independent of the research team at (650)-723-2480 or toll-free at 1-866-680-2906. You can also write to the [Stanford IRB, Stanford University, 1705 El Camino Real, Palo Alto, CA 94306](#).

\_\_\_\_\_

Email address

\_\_\_\_\_

\_\_\_\_\_

Parent's Full Name

\_\_\_\_\_

\_\_\_\_\_

Check the box below to indicate "Yes"

I consent to my child's participation in this study.

**>>**


Powered by Qualtrics 

Fig. 9 Consent page 1

**CONTACT INFORMATION:**

## Questions:

If you have any questions, concerns or complaints about this research, its procedures, risks, and benefits, contact a member from the research team:

Smart Primer team ([smartest.primer@gmail.com](mailto:smartest.primer@gmail.com)), Sherry Ruan ([ssruan@stanford.edu](mailto:ssruan@stanford.edu)), or the protocol Director James Landay ([landay@stanford.edu](mailto:landay@stanford.edu)).

## Independent Contact:

If you are not satisfied with how this study is being conducted, or if you have any concerns, complaints, or general questions about the research or your rights as a participant, please contact the Stanford Institutional Review Board (IRB) to speak to someone independent of the research team at (650)-723-2480 or toll-free at 1-866-680-2906. You can also write to the [Stanford IRB, Stanford University, 1705 El Camino Real, Palo Alto, CA 94306](#).

Email address

Parent's Full Name

Check the box below to indicate "Yes"

 I consent to my child's participation in this study.

&gt;&gt;

Powered by Qualtrics **Fig. 10** Consent page 2

## Stanford

**This form is for the child who will be completing the study to read and fill out.**

### 1. What will happen to me in this study?

We are studying how much kids like playing with our [Smart Primer](#) website. You will get to interact with an application on the website. There will be a survey and a quiz before and after the study. We will not reveal your quiz performance or the survey responses to anyone else.

### 2. Can anything bad happen to me?

You will be playing with our website at your own home from your home computer. Nothing bad will happen to you from playing. You may feel tired during the study, and it's ok to take breaks. If you don't want to play at any time, you can tell your parents or us to opt out.

### 3. Can anything good happen to me?

You will be playing with our website and may learn something new.

### 4. What happens if I get hurt?

Your parents will be there and they will make sure you get taken care of.

### 5. Who can I talk to about the study?

If you have any questions about the study or any problems to do with the study you can contact the student in charge of the study: Sherry Ruan ([ssruan@stanford.edu](mailto:ssruan@stanford.edu)), or the protocol Director James Landay at [landay@stanford.edu](mailto:landay@stanford.edu). If you have questions about the study but want to talk to someone else who is not a part of the study, you can call the Stanford Institutional Review Board (IRB) at (650)-723-5244 or toll-free at 1-866-680-2906. You can also write to the [Stanford IRB, Stanford University, 1705 El Camino Real, Palo Alto, CA 94306](#).

### 6. What if I do not want to do this?

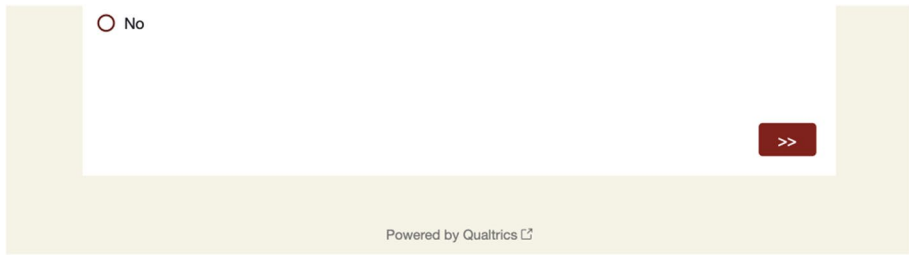
You can choose whether or not you want to do the study. If you don't want to, just say so and we can stop at any time. If you wish to end the study when playing the app, you can click on the opt out button to exit the app. If you wish to stop during the survey, you can choose not to submit the survey. Your compensation will not be affected if you opt out of the task but finish the survey.

Full Name

By clicking "Yes", you agree to participate in this study. If not, click "No".

Yes

Fig. 11 Consent page 3



**Fig. 12** Consent page 4

**Acknowledgements** This work was supported in part by a Stanford Hoffman-Yee grant, a gift from the TAL education group, and a NSF #2112926 grant.

**Author Contributions** SR, AN, WS, JH, JZ, MG, YL, KDN, CW, RY, JL, and EB conducted research. SR, AN, WS, YL, JL, and EB designed research. SR, AN, WS, and EB performed the analysis. SR, WS, AN, and EB wrote the manuscript.

**Funding** This work was supported in part by a Stanford Hoffman-Yee grant, a gift from the TAL education group, and a NSF #2112926 grant.

**Availability of data and material** Under the IRB approval, which was given for our studies which involved interacting with children, the data will not be available publicly. We can share anonymized data for replication and verification purposes after receiving a reasonable request.

**Code availability** The code used is available to the public via the links provided in the main text.

## Declarations

**Conflict of interest** The authors declare no conflict of interest.

**Ethical approval** The two experimental studies were approved by the Stanford Institutional Review Board.

**Consent to participate** A copy of the subject consent forms are provided in Figs. 9, 10, 11 and 12.

**Consent for publication** Not relevant.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Bassen, J., Balaji, B., Schaarschmidt, M., Thille, C., Painter, J., Zimmaro, D., Games, A., Fast, E., & Mitchell, J. C. (2020). Reinforcement learning for the adaptive scheduling of educational activities. In *CHI*, pp. 1–12
- Beal, C. R., Arroyo, I. M., Cohen, P. R., & Woolf, B. P. (2010). Evaluation of animalwatch: An intelligent tutoring system for arithmetic and fractions. *Journal of Interactive Online Learning*, 9(1), 1–14.
- Bird, S., Klein, E., & Loper, E. (2009). Natural language processing with Python: Analyzing text with the natural language toolkit. O'Reilly Media, Inc.



- Carey, E., Hill, F., Devine, A., & Szucs, D. (2017). The modified abbreviated math anxiety scale: A valid and reliable instrument for use with children. *Frontiers in Psychology*, 8, 11. <https://doi.org/10.3389/fpsyg.2017.00011>
- Carey, E., Hill, F., Devine, A., & Szűcs, D. (2017). The modified abbreviated math anxiety scale: A valid and reliable instrument for use with children. *Frontiers in Psychology*, 8, 11.
- Chi, M., VanLehn, K., Litman, D., & Jordan, P. (2011). Empirically evaluating the application of reinforcement learning to the induction of effective and adaptive pedagogical strategies. *User Modeling and User-Adapted Interaction*, 21(1), 137–180.
- Choe, K. W., Jenifer, J. B., Rozek, C. S., Berman, M. G., & Beilock, S. L. (2019). Calculated avoidance: Math anxiety predicts math avoidance in effort-based decision-making. *Science Advances*, 5(11), 1062.
- Corbett, A. (2001). Cognitive computer tutors: Solving the two-sigma problem. In *International Conference on User Modeling*, pp. 137–147. Springer
- de Barros, A., & Ganimian, A.J. (2021). Which students benefit from personalized learning? Experimental evidence from a math software in public schools in India
- Dietz, G., Pease, Z., McNally, B., & Foss, E. (2020). Giggle gauge: a self-report instrument for evaluating children's engagement with technology. In *Proceedings of the Interaction Design and Children Conference*, pp. 614–623
- Doroudi, S., Alevan, V., & Brunskill, E. (2019). Where's the reward? *International Journal of Artificial Intelligence in Education*, 29(4), 568–620.
- Facebook: Facebook React. <https://github.com/facebook/react>. Accessed: 2019-08-20 (2019)
- Hasura: Hasura GraphQL. <https://github.com/hasura/graphql-engine>. Accessed: 2019-08-20 (2019)
- Hendrycks, D., & Gimpel, K. (2016). Gaussian error linear units (gelus). arXiv preprint [arXiv:1606.08415](https://arxiv.org/abs/1606.08415)
- Levine, S., Finn, C., Darrell, T., & Abbeel, P. (2016). End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1), 1334–1373.
- Liu, Y., Swaminathan, A., Agarwal, A., & Brunskill, E. (2020). Off-policy policy gradient with stationary distribution correction. In *Uncertainty in Artificial Intelligence*, pp. 1180–1190. PMLR
- Lundberg, S.M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems* 30
- Mandel, T., Liu, Y.-E., Levine, S., Brunskill, E., & Popovic, Z. (2014). Offline policy evaluation across representations with applications to educational games. In *AAMAS*, vol. 1077
- Metelli, A.M., Papini, M., Faccio, F., & Restelli, M. (2018). Policy optimization via importance sampling. arXiv preprint [arXiv:1809.06098](https://arxiv.org/abs/1809.06098)
- Microsoft: Microsoft TypeScript. <https://github.com/microsoft/TypeScript>. Accessed: 2019-08-20 (2019)
- Nickow, A., Oreopoulos, P., & Quan, V. (2020). The impressive effects of tutoring on pre-k-12 learning: A systematic review and meta-analysis of the experimental evidence. working paper 27476. National Bureau of Economic Research
- Nie, A., Flet-Berliac, Y., Jordan, D., Steenbergen, W., & Brunskill, E. (2022). Data-efficient pipeline for offline reinforcement learning with limited data. *Advances in Neural Information Processing Systems*, 35, 14810–14823.
- Park, H. W., Grover, I., Spaulding, S., Gomez, L., & Breazeal, C. (2019). A model-free affective reinforcement learning approach to personalization of an autonomous social robot companion for early literacy education. *AAAI*, 33, 687–694.
- Pomerleau, D. (1990). Rapidly adapting artificial neural networks for autonomous navigation. *NeurIPS* 3
- Postgres: Postgres. <https://www.postgresql.org/>. Accessed: 2019-08-20 (2019)
- Projects, T.P.: Flask. <https://flask.palletsprojects.com/>. Accessed: 2021-03-03 (2010)
- Rowe, J.P., Lester, J.C. (2015). Improving student problem solving in narrative-centered learning environments: A modular reinforcement learning framework. In *International Conference on Artificial Intelligence in Education*, pp. 419–428. Springer
- Ruan, S., He, J., Ying, R., Burkle, J., Hakim, D., Wang, A., Yin, Y., Zhou, L., Xu, Q., AbuHashem, A., Dietz, G., Murnane, E.L., Brunskill, E., & Landay, J.A. (2020). Supporting children's math learning with feedback-augmented narrative technology. In *IDC*, pp. 567–580. <https://doi.org/10.1145/3392063.3394400>.
- Sammut, C., Hurst, S., Kedzier, D., & Michie, D. (1992). Learning to fly. In *Machine Learning Proceedings 1992*, pp. 385–393. Elsevier.
- Schaarschmidt, M., Mika, S., Fricke, K., & Yoneki, E. (2019). Rlgraph: Modular computation graphs for deep reinforcement learning. In *Proceedings of the 2nd Conference on Systems and Machine Learning (SysML)*
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint [arXiv:1707.06347](https://arxiv.org/abs/1707.06347)
- Shen, S., & Chi, M. (2016). Reinforcement learning: the sooner the better, or the later the better? In *UMAP*, pp. 37–44.

- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419), 1140–1144.
- Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pp. 3319–3328. PMLR
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221.
- Zhou, G., Azizsoltani, H., Ausin, M. S., Barnes, T., & Chi, M. (2022). Leveraging granularity: Hierarchical reinforcement learning for pedagogical policy induction. *International Journal of Artificial Intelligence in Education*, 32(2), 454–500.
- Zhou, G., Azizsoltani, H., Ausin, M.S., Barnes, T., & Chi, M. (2019). Hierarchical reinforcement learning for pedagogical policy induction. In *Artificial Intelligence in Education: 20th International Conference, AIED 2019, Chicago, IL, USA, June 25–29, 2019, Proceedings, Part I 20*, pp. 544–556. Springer

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Authors and Affiliations

Sherry Ruan<sup>2</sup> · Allen Nie<sup>1</sup> · William Steenbergen<sup>2</sup> · Jiayu He<sup>2</sup> · J. Q. Zhang<sup>2</sup> · Meng Guo<sup>1</sup> · Yao Liu<sup>2</sup> · Kyle Dang Nguyen<sup>1</sup> · Catherine Y. Wang<sup>1</sup> · Rui Ying<sup>2</sup> · James A. Landay<sup>1</sup> · Emma Brunskill<sup>1</sup> 

✉ Emma Brunskill  
ebrun@cs.stanford.edu

Sherry Ruan  
sherry.s.ruan@gmail.com

Allen Nie  
anie@stanford.edu

William Steenbergen  
w.a.t.steenbergen@gmail.com

Jiayu He  
jiauhe@alumni.stanford.edu

J. Q. Zhang  
jjequanz@stanford.edu

Meng Guo  
mguo19@stanford.edu

Yao Liu  
yao.liu.chn@gmail.com

Kyle Dang Nguyen  
kylen@stanford.edu

Catherine Y. Wang  
cyw339@stanford.edu

Rui Ying  
ruiying@stanford.edu

James A. Landay  
landay@stanford.edu

<sup>1</sup> Computer Science, Stanford University, Stanford, CA 94305, USA

<sup>2</sup> Work done while at Stanford University, Stanford, CA 94305, USA