

Adaptively Scaffolding Cognitive Engagement with Batch Constrained Deep Q-Networks

Fahmid Morshed Fahid¹[0000-0002-4802-3979], Jonathan P. Rowe¹[0000-0003-2038-9239], Randall D. Spain¹, Benjamin S. Goldberg²[0000-0002-1396-9887], Robert Pokorny³, and James Lester¹

¹ North Carolina State University, Raleigh NC 27695, USA
{ffahid, jprowe, rdspain, lester}@ncsu.edu

² U.S. Army CCDC - STTC, Orlando FL 32826, USA
benjamin.s.goldberg.civ@mail.mil

³ Intelligent Automation Inc. Rockville MD 20855, USA
bapokorny@gmail.com

Abstract. Scaffolding student engagement is a central challenge in adaptive learning environments. The ICAP framework defines levels of cognitive engagement with a learning activity in terms of four different engagement modes—Interactive, Constructive, Active, and Passive—and it predicts that increased cognitive engagement will yield improved learning. However, a key open question is how best to translate the ICAP theory into the design of adaptive scaffolding in adaptive learning environments. Specifically, should scaffolds be designed to require the highest levels of cognitive engagement (i.e., Interactive and Constructive modes) with every instance of feedback or knowledge component? To answer this question, in this paper we investigate a data-driven pedagogical modeling framework based on batch-constrained deep Q-networks, a type of deep reinforcement learning (RL) method, to induce policies for delivering ICAP-inspired scaffolding in adaptive learning environments. The policies are trained with log data from 487 learners as they interacted with an adaptive learning environment that provided ICAP-inspired feedback and remediation. Results suggest that adaptive scaffolding policies induced with batch-constrained deep Q-networks outperform heuristic policies that strictly follow the ICAP model without RL-based tailoring. The findings demonstrate the utility of deep RL for tailoring scaffolding for learner cognitive engagement.

Keywords: Deep Reinforcement Learning, Cognitive Engagement, ICAP, Adaptive Learning Environments.

1 Introduction

Adaptive learning environments provide scaffolding in the form of hints, feedback and remediation to improve learning experiences. Scaffolds offer temporary support to students as they learn, which is gradually faded as students gain knowledge and achieve mastery. Designing effective scaffolds is challenging. Determining how and when to

deliver scaffolding in different situations is critical to enabling effective learning experiences [22]. A key factor in adaptive scaffolding is the cognitive engagement of learners. Chi and Wylie [7] describe *cognitive engagement* as an “active learning” process that involves higher-order thinking (e.g., analyzing, synthesizing, evaluating). Ideally, adaptive scaffolding is designed to optimize students’ cognitive engagement, and by extension, enhance learning outcomes.

The ICAP framework provides a taxonomy for categorizing different modes of cognitive engagement: Interactive, Constructive, Active, and Passive [7]. ICAP predicts that learning activities requiring higher levels of cognitive engagement (e.g., peer dialogue, writing a summary) yield improved learning outcomes compared to activities that involve lower levels of cognitive engagement (e.g., listening passively, highlighting text). There is strong evidence in support of the ICAP theory, and it has been used to guide the design of lesson plans [6] and adaptive learning technologies [20], but it is less clear how to translate ICAP into the design of individual scaffolds. High levels of cognitive engagement require time and student motivation. A direct translation of ICAP may not be optimal for every scaffold and knowledge component in an adaptive learning environment. This raises a natural question: should ICAP be operationalized by *adaptively* scaffolding cognitive engagement, eliciting higher-order thinking at key moments with the aim of enhancing overall learning outcomes, and, if so, how should we devise models for adaptively scaffolding cognitive engagement?

Recent years have seen growing interest in using reinforcement learning (RL) to devise policies for scaffolding student learning in adaptive learning environments [5, 8]. Deep RL, which combines RL and deep neural networks, has shown particular promise for this task [1-3, 31]. Several studies have shown that deep RL techniques yield effective pedagogical models in adaptive learning environments [3, 15]. However, previous work has not systematically investigated methods for adaptively scaffolding cognitive engagement with deep RL techniques.

In this paper, we introduce a data-driven pedagogical modeling framework based on batch constrained deep Q-networks, a type of deep RL method, to induce policies for scaffolding cognitive engagement in adaptive learning environments. The policies drive ICAP-based feedback and remediation following instructional videos and embedded assessments in a learning environment for training operational command skills. The policies are induced using interaction log data from 487 learners as they engaged with the adaptive learning environment. We compare scaffolding policies induced with batch constrained deep Q-networks with heuristic policies that strictly follow the ICAP model without RL-based tailoring.

2 Related Work

RL provides a natural framework for inducing data-driven scaffolding models to improve student learning experiences. Wang conducted a study with 30 students learning software development concepts in a dialogue-based tutoring system and found that an RL-based teaching assistant was able to learn from its teaching experience and continuously improve its teaching strategies online [29]. Georgila and colleagues [10] found

that RL-based models fostered increased confidence among learners through adaptive scaffolding to support the development of interpersonal skills. Their results suggest that the induced policies matched, or outperformed heuristic scaffolding models designed by human experts. Similar findings have been reported in other studies and learning environments investigating RL-based pedagogical models [23, 33].

Over the past several years, deep RL techniques have shown significant promise for inducing scaffolding policies in adaptive learning environments. For example, Wang and colleagues [30] found that adaptively scaffolding student learning in a narrative-centered learning environment for middle school microbiology using deep RL models trained with simulated students outperformed baseline methods. Additional work has investigated offline deep RL methods, where RL models are trained with previously collected data rather than simulations to induce scaffolding policies. For example, Azizoltani and colleagues [4] found that inferring immediate rewards using Gaussian process estimation to train offline deep RL-based pedagogical models can significantly improve learning gains in students. To date, deep RL techniques have not been used to induce policies for adaptively scaffolding cognitive engagement with ICAP-inspired feedback and remediation in adaptive learning environments.

The ICAP framework predicts that as students become more actively engaged with learning materials, moving from passive to active to constructive to interactive behaviors, their learning will increase. Support for the ICAP framework has been found in a number of studies [18, 20, 32]. Mitrovic et al. [20] found that using interactive visualization and prompts to enforce constructive engagement in a video-based learning environment led to high levels of confidence and lower levels of frustration during the learning episode compared to students who engaged in passive learning behaviors. Few studies have investigated how adaptive ICAP-inspired scaffolding applied at a step-based or micro-loop level in adaptive learning environments supports student learning [26].

3 Dataset

To induce data-driven pedagogical models for delivering ICAP-inspired feedback and remediation, we utilize log data collected from an online study involving 487 learners (54% male, 42% female) recruited through Amazon Mechanical Turk who interacted with an adaptive learning environment for training operational command skills. The learning environment was built using the Generalized Intelligent Framework for Tutoring (GIFT), an open-source domain-independent framework for designing, deploying, and evaluating adaptive learning technologies [25]. The learning environment includes a series of instructional videos that cover core concepts and principles associated with operational command. Following each video, learners answered a series of multiple-choice questions. An incorrect response to a question prompted the learning environment to deliver ICAP-inspired feedback and remediation that required the learner to either (1) passively re-read a transcription of the video that was just presented in the lesson video, (2) re-read the transcription of the video and actively highlight the portion of text that answered the recall question that was just missed, or (3) re-read the video transcription and constructively summarize the answer to the question in their own

words. The learning environment did not have built-in support for the interactive mode of engaging with feedback and remediation, so that component of ICAP was omitted. The active and constructive remediation prompts included expert highlighting/summaries that asked students to self-evaluate the accuracy of their responses. The learning environment also included a “no remediation” prompt that provided learners with a simple feedback message stating they incorrectly answered the question.

After completing a remediation exercise, learners were presented again with the previously attempted question. If they answered the question correctly, then they advanced to the next question or video lesson. Learners continued to receive remediation, potentially of different types, until they correctly answered the question. The learning environment utilized a random policy to determine the type of ICAP-inspired remediation learners received each time they missed a question (irrespective of their number of attempts), although a software error caused passive and no remediation instances to be under-sampled.

In all, learners completed 39 embedded assessments, which were distributed across four units that typically take 1-2 hours in total to complete. The adaptive learning environment also included a set of web-based surveys designed to collect demographic information and a set of pre- and post-test items that measured student learning as a result of completing the course.

The resulting dataset included a total of 4,998 instances of ICAP-inspired feedback and remediation. On average, learners received 10 instances of remediation while completing the course ($SD=12.7$; $min=1$, $max=113$). Table 1 summarizes the distribution of remediation instances encountered throughout the course. A paired t-test showed that the pre-test scores ($M=4.18$, $SD=2.30$, $min=0$, $max=11$) and the post-test scores ($M=8.32$, $SD=2.96$, $min=0$, $max=12$) were significantly different ($p<0.001$), implying the adaptive learning environment improved knowledge of operational command concepts and skills among the participants.

Table 1. Distribution of ICAP-based remediation instances.

Remediation	Total	Chapter 1	Chapter 2	Chapter 3	Chapter 4
None	470 (9.40%)	155 (3.10%)	141 (2.82%)	141 (2.82%)	33 (0.66%)
Passive	445 (8.90%)	145 (2.90%)	136 (2.72%)	127 (2.54%)	37 (0.74%)
Active	2074 (41.50%)	684 (13.65%)	587 (11.74%)	639 (12.79%)	166 (3.32%)
Constructive	2009 (40.20%)	626 (12.53%)	606 (12.12%)	611 (12.22%)	166 (3.32%)

4 Adaptive Scaffolding with Batch Constrained Deep Q-Networks

In this section, we present a deep RL framework for creating policies to scaffold cognitive engagement in adaptive learning environments. Specifically, we describe our deep RL-based pedagogical model architecture, our approach to formalizing adaptive scaffolding as a Markov decision process, and a pair of metrics for evaluating policies for the delivering ICAP-inspired feedback and remediation.

4.1 Deep RL-based Pedagogical Model Architecture

To devise data-driven policies for adaptively scaffolding students' cognitive engagement, we used deep Q learning, a type of RL technique that leverages deep neural network-based function approximation to represent the values of input states [21]. Q-learning is a model-free RL algorithm where the goal is to learn an optimal policy π^* based on the optimal action-value function $Q^*(s, a)$ estimated from sample data without use of an explicit model of the task environment [27]. Starting from state $s \in \mathbb{S}$ and taking action $a \in \mathbb{A}$ while getting reward $r \in \mathbb{R}$, the Q values are defined as the expected cumulative reward following a policy π that generates a set of actions at each successive state. To reduce extrapolation errors often seen in offline RL, we utilized batch constrained deep Q-networks [9].

Deep Q-networks (DQNs) follow an off-policy learning approach that involves iteratively sampling from a finite experience buffer to greedily estimate Q values according to the Bellman equation. A loss function (Equation 1) is defined to train a deep neural network to estimate the model's Q-values:

$$Loss(\theta) = \mathbb{E}[(y - Q(s, a; \theta))^2] \quad (1)$$

where θ represents the set of weights in the neural network.

A variant of DQNs is the *Double DQN*, which uses two separate networks to reduce overestimation bias in the DQN by separating the action selection and action evaluation components of the model [11]. This provides improved stabilization and convergence while the model is trained. In Double DQNs, two neural networks with identical architectures are used, namely, the target network ($\bar{\theta}$) and the online network (θ). The online network is trained on every iteration while the target network is frozen for a fixed number of iterations. During training, the online network is used to select the next action $a' \in \mathbb{A}$ based on the next state $s' \in \mathbb{S}$, and the target network is used to evaluate the Q value of the action:

$$y = r + \gamma Q(s', \argmax_{a' \in \mathbb{A}} Q(s', a'; \bar{\theta}); \bar{\theta}) \quad (2)$$

where $\gamma \in [0, 1]$ is the discount factor that controls the contribution of future rewards.

DQNs are often used with an *experience replay buffer* to keep track of a finite set of recent training observations [21]. During training, transitions are sampled from the buffer randomly. Prioritized experience replay [24] prioritizes the sampling of transitions based upon the current temporal difference errors. This additional priority makes the network more data efficient [12] by ensuring quick convergence. Priority is calculated as follows:

$$t_{priority} = |\{r + \gamma Q(s', \argmax_{a' \in \mathbb{A}} Q(s', a'; \bar{\theta}); \bar{\theta})\} - \{Q(s, a; \theta)\}|^\omega \quad (3)$$

Here, $t_{priority}$ is the priority of a transition t and ω is a hyperparameter.

In batch RL, also known as offline RL, the experience replay buffer remains fixed. This approach is often necessary in RL applications in adaptive learning environments, where a training corpus is collected from students prior to employing RL, and additional data collection is not feasible during the RL process. With limited data, DQNs often

suffer from divergence issues due to extrapolating Q values outside of the data distribution. Batch constrained DQNs restrict such extrapolation errors by only allowing actions that are evident in the available data using a probabilistic sampling technique [9]:

$$y = r + \gamma Q(s', \text{argmax}_{a' | (a' | s') / \max_a \pi_b(\bar{a} | s') > \tau} Q(s', a'; \theta); \bar{\theta}) \quad (4)$$

Here, π_b is the policy used to collect the data and τ is a probability threshold.

Deep neural networks within batch constrained DQNs can be implemented using different neural architectures. In this work, we implement two commonly used architectures: fully connected (FC) layers and long short-term memory (LSTM) layers. In FC layers, each neuron is a perceptron that calculates a weighted sum of the input units to produce an output value through an activation function. All inputs are connected to all neurons in the first layer, all of the output of the first layer is fully connected to input neurons of the second layer, and so on until the final output layer.

LSTMs are a specialized version of recurrent neural networks that use long term temporal dependencies to avoid common issues in neural networks such as the vanishing and exploding gradient problems [13]. An LSTM unit consists of a memory cell state and three gates: a forget gate, an input gate, and an output gate. These pieces together control the flow of information during model training. Notably batch constrained DQNs with LSTM networks support sequential input representations, which enables them to keep track of (and forget) previous inputs and hidden states.

4.2 States, Actions and Reward

To formalize the task of inducing a policy for scaffolding cognitive engagement in an adaptive learning environment, we defined a Markov decision process, which involves controlling a set of actions \mathbb{A} based on some state $s \in \mathbb{S}$ to optimize the accumulation of reward $r \in \mathbb{R}$. Markov decision process provide a standardized mathematical representation for RL tasks. We define the state (\mathbb{S}), action (\mathbb{A}) and reward (\mathbb{R}) components of the Markov decision process as follows.

State (\mathbb{S}). We devise the state representation by extracting 31 features from learners' log data, which are divided into 3 groups: (1) survey features, (2) video playback features, and (3) remediation engagement features. The survey features include gender, age, education level, content familiarity, domain interest, and pre-test score on a content knowledge assessment. Four video playback features are extracted: time spent on the last video, average time spent on videos, whether learners received automated feedback about their time spent on the last video, and time spent on the feedback. Twenty remediation features are extracted: the previous type of remediation delivered, the total number of remediation instances delivered for each ICAP category, the average time spent engaged in the remediation activity, the average time spent on all previous remediation activities, and features reflecting how long learners spent answering the recall questions. We normalize each feature to range between $[0,1]$ to improve stability when training the DQNs. Batch constrained DQNs require a discrete state space to calculate the probability of each state-action pair. Therefore, we cluster the set of state-action pairs into 5 groups using k-means clustering. We visually select the number of clusters using

the distortion elbow method [17]. These clusters are used when sampling for the batch constrained DQN action probabilities.

Action (A). At each pedagogical decision point (i.e., after a missed question), there are 4 possible actions (i.e., ICAP-based remediations) that can be selected: constructive (re-read the video transcription and summarize), active (re-read the video transcript and highlight relevant section), passive (re-read the video transcript), or no remediation.

Reward (R). Each participant completed a 12-item pre- and post-test to assess content knowledge about operational command concepts. We use learners' pre- and post-test scores to calculate normalized learning gains (NLG) associated with each sequence of ICAP-inspired remediation instances [19]. Note that, for each student, we will only have a single NLG value at the end of their episode, which is the delayed reward [4]. Our reward value is a real number and ranges between 0 to 100.

4.3 Evaluation Metrics

In batch constrained DQNs, the optimal policy (π^*) usually has a significantly different distribution of state-action pairs than the behavioral policy (π_b) that was used to collect the training data. Performing RL policy evaluation with data collected under a different policy is known as *off-policy evaluation* [28]. We use two evaluation metrics, Expected Cumulative Reward and Doubly Robust.

Expected cumulative reward (ECR) computes the average expected reward associated with a particular policy π beginning at the initial state $s_i \in \mathbb{S}$ in a given dataset D of the RL task. Specifically, ECR reports the average Q value over all initial states as follows:

$$ECR(D) = \frac{1}{N} \sum_{i=1}^N \max_{a \in \mathbb{A}} Q(s_i, a) \quad (5)$$

where N is the total number of episodes and s_i is an initial state for the i^{th} episode.

Doubly Robust (DR) evaluation [14] is an alternative technique that combines the low variance of importance sampling estimation and the low biases of model-based estimation into a single metric according to the following equation:

$$DR(D) = \frac{1}{N} \sum_{i=1}^N \sum_{t=0}^{\infty} \gamma^t \prod_{l=0}^t \frac{\pi(a_l|s_l)}{\pi_b(a_l|s_l)} \left(R_t^i - Q(s_t^i, a_t^i) \right) + \gamma^t \prod_{l=0}^{t-1} \frac{\pi(a_l|s_l)}{\pi_b(a_l|s_l)} V(s_t^i) \quad (6)$$

The Q function and the value function (V) are based on a given policy π . Doubly Robust provides unbiased estimates if a model is accurate and/or provides low variance estimates if the behavior policy is known.

5 Results and Discussion

To devise deep RL-based policies for adaptively scaffolding cognitive engagement, we created pedagogical models with batch constrained double DQNs (BCQs) using prioritized experience replay buffers. All models and analyses were implemented in Python using the Scikit-learn and Keras packages with Tensorflow backend. We select $\tau =$

10% for the batch constrained hyperparameter because the training data contains 8.9% passive remediations and 9.4% no remediations, resulting in any τ greater than 12% never applying constrained sampling and any value below 8% always enforcing random sampling. We use $\omega = 0.5$ as our priority exponent following prior work [12] and use $\gamma = 0.95$ with a minibatch size of 128. We copy the parameters from the online network to the target network every 100 epochs. All models are trained for 20,000 epochs, repeated 5 times with different random seeds. We do not split our dataset into a training and testing set for validation as it is not necessary in batch RL [28].

We compare two alternative neural network architectures in the BCQ models: FC networks and LSTM networks. The FC BCQ models utilize four fully connected layers with 128 units per layer using ReLU activation functions. The LSTM BCQ models utilize 3 LSTM layers and a fully connected layer at the end, each with 128 units using ReLU activation functions. The output layer always uses a linear activation function to output Q values. For both architectures, a learning rate of 0.001 was utilized with the Adam optimizer [16] and L2 regularization.

For both BCQ models we explore three types of input: input with only the current state (FC-1 or LSTM-1), input with the current state and the previous state (FC-2 or LSTM-2), and input with the current state and previous two states (FC-3 or LSTM-3). In the FC BCQ models, the input states are concatenated and encoded as a single observation. In LSTM BCQ, the input states are provided sequentially.

We include three baseline models for the purpose of comparison in our analyses. All baseline models use a similar architecture as the FC BCQ models, but instead of following a greedy approach to action selection based upon the Bellman equation during model training (i.e., select the action with the maximum Q value in the next state), each baseline model takes a predetermined action while learning the Q functions as follows:

$$y = r + \gamma Q(s', a^{base}; \bar{\theta}) \quad (7)$$

Here a^{base} is a constructive remediation for the constructive baseline model, and a^{base} is no remediation for the no remediation baseline model. For the random baseline, a^{base}

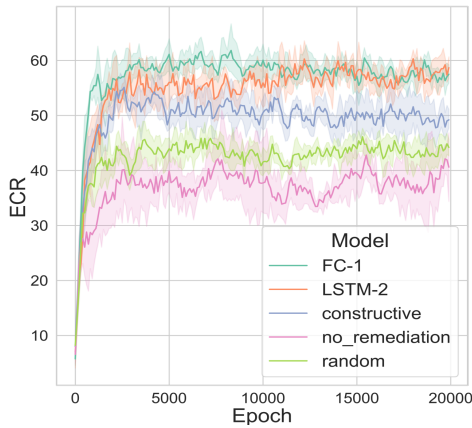


Fig. 1. Comparison between BCQ and ICAP-inspired baseline models.

involves selecting an ICAP-inspired remediation according to a random policy. We select these baselines to serve as heuristic-based models that strictly follow ICAP without RL-based tailoring. τ is set to 0 for the baseline models.

We investigated the learning curves of different models based upon their ECR values. We found LSTM-2 BCQ performs slightly better than the LSTM-1 and LSTM-3 models, whereas little difference was observed between the three FC BCQ models. Based upon these findings, we focus our remaining analyses on FC-1 and LSTM-2.

As shown in Fig. 1, FC-1 and LSTM-2 both perform better than the three baseline models that strictly follow ICAP-inspired heuristics. Among the baselines, we observe that a random policy yields the lowest ECR and a constructive policy yields the highest.

To investigate the models' behavior after they converge to a stable policy, we examined their performance during the last 2000 training epochs (10% of the total training). Results are shown in Fig. 2. We observed that FC-1 and LSTM-2 scored higher than all other baselines in terms of both the ECR and DR evaluation metrics. Running a pairwise Tukey HSD test reveals significant differences ($p\text{-value} < 0.001$) between all models except between FC-1 and LSTM-2 for both the ECR and DR metrics. These results suggest there are no observed significant differences between the FC-1 BCQ and LSTM-2 BCQ models, but both significantly outperform all of the baselines.

Next, we examined the number of times each type of ICAP-inspired remediation was selected by the different models. Once again, we use the last 2000 epochs for this analysis. For FC-1, we find that the most frequently selected action is Constructive ($M=2195$, $SD=573$), with Active being the second most frequent ($M=1183$, $SD=450$) and Passive being the third ($M=962$, $SD=674$). (All pairwise comparison p -values are less than 0.05). For LSTM-2, there are no significant differences observed between the number of Active ($M=1041$, $SD=555$), Passive ($M=1041$, $SD=618$) and No Remediation ($M=1003$, $SD=566$) actions. However, both models recommend Constructive remediation significantly more often than other types of remediation ($p=0.001$).

To interpret these results, it is useful to revisit our original research question: should ICAP be operationalized by *adaptively* scaffolding cognitive engagement, and, if so, how should we devise models for adaptively scaffolding cognitive engagement? In this analysis, higher ECR and DR values suggest a remediation policy has the potential to yield increased student normalized learning gains. Our findings indicate that the constructive heuristic policy performs the same or better than the no remediation and random policies, respectively (Fig. 1 and Fig. 2). This is consistent with the ICAP model; feedback and remediation that elicits higher cognitive engagement yields higher learning gains than policies that elicit lower or randomized cognitive engagement. Upon




Fig. 2. Performance comparison between BCQ and baseline models during the last 2,000 training epochs.

comparison with adaptive ICAP-inspired remediation policies (i.e., FC-1 BCQ and LSTM-2 BCQ), we find that the adaptive policies perform better than the non-adaptive constructive policy on the ECR and DR metrics (Fig. 1 and Fig. 2). This suggests that a deep RL-based adaptive approach to operationalizing ICAP to design scaffolding policies has promise for promoting increased learning compared to non-adaptive scaffolding policies. Notably, we see that the adaptive policies select constructive remediation more frequently than other types, which is further consistent with ICAP (Fig. 2).

There are limitations of the work. Most notably, the ECR and DR metrics are calculated from our existing dataset, which was collected according to a random policy. It is possible that the distribution of our data may differ from the true distribution when deep RL-based remediation policies are implemented in a run-time setting. This highlights the importance of future work implementing adaptive ICAP-inspired remediation policies at run-time and evaluating their impact on student learning outcomes. Despite these limitations, the results point toward the potential for improving student learning outcomes by combining ICAP with RL-based tailoring using batch constrained DQNs.

6 Conclusions and Future Work

Scaffolding cognitive engagement is a key challenge in adaptive learning environments. The ICAP framework predicts that increased cognitive engagement will yield improved learning, but it is unclear how to best translate the guidance provided by this theory into the design of scaffolding in adaptive learning environments. We utilized batch constrained deep Q-networks to induce policies for presenting learners with ICAP-inspired scaffolding in an adaptive learning environment. Empirical analysis of converged RL policies indicates that batch constrained deep Q-networks yield adaptive scaffolding policies that outperform heuristic-based policies which exclusively select constructive scaffolding, no scaffolding, or scaffolding at random. Policies induced with batch constrained deep Q-networks also select constructive scaffolding more frequently than active, passive, or no scaffolding. These results (1) support the ICAP framework, and (2) suggest that adaptively scaffolding cognitive engagement using deep RL-induced policies shows promise for optimizing student learning outcomes.

There are several promising directions for future work. First, future research should investigate the impact of additional reward and state features in deep RL-based policies to identify their impact on student learning and engagement. Second, it will be instructive to examine how multimodal data such as video-based analysis of student engagement can be used to augment RL-induced policies for scaffolding student engagement. Finally, it will be important to evaluate the RL policies by implementing them in a run-time learning environment and investigating their impact on student learning outcomes.

Acknowledgements. The research described herein has been sponsored by the U.S. Army Research Laboratory under cooperative agreement W911NF-15-2-0030. The statements and opinions expressed in this article do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

1. Ai, F., Chen, Y., Guo, Y., Zhao, Y., Wang, Z., Fu, G.: Concept-aware deep knowledge tracing and exercise recommendation in an online learning system. In: Proceedings of the 12th International Conference on Educational Data Mining, pp. 240–245. (2019).
2. Ausin, M.S., Maniktala, M., Barnes, T., Chi, M.: Exploring the impact of simple explanations and agency on batch deep reinforcement learning induced pedagogical policies. In: Proceedings of the 21st International Conference on Artificial Intelligence in Education, pp. 472–485. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-030-52237-7_38
3. Ausin, M.S., Azizsoltani, H., Barnes, T., Chi, M.: Leveraging deep reinforcement learning for pedagogical policy induction in an intelligent tutoring system. In: Proceedings of the 12th International Conference on Educational Data Mining, pp. 168–177 (2019).
4. Azizsoltani, H., Jin, Y.: Unobserved is not equal to non-existent: Using Gaussian processes to infer immediate rewards across contexts. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence, pp. 1974–1980. (2019). <https://doi.org/10.24963/ijcai.2019/273>
5. Chi, M., VanLehn, K., Litman, D.: Do micro-level tutorial decisions matter: Applying reinforcement learning to induce pedagogical tutorial tactics. In: Proceedings of 10th International Conference on Intelligent Tutoring Systems, pp. 224–234. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13388-6_27
6. Chi, M.T.H., Adams, J., Bogusch, E.B., Bruchok, C., Kang, S., Lancaster, M., Levy, R., Li, N., McEldoon, K.L., Stump, G.S., Wylie, R.: Translating the ICAP theory of cognitive engagement into practice. *Cognitive Science* 42(6), 1777–1832 (2018). <https://doi.org/10.1111/cogs.12626>.
7. Chi, M.T.H., Wylie, R.: The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist* 49(4), 219–243 (2014). <https://doi.org/10.1080/00461520.2014.965823>.
8. Doroudi, S., Aleven, V., Brunskill, E.: Where’s the reward? A review of reinforcement learning for instructional sequencing. *International Journal of Artificial Intelligence in Education* 29(4), 568–620 (2019). <https://doi.org/10.1007/s40593-019-00187-x>
9. Fujimoto, S., Meger, D., Precup, D.: Off-policy deep reinforcement learning without exploration. In: Proceedings of the 36th International Conference on Machine Learning, pp. 2052–2062 (2019).
10. Georgila, K., Core, M.G., Nye, B.D., Karumbaiah, S., Auerbach, D., Ram, M.: Using reinforcement learning to optimize the policies of an intelligent tutoring system for interpersonal skills training. In: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems, pp. 737–745, IFAAMAS, Richland, SC (2019). <https://dl.acm.org/doi/abs/10.5555/3306127.3331763>
11. Van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: Proceedings of the 30th AAAI Conference on Artificial Intelligence, pp. 2094–2100. (2016).
12. Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., Silver, D.: Rainbow: Combining improvements in deep reinforcement learning. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence, pp. 3215–3222. (2018).
13. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation*. 9(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>

14. Jiang, N., Li, L.: Doubly robust off-policy value evaluation for reinforcement learning. In: Proceedings of the 33rd International Conference on Machine Learning, pp. 652–661. (2016).
15. Ju, S., Zhou, G., Barnes, T., Chi, M.: Pick the moment: Identifying critical pedagogical decisions using long-short term rewards. In: Proceedings of the 13th International Conference on Educational Data Mining, pp. 126–136. (2020).
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980. (2014).
17. Kodinariya, T.M., Makwana, P.R.: Review on determining number of cluster in K-Means clustering. *International Journal of Advance Research in Computer Science and Management Studies*, 1(6), 90–95 (2013).
18. Lim, J., Ko, H., Yang, J.W., Kim, S., Lee, S., Chun, M.S., Ihm, J., Park, J.: Active learning through discussion: ICAP framework for education in health professions. *BMC medical education*. 19(1), Article 47 (2019). <https://doi.org/10.1186/s12909-019-1901-7>
19. Marx, J.D., Cummings, K.: Normalized change. *American Journal of Physics*. 75(1), 87–91 (2007). <https://doi.org/10.1119/1.2372468>
20. Mitrovic, A., Gordon, M., Piotrkowicz, A., Dimitrova, V.: Investigating the effect of adding nudges to increase engagement in active video watching. In: Proceedings of the 20th International Conference on Artificial Intelligence in Education, pp. 320–332. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-030-23204-7_27
21. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovsky, G., Petersen, S.: Human-level control through deep reinforcement learning. *Nature*. 518(7540), 529–533 (2015). <https://doi.org/10.1038/nature14236>
22. van de Pol, J., Volman, M., Oort, F., Beishuizen, J.: The effects of scaffolding in the classroom: Support contingency and student independent working time in relation to student achievement, task effort and appreciation of support. *Instructional Science*. 43(5), 615–641 (2015). <https://doi.org/10.1007/s11251-015-9351-z>
23. Sawyer, R., Rowe, J., Lester, J.: Balancing learning and engagement in game-based learning environments with multi-objective reinforcement learning. In: Proceedings of the 18th International Conference on Artificial Intelligence in Education, pp. 323–334. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-61425-0_27
24. Schaul, T., Quan, J., Antonoglou, I., Silver, D.: Prioritized experience replay. arXiv preprint arXiv:1511.05952. (2015).
25. Sottolare, R.A., Brawner, K.W., Goldberg, B.S., Holden, H.K.: The generalized intelligent framework for tutoring (GIFT). US Army Research Laboratory–Human Research & Engineering Directorate (ARL-HRED). Orlando, FL (2012).
26. Spain, R., Rowe, J., Goldberg, B., Pokorny, R., Lester, J.: Enhancing learning outcomes through adaptive remediation with GIFT. In: Proceedings of the Interservice/Industry Training, Simulation and Education Conference. Paper No. 19275 (2019).
27. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. 2nd edn. MIT Press, Cambridge, MA (2018).
28. Thomas, P., Brunskill, E.: Data-efficient off-policy policy evaluation for reinforcement learning. In: Proceedings of the 33rd International Conference on Machine Learning, pp. 2139–2148 (2016).
29. Wang, F.: Reinforcement learning in a POMDP based intelligent tutoring system for optimizing teaching strategies. *International Journal of Information and Education Technology*. 8(8), 553–558 (2018). <https://doi.org/10.18178/ijiet.2018.8.8.1098>

30. Wang, P., Rowe, J., Min, W., Mott, B., Lester, J.: High-fidelity simulated players for interactive narrative planning. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence, pp. 3884–3890 (2018). <https://doi.org/10.24963/ijcai.2018/540>
31. Wang, P., Rowe, J.P., Min, W., Mott, B.W., Lester, J.C.: Interactive narrative personalization with deep reinforcement learning. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence, pp. 3852–3858 (2017). <https://doi.org/10.24963/ijcai.2017/538>
32. Wiggins, B.L., Eddy, S.L., Grunspan, D.Z., Crowe, A.J.: The ICAP active learning framework predicts the learning gains observed in intensely active classroom experiences. *AERA Open*. 3(2), 1–14 (2017). <https://doi.org/10.1177/2332858417708567>
33. Zhou, G., Yang, X., Azizsoltani, H., Barnes, T., Chi, M.: Improving student-system interaction through data-driven explanations of hierarchical reinforcement learning induced pedagogical policies. In: Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization, pp. 284–292. ACM, New York (2020). <https://doi.org/10.1145/3340631.3394848>