

Closed-loop Teaching via Demonstrations to Improve Policy Transparency

Michael S. Lee, Reid Simmons, Henny Admoni
Robotics Institute, Carnegie Mellon University
{ml5, rsimmons, hadmoni}@andrew.cmu.edu

Abstract

Demonstrations are a powerful way of increasing the transparency of AI policies. Though informative demonstrations may be selected *a priori* through the machine teaching paradigm, student learning may deviate from the preselected curriculum *in situ*. This paper thus explores augmenting a curriculum with a closed-loop teaching framework inspired by principles from the education literature, such as the zone of proximal development and the testing effect. We utilize tests accordingly to close to the loop and maintain a novel particle filter model of human beliefs throughout the learning process, allowing us to provide demonstrations that are targeted to the human’s current understanding in real time. A user study finds that our proposed closed-loop teaching framework reduces the regret in human test responses by 43% over a baseline.

1 Introduction

Much progress has been made in obtaining complex and capable AI policies through reinforcement learning (e.g. [Brohan *et al.*, 2023]). Ensuring the transparency (i.e. understandability and predictability [Endsley, 2017]) of these policies in all scenarios is key to calibrating the expectations of developers and end-users toward proper usage; however, this remains a challenge [Wells and Bednarz, 2021].

One effective way to increase policy transparency is through demonstrations of the policy, which can be selected through a *machine teaching* [Zhu, 2015] paradigm that selects the minimal set of examples (e.g. demonstrations) that will help a student comprehend a concept (e.g. a policy) given their learning model. Though machine teaching can assist in selecting a principled curriculum of demonstrations *a priori*, student learning may deviate from the modeled learning trajectory *in situ*. In prior work by Lee et al. [2022], machine teaching-selected demonstrations improved human performance on post hoc tests examining understanding of later-demonstrated concepts but decreased performance on post hoc tests examining understanding of early-demonstrated concepts, suggesting perhaps that the curriculum moved too quickly past the early concepts without *in situ* testing to provide additional instruction as necessary.

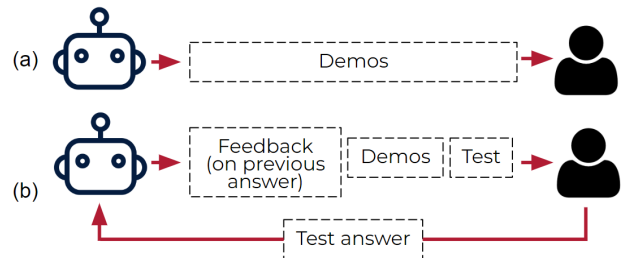


Figure 1: (a) Previous works aim to improve policy transparency via a set of demonstrations selected *a priori*, but student learning may deviate from the expected trajectory. (b) We propose a closed-loop teaching framework using tests and feedback to detect and correct for such deviations *in situ*.

Thus, our key idea is to complement a curriculum of machine teaching-selected demonstrations with a closed-loop teaching framework inspired by the education literature to provide tailored instruction in real time (Fig 1). A guiding educational concept is teaching in the *zone of proximal development* (ZPD) or “Goldilocks zone” [Hattie and Clarke, 2018; Vygotsky, 1980], which suggests that the examples provided to the learner should not be too easy nor too difficult, given their current understanding. However, the ZPD often changes at different rates for different students based on their personal learning rate, which must be assessed periodically through testing. We inform the testing cadence with the educational concept of the *testing effect* [Roediger III and Karpicke, 2006], which predicts an increase in learning outcomes when a portion of the teaching budget is devoted to testing the student (leveraging testing not only as a tool for assessment but also for teaching). And by incorporating tests and feedback in a closed teaching loop, we maintain an up-to-date model of human beliefs and promote demonstrations that are provided at the right level of difficulty *in situ*.

To illustrate the utility of our closed-loop teaching framework, consider a robot that increases the transparency of its reward function and policy to a human using demonstrations, tests, and feedback (Fig. 2). The robot’s objective is to deliver a package to the destination, whose reward function balances traveling through difficult terrain, like mud, and reducing the number of actions it takes. To convey its reward function, the robot first provides a human with the demonstration in Fig. 2a. Because the robot takes a two-action detour to avoid the mud instead of going through it, the human may infer that the

robot associates mud with a negative reward.

The robot considers what to demonstrate next to convey more information regarding its reward function. Importantly, it knows that the human likely considers mud as costly from the first demonstration, but does not know *how* costly. For instance, the human may counterfactually believe that the robot would take a four-action detour when faced with two mud patches (Fig. 2b). However, the robot knows that its ratio of mud to action reward is -3 to -1 and that consequently, it would simply go through the mud in Fig. 2b to maximize its reward. Seeing how its direct path meaningfully differs from the human’s likely detouring counterfactual (i.e. an alternative, potentially suboptimal behavior), the robot considers this to be an informative next demonstration to provide that targets the human’s ZPD – providing a meaningful yet incremental update to the human belief through an additional unit of information that upper-bounds the cost of mud.

The robot then follows the two demonstrations with a diagnostic test that simultaneously challenges the human to apply their learned knowledge and reveals whether the robot’s current model of the human’s beliefs needs to be corrected (Fig. 2c). If the human answers incorrectly, the robot may provide feedback, a remedial demonstration, then a sequence of remedial tests and feedback until the human demonstrates concept mastery, inspired by the testing effect (Fig. 2e-h). Importantly, the robot continues to update its model of the human’s beliefs according to the test answer and throughout the remedial interactions to consider the right counterfactuals when estimating the informativeness of future demonstrations.

The above interaction demonstrates the importance of maintaining a calibrated model of the human’s beliefs through closed-loop testing, which can help select demonstrations that are within the human’s ZPD. Our contributions are thus as follows: First, a closed-loop teaching framework that provides demonstrations, tests, and feedback based on insights from the education literature. Second, a particle filter model of human beliefs that supports iterative updates and a calibrated prediction of the counterfactuals likely considered by the human for each demonstration that could be provided. Third, a user study that finds that our framework reduces the regret of human test responses by 43% over a baseline.

2 Related Work

Explainable RL: The field of explainable reinforcement learning (RL) focuses on assisting humans in understanding the decision making of RL agents. Recent surveys [Milani *et al.*, 2023; Puiutta and Veith, 2020; Wells and Bednarz, 2021] highlight a variety of approaches, such as approximating a black box RL policy via an interpretable model (e.g. a decision tree [Silva *et al.*, 2020]), using saliency maps to highlight features of a state used for decision making [Greydanus *et al.*, 2018], and identification of critical training points (e.g. for estimating Q-values [Gottesman *et al.*, 2020]). In this work, we focus on a complementary direction to those outlined above – conveying an understanding of an agent’s overall behavior through representative examples.

Policy Summarization: Policy summarization aims to provide a global understanding of a policy to a user through

example state-action pairs [Amir *et al.*, 2019], which can aid in transparency. One approach relies on heuristics such as entropy or differences in Q-values to select states and actions to show [Huang *et al.*, 2018; Amir and Amir, 2018]. We instead build on the second approach based on machine teaching [Zhu, 2015], which we highlight below.

Lee *et al.* model human learning from agent demonstrations as resembling inverse reinforcement learning (IRL), and leverage human teaching techniques such as scaffolding [2021] and principles from cognitive science such as counterfactual reasoning [2022] to provide demonstrations that incrementally provide information on the agent’s underlying reward function. However, these methods model the human learner as using exact IRL [Ng and Russell, 2000], which is unable to gracefully handle conflicting information (e.g. knowledge assumed to be learned but failed to be demonstrated during testing). Furthermore, they utilize tests for assessment only after having provided demonstrations. We build on this line of work by proposing a Bayesian model of human beliefs in the form of a particle filter and also utilizing intermittent testing to simultaneously maintain an up-to-date model of human beliefs and provide targeted instruction.

Huang *et al.* [2019] also use Bayesian IRL [Ramachandran and Amir, 2007] to model human learning from agent demonstrations, but only update the relative probabilities of a static set of reward beliefs with each additional demonstration. We instead allow for resampling [Li *et al.*, 2013] of the beliefs within our particle filter to more efficiently approximate the posterior distribution of human beliefs. Finally, Qian and Unhelkar [2022] also explore interactive policy summarization where they allow humans to request specific demonstrations from an agent. They find that a hybrid strategy of AI-selected and human-selected demonstrations yields the best objective and subjective results; our proposed approach could supply the former demonstrations in their framework.

3 Technical Background

This section provides the background for selecting informative demonstrations for a (human) learner using IRL-like reasoning to infer a reward function underlying demonstrations.

Markov decision process: The agent models its world as an instance (indexed by i) of a Markov decision process, MDP_i , comprised of sets of states \mathcal{S}_i and actions \mathcal{A} , a transition function T_i , reward function R , discount factor γ , and initial state distribution S_i^0 . We refer to a group of related MDP instances as a *domain* (described below) and $\mathcal{S} : \bigcup_i \mathcal{S}_i$ is the union over all of their states. An optimal trajectory ξ^* is a sequence of (s_i, a, s'_i) tuples that follow the agent’s optimal policy π_i^* . In line with prior work [Abbeel and Ng, 2004], reward R is represented as a weighted linear combination of reward features ϕ : $R = \mathbf{w}^{*\top} \phi(s, a, s')$. Finally, we assume the human is aware of the full MDP apart from weights \mathbf{w}^* .

A domain is a group of MDPs that share R , \mathcal{A} , and γ but differ in T_i , \mathcal{S}_i , and S_i^0 . For example, all MDPs in the delivery domain share the same R even though they may contain different mud patches (Figs. 2a and 2b). Thus through IRL, all demonstrations within a domain will support inference over a common \mathbf{w}^* . We simplify the notation such that

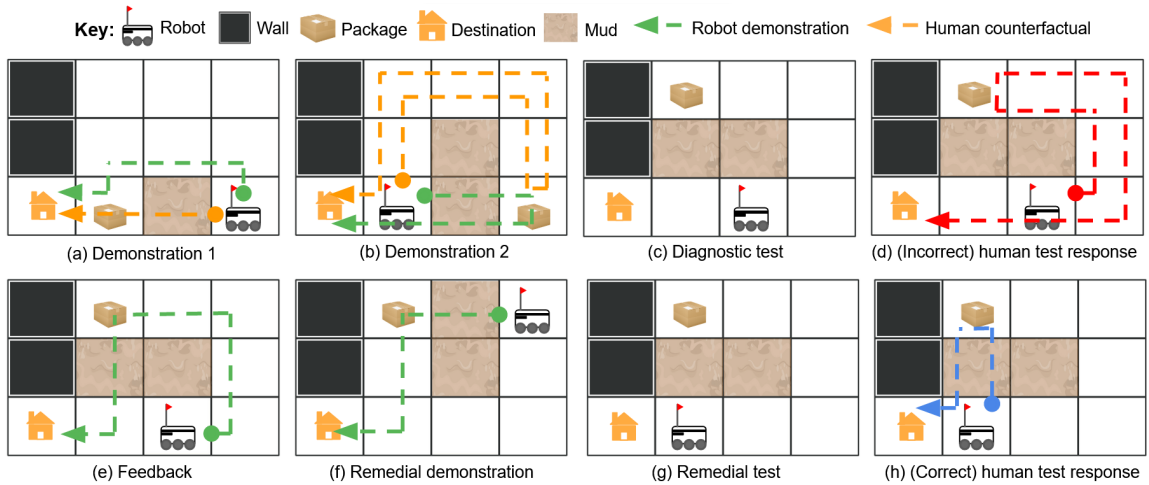


Figure 2: Sample teaching sequence for a batch of KCs on mud cost. (a) First demonstration (green) contrasts with a counterfactual alternative likely considered by a human (orange), which conveys that mud is costly. (b) Second demonstration lowerbounds mud cost. (c) Human is asked to predict the robot’s behavior in a test. (d) Incorrect response suggests that the demonstration was not understood. (e) Human is given the correct response as feedback. (f) Remedial demonstration is provided to target the misunderstanding. (g) Human is given a remedial test. (h) Correct answer suggests understanding.

π^* refers to any optimal policy within a domain, and ξ^* refers to a demonstration (dropping the corresponding MDP).

Machine teaching for policies: Our objective in selecting an informative curriculum of demonstrations for conveying π^* is captured by the machine teaching framework for policies [Lage *et al.*, 2019]. We aim to select a set of demonstrations that helps a human, who is assumed to use IRL-like reasoning [Jara-Ettinger, 2019], approximate w^* and then perhaps use planning [Shteingart and Loewenstein, 2014] to recover π^* . Thus, the objective reduces to selecting demonstrations that are informative at conveying w^* , which can be measured using behavior equivalence classes.

Behavior equivalence class: The *behavior equivalence class* (BEC) of a demonstration is the set of reward functions under which the demonstration is still optimal.

For a reward function that is a weighted linear combination of features, the BEC of a demonstration ξ^* of π^* is defined as the half-space [Lee *et al.*, 2022] formed by the exact IRL equation [Ng and Russell, 2000]

$$\text{BEC}(\xi^*|\pi^*, \pi_w) := \mathbf{w}^{*\top} (\mu_{\pi^*}^s - \mu_{\pi_w}^s) \geq 0, s = \xi^*(0), \quad (1)$$

where $\mu_{\pi}^s = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t \phi(s_t) | \pi, s_0 = s]$ is the vector of reward feature counts accrued from starting in s and following π after (π_w is the optimal policy under reward weight w) and $\xi^*(0)$ is the first state of ξ^* . Any demonstration can be converted into a constraint on w^* using Eq. 1 and a candidate belief w . Importantly, each constraint can be considered a *knowledge component* (KC) [Koedinger *et al.*, 2012] which captures a characteristic of the reward function (e.g. a trade-off between the underlying reward feature weights).

Consider again the delivery domain, which has binary reward features $\phi = [\textit{traversed mud}, \textit{battery recharged}, \textit{action taken}]$, $w^* \propto [-3, 3.5, -1]^1$. We assume that the human begins with a prior that the weight of the “action taken” feature

¹In practice, we require $\|w^*\|_2 = 1$ to bypass both the scale invariance of IRL and the degenerate all-zero reward function.

is negative (e.g. a bias toward the shortest path, Fig. 3a). The demonstration in Fig. 3b yields the constraint (or KC) in Fig. 3c, which indicates that $w_0^* \leq 2w_2^*$ (i.e. mud is at least twice as costly as an action), since two actions were taken to detour around the mud rather than counterfactually going through it (the optimal trajectory for a candidate belief that considers mud to be slightly negative, neutral, or slightly positive).

4 Methods

The example of the delivery robot in Section 1 highlights the importance of maintaining an up-to-date model of human beliefs and likely counterfactuals when selecting a demonstration. In this section, we propose a particle filter-based model of human beliefs amenable to iterative Bayesian updates and sampling for counterfactual reasoning, where each particle represents a potential human belief regarding the agent’s reward function. We then leverage this model in a closed-loop teaching framework that leverages insights from the education literature to select demonstrations that target gaps identified through testing.

4.1 Particle Filter Human Model

The particle filter routines detailed in the following paragraphs come together in Alg. 1.

Updating Particle Positions and Weights: Assume a set of particles, defined by their positions and associated weights $\{\mathbf{x}_t, \tilde{w}_t\}$. Without loss of generality, assume that a demonstration or test response is provided at each time step t . Each demonstration generates multiple constraints by comparing the demonstration against possible counterfactual trajectories and each incorrectly answered test will generate a single constraint by comparing the true test answer against the incorrect answer, both through Eq. 1. Each constraint generated via a demonstration or a test response is a half-space constraint, with one side being *consistent* with the demonstration or test response and the other side being *inconsistent*.




Figure 3: Example sequence on how a demonstration updates a particle filter model of human beliefs. The robot reward function is shown as a red dot, and the constraint consistent with the demonstration is shown in all plots for reference. (a) Particles before demonstration (prior). (b) Demonstration shown to the human, alongside a counterfactual that considers mud to be slightly negative or positive, or neutral. (c) The constraint (Eq. 1) consistent with the demonstration that conveys that mud must be at least twice as costly as an action, visualized with the uniform distribution portion of the custom distribution (Fig. 4) used to update particle weights. (d) Particles after demonstration (posterior).

Algorithm 1 Particle Filter for Modeling Human Beliefs

```

1: Initialize particles  $x_0^{(i)} \sim p(x_0)$  for  $i = 1, \dots, N$ 
2: for  $t = 1, \dots, T$  do
3:   // Update filter given new demonstration or test at  $t$ 
4:   for  $i = 1, \dots, N$  do
5:     Compute weight  $\tilde{w}_t^{(i)} = \check{w}_{t-1}^{(i)} \cdot p(x_t^{(i)} | y_t)$ 
6:   end for
7:   if  $\sum_{j=1}^N \tilde{w}_t^{(j)} < \check{w}_{threshold}$  then
8:     Perform a particle filter reset
9:   end if
10:  Normalize weights  $\tilde{w}_t^{(i)} = \frac{\tilde{w}_t^{(i)}}{\sum_{j=1}^N \tilde{w}_t^{(j)}}$ 
11:  Compute effective sample size  $n_{eff} = \frac{1}{\sum_{i=1}^N (\tilde{w}_t^{(i)})^2}$ 
12:  if  $n_{eff} < N_{threshold}$  then
13:    Resample  $x_t^{(i)}$  with probabilities  $\tilde{w}_t^{(i)}$  using KLD resampling
14:  end if
15: end for

```

Each constraint y_t can then be translated into a probability distribution $p(x_t | y_t)$ that can be used to update the weights of each particle (Fig. 3). We propose a custom probability distribution $p(x_t | y_t)$ that translates each constraint into a combination of a uniform distribution that aligns with the consistent half-space of the constraint and a von Mises-Fisher distribution (a generalization of the Gaussian distribution on a sphere [Dhillon and Sra, 2003]) whose mean direction aligns with the inconsistent half-space (Fig. 4). The uniform distribution asserts that any particle lying on the consistent half-space is equally valid for that demonstration, whereas the Von-Mises Fisher distribution asserts that a particle is exponentially less likely to have generated that demonstration as you move away from the consistent side of the constraint. The resulting probability density function (pdf) of the custom distribution is given in Eq. 2, with the normalizing constant c_1 that ensures that the pdf sums to 1 (Eq. 3), and the scaling constant c_2 that matches the probability of the Von-Mises Fisher distribution to that of the uniform distribution at meeting point of the two distributions (Eq. 4). Though the custom distribution naturally generalizes to higher dimensions, the particles in our domains with three reward features are constrained to the 2-sphere and the pdf is specified accordingly.




Figure 4: Cross-section of the spherical probability density function used to update particle weights given a constraint generated from a demonstration (Eq. 1).

$$f(x, \mu, \kappa) = \begin{cases} \frac{1}{2\pi c_1}, & \mu^\top x \geq 0 \\ \frac{c_2 \kappa e^{\kappa \mu^\top x}}{2c_1 \pi (e^\kappa - e^{-\kappa})}, & \mu^\top x < 0 \end{cases} \quad (2)$$

$$c_1 = \frac{1}{c_2 \int_0^\pi \int_{\frac{\pi}{2}}^{\frac{3\pi}{2}} \frac{\kappa e^{\kappa \cos(\theta) \cdot \sin(\phi) \sin(\phi)}}{2\pi (e^\kappa - e^{-\kappa})} d\theta d\phi + 0.5} \quad (3)$$

$$c_2 = \frac{1}{4\pi f(y, \mu, \kappa)}, \forall y \text{ s.t. } \mu^\top y = 0 \quad (4)$$

Sampling Human Beliefs: Given a running particle filter model, we may sample human beliefs in order to do counterfactual reasoning over how the human may interpret each demonstration that could be shown. We first run systematic resampling on a copy of the particles to downselect to a candidate set, favoring those that are higher weighted. We then rely on the 2-approximation algorithm [Hochbaum and Shmoys, 1985] to greedily select k distributed samples such that the maximum distance from any particle in the candidate set to one of the k samples is minimized. The algorithm iteratively picks the particle with the largest distance to the already selected samples as the next sample; this heuristic ensures that the maximum distance from any particle to any of the selected samples is never worse than twice the optimal. As nearby particles are likely to generate similar counterfactuals, we sample beliefs that are approximately spread out.

Due to space constraints, please find practical tips on how to resample the particle filter to combat sample degeneracy and impoverishment (line 13 of Alg. 1), as well as how to reset the particle filter if it receives heavily conflicting information (line 8 of Alg. 1) in the supplementary material.

4.2 Closed-loop Teaching

With a particle filter model of human beliefs amenable to iterative updates via demonstrations and tests, we now formulate a closed-loop teaching framework for conveying an agent’s reward function to a human. As we walk through the framework conceptualized in Fig. 5, we highlight the principles from the education literature that guide the design. A sample rollout of a teaching sequence is shown in Fig. 2, which serves as a visual correspondence to the algorithmic characterization of the framework in the supplementary material.

We first leverage feature and counterfactual scaffolding from Lee et al. [2022] to select KCs that incrementally increase in information across an increasing subset of features (e.g. mud vs action cost, recharging vs action cost, then trade-offs between all three). This set of KCs guides the machine teaching selection of the *curriculum* of demonstrations that can be used to teach the agent reward function to a human.

We begin the loop by taking a single batch of related KCs that define a *lesson* (e.g. bounds on mud cost) and providing it to the *demonstrator* (Fig. 5) to select demonstrations from the curriculum that convey these KCs. Specifically, we utilize counterfactual reasoning [Lee et al., 2022] to select demonstrations that are informative with respect to the counterfactuals likely considered by the human. We simultaneously leverage the educational principles of the *ZPD* [Vygotsky, 1980] to provide a sequence of demonstrations that provide information incrementally, i.e. demonstrations that convey one new constraint at a time (such as first providing a lower-bound on the mud cost, then later an upper-bound).

After the demonstrations have been provided, the *tester* selects *diagnostic tests* that will verify whether the human has learned the KCs in the lesson. These diagnostic tests optimize for visual dissimilarity from the teaching demonstrations and visual complexity (i.e. increasing distracting visual clutter) [Lee et al., 2022] to challenge the learner.

For each diagnostic test that is answered incorrectly, the *evaluator* will provide immediate *feedback* to the human on how their answer differed from the correct one, inspired by findings that immediate feedback on errors leads to better learning outcomes [Koedinger et al., 2013]. And for each diagnostic test that is answered incorrectly, a remedial demonstration that most closely conveys the missed KC with visual simplicity [Lee et al., 2021] will be provided to focus on the concept being taught, along with a remedial test with visual complexity to challenge the learner in demonstrating the missed KC. We note that this missed KC is determined by comparing the human’s test answer with the optimal test answer; while it may or may not be the same as one of the KCs originally contained in the lesson, it best addresses the human’s current misunderstanding. If the human also gets the remedial test wrong, the switch in Fig. 5 (labeled ‘S’) flips, and the *tester* and *evaluator* will continue to provide only visually dissimilar and complex remedial tests with cor-

responding feedback (but no additional demonstrations) until the human shows understanding of each iteration’s missed KC. This is motivated by the testing effect [Roediger III and Karpicke, 2006], which supports using tests not only for assessment but also for teaching and increasing learning outcomes. Note that for each demonstration provided or test response received throughout this learning process, we update the particle filter model of the human’s beliefs. And we utilize the particle filter model to consider the counterfactuals the human is likely to consider for each potential remedial demonstration or remedial test in order to select the one that will best convey or test the missed KC for the human. Once all of the missed KCs for this lesson have been demonstrated via correct remedial test responses, a fresh batch of KCs (i.e. a new lesson) is pulled from the KC bank and the switch flips upward to provide demonstrations again.

Alternatively, if all diagnostic tests in this lesson had been correctly answered initially, a fresh batch of KCs would have been pulled from the KC bank to begin the next lesson directly without remedial instruction.

When all lessons have been taught, the human’s subsequent knowledge can be evaluated on a held-out set of tests in which they predict the policy in previously unseen environments.

5 User Study

We ran an online user study exploring whether our proposed closed-loop teaching method improves the transparency of an agent’s policy to a human. The study involved participants learning about the agent policy in two domains through a combination of demonstrations, tests, and feedback and predicting the agent’s behavior in new test environments.

The within-subject variable was *domain*, which consisted of the following two conditions. In the *delivery* domain, the agent is penalized for moving out of mud and rewarded for recharging. In the *skateboard* domain, the agent is rewarded each time it moves with the skateboard (e.g. riding is efficient) or traverses through a designated path (see Fig. 6). Thus each domain consists of two unique reward features and one shared feature that penalizes each action. The *skateboard* domain was designed to be more challenging than the *delivery* domain (confirmed through pilot studies and study results), as the value of the skateboard depends both on the distance to the skateboard and subsequent distance to the goal.

The between-subjects variable was *feedback loop* with the following three conditions. *Open* feedback loop followed prior work by Lee et al. [2022] in utilizing counterfactual reasoning to select a set of informative demonstrations a priori that incrementally decreased in cumulative BEC area (i.e. a model of human beliefs), one KC at a time. *Partial* feedback loop additionally provided a diagnostic test after each lesson and provided feedback as necessary, while the *full* feedback loop additionally provided a remedial demonstration and remedial tests until the KC in question was correctly applied in a remedial test. For a fair comparison, each condition showed the same median number of demonstrations and tests (11 for delivery and 22 for skateboard).

The user study consisted of two trials, with each trial comprising a teaching portion and a testing portion in one domain.

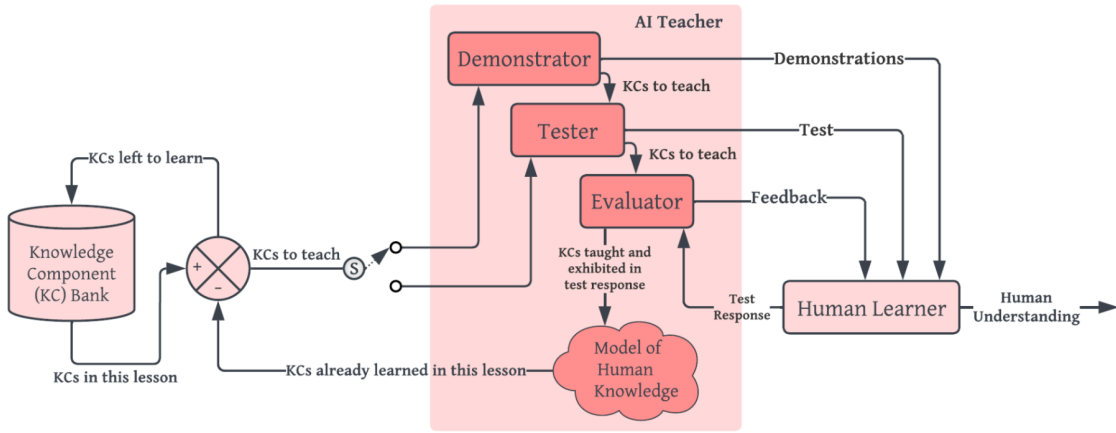


Figure 5: Proposed closed-loop teaching framework. Knowledge components (KCs) are passed to the AI teacher as a lesson. The demonstrator generates demonstrations that convey the KCs, the tester provides test(s), and the evaluator analyzes the test response(s), provides feedback on its correctness, and updates the model of human knowledge. If the human fails to learn a KC through two rounds of demonstrations and tests, the switch (labeled ‘S’) flips such that only tests and feedback are provided until an understanding of the remaining KCs is demonstrated through correct responses.




Figure 6: Two domains designed for user study, (a) delivery, (b) skateboard. The semantics of the objects were hidden using arbitrary shapes and colors.

During teaching, participants were first explicitly informed of the reward features of the domain. Then they inferred the corresponding reward weights by watching demonstrations and perhaps undergoing diagnostic tests, corrective feedback, and further remedial instruction depending on their assigned feedback loop condition. For every interaction, participants indicated whether it improved their understanding of the agent’s policy via a Likert scale. At the end of the teaching session, participants were asked to rate their level of focused attention, the usability of their assigned teaching condition, and their understanding of the agent’s policy via Likert scales. During testing, participants were tasked with predicting the agent’s optimal trajectory in six unseen test environments in random order, which were selected according to prior work [Lee *et al.*, 2022] to comprise two low, medium, and high difficulty environments each. We tested the following hypotheses (H1–H4) using the measures (M1–M4) below.

H1: (a) The test responses will be best for *full* feedback loop, then *partial*, then *open*. (b) *Delivery* will result in better test responses over *skateboard*.

H2: (a) Focused attention and perceived usability will be highest for *full* feedback loop, then *partial*, then *open*. (b) *Delivery* will result in higher focused attention and perceived usability over *skateboard*.

H3: (a) Improvement ratings will be highest for *full* feedback loop, then *partial*, then *open*. (b) *Delivery* will result in higher

improvement ratings over *skateboard*.

H4: (a) Understanding ratings will be highest for *full* feedback loop, then *partial*, then *open*. (b) *Delivery* will result in higher understanding ratings over *skateboard*.

M1. Test response: The reward of the human’s test response, measuring the human’s ability to predict the agent’s policy.

M2. Focused attention and perceived usability: We adapted the User Engagement Scale short form [O’Brien *et al.*, 2018] to ask six questions targeting focused attention and perceived usability, each answered with a 5-pt Likert scale.

M3. Improvement: “Did this interaction improve your understanding of the game strategy [i.e. agent policy]?”, answered with a 5-pt Likert scale.

M4. Understanding: “Do you feel that you now understand the game strategy?”, answered with a 5-point Likert scale.

6 Results and Discussion

We collected data from 206 participants using Prolific. Participants were roughly 70% male, 28% female, 1% non-binary, and 1% preferred not to disclose, and ages varied from 18 to 67 ($M = 32.49$, $SD = 11.15$). The recruitment process and study were approved by Carnegie Mellon University’s Institutional Review Board. In the *full* feedback loop condition, we removed data from one participant who did not miss any diagnostic tests during teaching (thus did not see any remedial instruction), and one outlier participant whose total number of interactions exceeded 3 standard deviations of the mean number of interactions in this condition (as repeated failures of similar remedial tests suggested lack of attention). This left 68 participants in each between-subjects condition.

H1: We measured the degree of suboptimality of human test responses using regret, i.e. the difference between rewards of human and optimal test responses. A two-way mixed ANOVA indicated a significant effect of feedback loop on regret ($F(2, 201) = 3.65, p = .03$). Tukey analyses revealed that *full* ($M = 0.24$) had 43% lower regret over *open* ($M = 0.42, p = .027$), with *partial* sitting in between with no significant difference to either ($M = 0.29$,

Fig. 7a). The ANOVA also indicated a significant effect of domain on regret ($F(1, 201) = 50.75, p \leq .001$), where a t-test revealed a significant difference between the regret between *delivery* ($M = 0.18$) and *skateboard* ($M = 0.45$), $t(406) = -5.792, p < .001$.

The ANOVA also indicated an interaction effect ($F(2, 201) = 3.45, p = .03$) between feedback loop and domain. In the *skateboard* domain, Tukey analyses revealed that *full* ($M = 0.33$) had significantly lower regret over *open* ($M = 0.62, p = .014$),

H1a is partially supported. Though the regret for *partial* sat in between *full* and *open* as expected (being an intermediary between those two conditions), it was not significantly different from either. However, *full* did indeed significantly outperform *open* even with the same median number of interactions – highlighting the importance of tailoring the content and interaction type to the human in situ. Finally, the interaction effect reveals that the difference between *full* and *open* on regret is largely driven by the *skateboard* domain, suggesting perhaps that the benefit of the proposed fully closed-loop teaching scheme is greater for more challenging domains. *H1b is supported.* *Delivery* resulted in a significantly lower regret over *skateboard*, as expected.

H2: A two-way mixed ANOVA did not find a significant effect of feedback loop ($F(2, 201) = 1.56, p = 0.21$), nor domain ($F(1, 201) = 0.38, p = .54$) on focused attention, nor an interaction effect between feedback loop and domain on focused attention ($F(2, 201) = 1.90, p = .15$). A two-way mixed ANOVA found a significant effect of domain on perceived usability ($F(1, 201) = 85.77, p < .001$). A t-test revealed a significant difference in the perceived usability ratings of *delivery* ($M = 3.57$) and *skateboard* ($M = 2.89$), $t(406) = 6.562, p < .001$. Finally, a two-way mixed ANOVA also found an interaction effect between feedback loop and domain on perceived usability ($F(2, 201) = 6.17, p = .003$), where Tukey revealed a significant difference between *partial* ($M = 2.64$) and *open* ($M = 3.21$) for *skateboard* ($p = .006$, Fig. 7b). A main effect of feedback loop on perceived usability was not found ($F(2, 201) = 2.06, p = .13$).

H2a is not supported. Though no main effects were found for feedback loop on focused attention or perceived usability, analysis of interaction effects in the *skateboard* domain interestingly reveals that *partial* feedback loop is rated less usable than *open* loop. A number of people in *partial* noted that they wanted more demonstrations to clear up confusion, and we hypothesize that perhaps it can be frustrating to have diagnostic tests highlight gaps in understanding without providing further instruction (as in the case of *full*) or simply providing additional instruction without highlighting gaps in understanding (as in the case of *open*). *H2b is partially supported.* The trend of the domain differences continues with *delivery* yielding significantly higher ratings of perceived usability over *skateboard*, though no difference was found between the domains for focused attention.

H3: As participants gave an *improvement* rating for each demonstration or test, a mean is more descriptive than a median for each participant and for each domain, and we use parametric analyses accordingly. A two-way mixed ANOVA

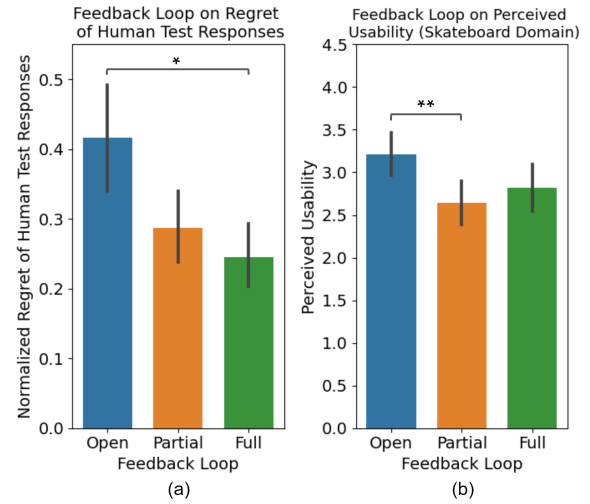


Figure 7: (a) *Full* closed-loop teaching yields lower regret for human tests responses than *open* across domains (lower is better). (b) *Partial* yields lower ratings on perceived usability (higher is better) than *open* in the skateboard domain. Error bars indicate 95% confidence intervals.

indicated a significant effect of domain on improvement ($F(1, 201) = 32.17, p < .001$). A t-test revealed that the teaching in *delivery* ($M = 3.38$) was rated to yield higher improvement than in *skateboard* ($M = 3.12$), $t(406) = 3.001, p = .003$. The ANOVA did not indicate a significant effect of feedback loop ($F(2, 201) = 1.54, p = .22$) nor a significant interaction effect ($F(2, 201) = 1.23, p = .29$) between feedback loop and domain.

H3a is not supported. Feedback loop did not impact ratings of improvement. *H3b is supported.* The ratings suggest that participants learned more overall about the *delivery* domain than the *skateboard* domain.

H4: A Kruskal-Wallis H test did not reveal a statistically significant effect of feedback loop on ratings of understanding ($p = .41$). However, a Wilcoxon signed-rank test showed a statistically significant difference in ratings of understanding between *delivery* and *skateboard* domains ($Z = -6.474, p < .001$). The mean rating on understanding for *delivery* was 3.90 and the mean rating for *skateboard* was 3.34.

H4a is not supported. Feedback loop did not impact ratings of understanding. *H4b is supported.* The ratings support a difference in the difficulty of the two domains.

7 Conclusion

Machine teaching provides a principled framework for selecting demonstrations *a priori* that increases the transparency of AI policies to humans; however, individuals may differ in their learning trajectories *in situ*. We thus augment a curriculum of preselected demonstrations with a novel closed-loop teaching framework inspired by key concepts from the education literature to provide tailored instruction. A user study finds that our teaching framework consisting of demonstrations, tests, feedback, and remedial instruction reduces the regret in human test responses by 43% over a baseline.

References

- [Abbeel and Ng, 2004] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, 2004.
- [Amir and Amir, 2018] Dan Amir and Ofra Amir. Highlights: Summarizing agent behavior to people. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2018.
- [Amir et al., 2019] Ofra Amir, Finale Doshi-Velez, and David Sarne. Summarizing agent strategies. *Autonomous Agents and Multi-Agent Systems*, Sep 2019.
- [Brohan et al., 2023] Anthony Brohan, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Alexander Herzog, Daniel Ho, Julian Ibarz, Alex Irpan, Eric Jang, Ryan Julian, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, 2023.
- [Dhillon and Sra, 2003] Inderjit S Dhillon and Suvrit Sra. Modeling data using directional distributions. Technical report, Citeseer, 2003.
- [Endsley, 2017] Mica R Endsley. From here to autonomy: lessons learned from human–automation research. *Human factors*, 2017.
- [Gottesman et al., 2020] Omer Gottesman, Joseph Futoma, Yao Liu, Sonali Parbhoo, Leo Celi, Emma Brunskill, and Finale Doshi-Velez. Interpretable off-policy evaluation in reinforcement learning by highlighting influential transitions. In *ICML*, 2020.
- [Greydanus et al., 2018] Samuel Greydanus, Anurag Koul, Jonathan Dodge, and Alan Fern. Visualizing and understanding atari agents. In *International conference on machine learning*. PMLR, 2018.
- [Hattie and Clarke, 2018] John Hattie and Shirley Clarke. *Visible learning: feedback*. Routledge, 2018.
- [Hochbaum and Shmoys, 1985] Dorit S Hochbaum and David B Shmoys. A best possible heuristic for the k-center problem. *Mathematics of operations research*, 1985.
- [Huang et al., 2018] Sandy H Huang, Kush Bhatia, Pieter Abbeel, and Anca D Dragan. Establishing appropriate trust via critical states. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [Huang et al., 2019] Sandy H Huang, David Held, Pieter Abbeel, and Anca D Dragan. Enabling robots to communicate their objectives. *Autonomous Robots*, 43(2):309–326, 2019.
- [Jara-Ettinger, 2019] Julian Jara-Ettinger. Theory of mind as inverse reinforcement learning. *Current Opinion in Behavioral Sciences*, 29:105–110, 2019.
- [Koedinger et al., 2012] Kenneth R Koedinger, Albert T Corbett, and Charles Perfetti. The knowledge-learning-instruction framework: Bridging the science–practice chasm to enhance robust student learning. *Cognitive science*, 2012.
- [Koedinger et al., 2013] Kenneth R Koedinger, Julie L Booth, and David Klahr. Instructional complexity and the science to constrain it. *Science*, 2013.
- [Lage et al., 2019] Isaac Lage, Daphna Lifschitz, Finale Doshi-Velez, and Ofra Amir. Exploring computational user models for agent policy summarization. In *International Joint Conference on Artificial Intelligence*, 2019.
- [Lee et al., 2021] Michael S Lee, Henny Admoni, and Reid Simmons. Machine teaching for human inverse reinforcement learning. *Frontiers in Robotics and AI*, 2021.
- [Lee et al., 2022] Michael S Lee, Henny Admoni, and Reid Simmons. Reasoning about counterfactuals to improve human inverse reinforcement learning. In *IROS*. IEEE, 2022.
- [Li et al., 2013] Tiancheng Li, Shudong Sun, and Tariq Pervez Sattar. Adapting sample size in particle filters through kld-resampling. *Electronics Letters*, 2013.
- [Milani et al., 2023] Stephanie Milani, Nicholay Topin, Manuela Veloso, and Fei Fang. Explainable reinforcement learning: A survey and comparative review. *ACM Comput. Surv.*, aug 2023.
- [Ng and Russell, 2000] Andrew Y Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *International Conf. on Machine Learning*, 2000.
- [O’Brien et al., 2018] Heather L O’Brien, Paul Cairns, and Mark Hall. A practical approach to measuring user engagement with the refined user engagement scale (ues) and new ues short form. *International Journal of Human-Computer Studies*, 2018.
- [Puiutta and Veith, 2020] Erika Puiutta and Eric MSP Veith. Explainable reinforcement learning: A survey. In *International cross-domain conference for machine learning and knowledge extraction*. Springer, 2020.
- [Qian and Unhelkar, 2022] Peizhu Qian and Vaibhav Unhelkar. Evaluating the role of interactivity on improving transparency in autonomous agents. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 1083–1091, 2022.
- [Ramachandran and Amir, 2007] Deepak Ramachandran and Eyal Amir. Bayesian inverse reinforcement learning. In *IJCAI*, volume 7, pages 2586–2591, 2007.
- [Roediger III and Karpicke, 2006] Henry L Roediger III and Jeffrey D Karpicke. The power of testing memory: Basic research and implications for educational practice. *Perspectives on psychological science*, 2006.
- [Shteingart and Loewenstein, 2014] Hanan Shteingart and Yonatan Loewenstein. Reinforcement learning and human behavior. *Current Opinion in Neurobiology*, 25:93–98, 2014.
- [Silva et al., 2020] Andrew Silva, Matthew Gombolay, Taylor Killian, Ivan Jimenez, and Sung-Hyun Son. Optimization methods for interpretable differentiable decision trees applied to reinforcement learning. In *Intern. conference on artificial intelligence and statistics*, 2020.

- [Vygotsky, 1980] Lev Semenovich Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard university press, 1980.
- [Wells and Bednarz, 2021] Lindsay Wells and Tomasz Bednarz. Explainable ai and reinforcement learning—a systematic review of current approaches and trends. *Frontiers in AI*, 2021.
- [Zhu, 2015] Xiaojin Zhu. Machine teaching: an inverse problem to machine learning and an approach toward optimal education. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 4083–4087, 2015.