

HOPE: Human-Centric Off-Policy Evaluation for E-Learning and Healthcare

Ge Gao

Department of Computer Science
North Carolina State University
Raleigh, NC, USA
ggao5@ncsu.edu

Markel Sanz Ausin

Department of Computer Science
North Carolina State University
Raleigh, NC, USA

Song Ju

Department of Computer Science
North Carolina State University
Raleigh, NC, USA
sju2@ncsu.edu

Min Chi

Department of Computer Science
North Carolina State University
Raleigh, NC, USA
mchi@ncsu.edu

ABSTRACT

Reinforcement learning (RL) has been extensively researched for enhancing human-environment interactions in various human-centric tasks, including e-learning and healthcare. Since deploying and evaluating policies online are high-stakes in such tasks, off-policy evaluation (OPE) is crucial for inducing effective policies. In human-centric environments, however, OPE is challenging because the underlying state is often unobservable, while only aggregate rewards can be observed (students' test scores or whether a patient is released from the hospital eventually). In this work, we propose a *human-centric OPE (HOPE)* to handle partial observability and aggregated rewards in such environments. Specifically, we reconstruct immediate rewards from the aggregated rewards considering partial observability to estimate expected total returns. We provide a theoretical bound for the proposed method, and we have conducted extensive experiments in real-world human-centric tasks, including sepsis treatments and an intelligent tutoring system. Our approach reliably predicts the returns of different policies and outperforms state-of-the-art benchmarks using both standard validation methods and human-centric significance tests.

KEYWORDS

Off-policy evaluation; reinforcement learning; POMDP; aggregated rewards; human-centric environments; e-learning; healthcare

ACM Reference Format:

Ge Gao, Song Ju, Markel Sanz Ausin, and Min Chi. 2023. HOPE: Human-Centric Off-Policy Evaluation for E-Learning and Healthcare. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 10 pages.

1 INTRODUCTION

Off-policy evaluation (OPE) sits at the epicenter of offline Reinforcement Learning (RL) research [10, 24, 37, 52], which estimates the performance of an RL-induced policy leveraging prior knowledge obtained from historical data. OPE is especially important in

human-centric environments where the execution of a bad policy can be costly and dangerous. For example, in healthcare, an over-estimated policy could be ineffective and even increase mortality by wrong treatments. There are at least **two main challenges** for OPE to work with human-centric environments. One is that the underlying state of a human-centric environment is usually **unobservable** [7, 33], resulting in a partially observable Markov decision process (POMDP). Specifically, the observations may not contain sufficient information to fully reconstruct the underlying states. For example, the population considered in the environment could change over time, and their characteristics may be varied across different cohorts. In e-learning, student backgrounds can vary from semester to semester; similarly, in healthcare, patient demographics can change in hospitals across different regions [27, 42]. **The other challenge** is that **only an aggregated (or delayed) reward can be observed** after a certain period of time, where all immediate rewards in between are often missing. The most appropriate rewards in e-learning and healthcare are student learning performance and patient outcomes, which are typically unavailable until the entire trajectory is complete. This is due to the complex natures of both learning and disease progression, which make it difficult to assess students' learning or patient health states moment by moment. More importantly, many instructional or medical interventions that boost short-term performance may not be effective over the long term. Different from delayed rewards in classic mouse-in-the-maze situations where agents receive insignificant rewards along the path and a significant reward in the final goal state (the food), in e-learning and healthcare, there are immediate rewards along the way but they are often **unobservable**.

Prior OPE work [3, 16, 24, 37] have achieved outstanding performance in simulated environments such as Mujoco [55]. However, one may not be able to directly apply such methods toward human-centric environments, since immediate rewards are missing and the environment is partially observable. A possible way is assuming the rewards are sparse that only indicate whether a task is completed partially or fully. Sparse rewards typically correspond to the attainment of some particular tasks such as a robot attaining designated waypoints which provide little feedback for immediate steps [46]. In many human-centric tasks, interaction outcomes can be gradually improved step-by-step and the immediate reward on

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

each step can be meaningful by itself. For example, students’ cognitive outcomes and learning performance are gradually improved during interaction with the intelligent tutor in college [34]. It has been found that immediate rewards are more effective than sparse rewards, toward evaluating decision outcomes [5]. Moreover, in human-centric environments, policy performance may be correlated with the horizon of the environment. For example, medical interventions that have shown good short-term performance may not be effective over the long-term [5]. Consequently, it is important to reconstruct immediate rewards for OPE to work effectively in human-centric environments.

On the other hand, prior evaluation metrics for OPE are generally error metrics (e.g., absolute error, rank correlations) proposed by Fu et al. [12], Voloshin et al. [60], while human-centric research often emphasizes the need for *statistical significance test in empirical study* [20, 64]. For example, rank correlation summarizes the performance of a set of policies’ relative rankings using averaged returns. Statistical significance tests can tell how likely the relationship we have found is due only to random chance across samples, and they are commonly employed and easier to be interpreted by domain experts [20, 25]. The framework for validating OPE in human-centric environments may need to be extended beyond error metrics.

In this work, we propose a **human-centric OPE (HOPE)** approach to tackle the two challenges above. Specifically, it first reconstructs immediate rewards from the aggregated rewards. Then, importance sampling is used to process the reconstructed rewards and estimate expected total returns. Any existing OPE method can be used jointly with the reconstructed rewards to estimate expected returns. We validate HOPE on two typical real-world human-centric tasks in healthcare and e-learning, *i.e.* sepsis treatments and an intelligent tutoring system (ITS), and extend existing validation frameworks with significance tests for human-centric environments. To summarize, our work has at least three main contributions:

- To the best of our knowledge, HOPE is the first OPE approach that tackles both partial observability and aggregated rewards. We also provide theoretical bound for HOPE.
- HOPE is extensively validated through real-world environments in real-world human-centric environments including e-learning and healthcare. The results show that our approach outperforms the state-of-the-art OPE approaches.
- We introduce significance tests on top of existing OPE validation frameworks, to facilitate comprehensive study and comparison of OPE methods in human-centric environments.

2 RELATED WORK

In e-learning and healthcare, RL has been widely investigated to learn policies from historical user interaction data [18, 39]. However, deploying and evaluating policies online are high stakes in such domains, as a poor policy can be fatal to humans in healthcare. It’s thus crucial to propose effective OPE methods for human-centric environments. OPE is used to evaluate the performance of a *target* policy given historical data drawn from (alternative) *behavior* policies. Classic methods, such as expected cumulative reward (ECR) [51], have been employed in real-world applications such as e-learning [9]. However, ECR is not statistically consistent, which

is a significant concern for high-stakes domains [33]. In practice, researchers have found that selected policies based on ECR were even not effective compared to random policies with real students in terms of human-centric significant test [9, 48].

Various contemporary OPE methods have been proposed and can be divided into three general categories [60]: 1) Inverse Propensity Scoring (IPS) [10, 44]; 2) Direct Methods [24, 30, 31, 37, 57, 61]; 3) Hybrid Methods [24, 52]. IPS has been widely investigated in statistics [43] and RL [44], with the key idea to reweigh the rewards in historical data using the importance ratio between β and π . IPS yields consistent estimates and it has several variations including IS [44], WIS [44], PDIS [44], and PHWIS [10], etc. Direct Methods directly estimate the value functions of the evaluation policy. For example, FQE [30] is functionally a policy evaluation counterpart to batch Q-learning. DualDICE [37] estimates the discounted stationary distribution ratios, which measure the likelihood that π will experience the state-action pair normalized by the probability with which the state-action pair appears in the off-policy data. Hybrid Methods combine aspects of both IPS and direct methods. For example, DR [24] is an unbiased estimator leveraging a direct method to decrease the variance of the unbiased estimates produced by IS. These OPEs have achieved desirable performance in simulated environments. Recently, a few approaches have proposed OPE targeting some real-world tasks such as robotic grasping [8, 17, 22, 40]. The proposed methods generally assume the rewards are observable from the environment or sparse. It lacks investigation into the effectiveness of applying the existing OPE methods in real-world human-centric environments.

Recently, researchers have recognized the partial observability of human-centric environment and proposed OPE for POMDPs [6, 50, 58]. They assume that the underlying states may not exist and treat them as confounding for policy evaluation. For example, Tennenholtz et al. propose a method for POMDP with unobserved confounding and compare their method to IS in carefully-designed synthetic environments [50]. The results show that their proposed method can outperform IS under certain levels of confounding. In practice, such as in our intelligent tutoring experiment with real students, confounding is agnostic and unable to control, accompanied by missing immediate rewards. Our goal is to utilize limited observable information from real-world e-learning and healthcare environments, and effectively estimate policy performance with both partial observability and missing immediate rewards.

3 HUMAN-CENTRIC OPE (HOPE)

Problem Definition: We consider the human-centric environment as a **partially observable Markov decision process (POMDP)**, which is a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{P}, \Omega, R, \gamma)$. The state space \mathcal{S} is assumed *unknown*. Moreover, \mathcal{A} represents the action space, \mathcal{O} is the observation space, and \mathcal{P} defines transition dynamics from the current state to the next states given an action. The observation model $p(o|s) \sim \Omega$ is also assumed *unknown*. R is the immediate reward function, and $r_t = R(s_t, a_t)$ are the *immediate* rewards. $\gamma \in (0, 1]$ is discount factor. Each episode has a finite horizon T . In general, at time-step t , the agent is given an observation $o_t \in \mathcal{O}$ by the environment, then chooses an action $a_t \in \mathcal{A}$ following some policy $\pi : \mathcal{O} \rightarrow \mathcal{A}$. Then, the environment provides the immediate

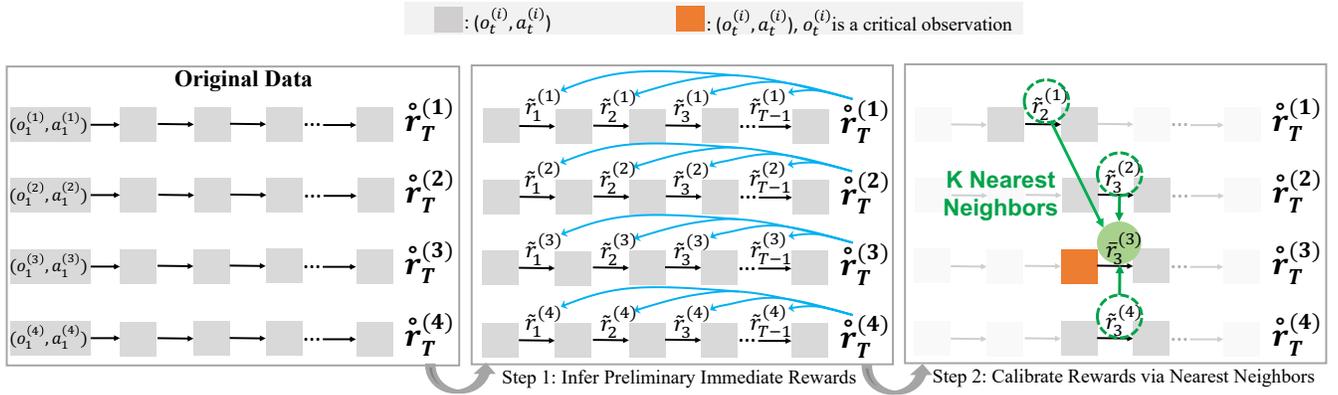


Figure 1: Illustration of reconstructing immediate rewards. First, the preliminary rewards are inferred from aggregated rewards. Then, the immediate rewards on critical observations are calibrated by averaging from nearest neighbors.

reward r_t and next observation o_{t+1} . In human-centric tasks, it is common that the immediate rewards r_t are unknown for most $t \in [1, T]$ in a trajectory. However, we assume that an aggregated reward $\hat{r}_t = \sum_{i=t-W}^t \gamma^{i-t+W+1} r_i$ is issued by the environment every W steps, quantifying the summed immediate rewards obtained between $(t-W)$ -th and t -th steps.

Off-Policy Evaluation in POMDP: the goal of OPE under POMDP is to estimate the expected total return over the *target* policy π , $V^\pi = \mathbb{E}[\sum_{t=1}^T \gamma^{t-1} r_t | a_t \sim \pi]$, using historical data collected over a *behavioral* policy $\beta \neq \pi$ deployed to the environment. Specifically, the historical data $\mathcal{D} = \{\dots, [\dots, (o_t, a_t, r_t, o'_t), \dots]^{(i)}, \dots\}_{i=1}^N$ consist of a set of N trajectories, where each trajectory is denoted as $\tau^{(i)}$.

Figure 1 illustrates how HOPE addresses the two major challenges of OPE in human-centric environments: partial observability and aggregated rewards. Specifically, it first infers preliminary immediate rewards from aggregated rewards (Step 1), and reconstructs immediate rewards r_t using nearest neighbors from trajectories \mathcal{D} (Step 2), which takes into account the uncertainty introduced by partial observability. The intuition of Step 2 follows findings presented in previous research [15, 26, 62], where humans with similar underlying states can present similar behaviors. Then, we use a general method, weighted importance sampling, to process the reconstructed rewards and estimate the expected total return V^π , since it is the most straightforward approach and can help isolate the source of improvements brought in by the immediate rewards reconstruction framework we propose.

3.1 Reconstruct Immediate Rewards via Nearest Neighbors

In human-centric environments, immediate rewards are generally not available, while the aggregated rewards themselves may not provide full information that can be directly used to estimate the expected total return of a target policy. Although prior works has noted that estimating immediate rewards can provide useful information toward training RL policies [5, 35], it has not been extensively investigated in the context of OPE; since prior OPE works generally assume immediate rewards are observable from the environment

or sparse. In many human-centric tasks, interaction outcomes can be gradually improved step-by-step and each immediate reward on each step can represent meaningful information by itself. Therefore, we propose to reconstruct immediate rewards from historical data, aiming to enrich the information provided for OPE.

In the OPE problem we consider, the underlying state is unobservable due to POMDP. And the aggregated rewards are assumed to be obtained at the end of an episode following $\hat{r} = \sum_{t=1}^T \gamma^t r_t$, with T being the horizon of the environment. In prior work, function approximation can be used to infer immediate rewards r_t from the aggregated rewards \hat{r} [5]. Consequently, we first produce rough per-step *preliminarily* reconstructed reward \tilde{r} taking as inputs observations and actions, $\tilde{r} = f(o_t, a_t | \theta)$, and trained toward, *i.e.*,

$$\min_{\theta} l(\theta) = \frac{1}{N} \sum_{i=1}^N \left(\hat{r} - \sum_{t=1}^T \gamma^{t-1} \tilde{r} \right)^2. \quad (1)$$

We keep the objective straightforward and use a standard training method in [4], so that the source of performance improvements can be isolated. For conciseness, we refer to \tilde{r} as *preliminary immediate rewards* throughout the rest of the paper. In our experiments described in Section 4.5, we compare it to treating the rewards as sparse, and the results present the effectiveness of the preliminary immediate rewards reconstruction procedure.

Consider that the mapping $f : \mathcal{O} \times \mathcal{A} \rightarrow \mathbb{R}$ learned above may not perfectly reconstruct the immediate reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, since the observation model $p(o|s)$ and state space \mathcal{S} are both assumed unknown. Previous research have found that humans with similar underlying state can be observed similar behaviors [62, 63], and some observations may provide critical information over others [56]. We follow such intuitions and introduce a remedy that uses \tilde{r}_t and its nearest neighbors, from trajectories that are similar to each other, to reconstruct immediate rewards r_t more accurately.

Specifically, to capture the information pertaining to the underlying states, we define the reconstructed immediate reward \hat{r}_t considering both average rewards associated with neighboring trajectories and observations that may provide more critical information over others, by which we call *critical observations*, *i.e.*,

$$\hat{r}_t = \mathbb{1}(o_t \in O^*) \bar{r}_t + \mathbb{1}(o_t \notin O^*) \tilde{r}_t, \quad (2)$$

where O^* is the set of critical observations, $\mathbb{1}(\cdot)$ is indicator function.

We first introduce how to identify the critical observations. Given that Q -functions, $Q^\pi(o, a)$, representing the expected return of taking action a at observation o , the difference of Q^π between any two actions (over the same o) can be used to quantify the magnitude of the difference in the final outcomes. As a result, we choose to leverage such Q -difference toward determining critical observations for OPE. A formal definition of a critical observation is given below.

DEFINITION 1 (CRITICAL OBSERVATION). *In a discrete action POMDP, observation o is a critical observation if there exists a constant h such that the maximum difference of $Q^\pi(o, a)$, $a \in \mathcal{A}$ is greater than h , i.e.,*

$$\max_{a', a'' \in \mathcal{A}} (Q^\pi(o, a') - Q^\pi(o, a'')) > h. \quad (3)$$

In Fig. 2, we illustrate how critical observations are identified using the threshold h . The need of critical observations is further justified by ablation study in Section 4.5.

Denote $\hat{r}_t^{(i)}$ as the reconstructed immediate reward at time-step t on trajectory $\tau^{(i)}$. We define the averaged reward $\bar{r}_t^{(i)}$ associated with critical observation $o_t^{(i)}$ as an average of preliminary immediate rewards \tilde{r} that occur at neighboring events following

$$\bar{r}_t^{(i)} = \frac{1}{K} \sum_k \tilde{r}_{t'}^{(k)}, \quad (4)$$

where the summation is performed over all $\tilde{r}_{t'}^{(k)}$ and K is the total number of nearest neighbours pertaining to the algorithm 1. Specifically, we define that two trajectories are neighbors if they have similar visitation distributions over the observation and action spaces, following

$$d(\tau^{(i)}, \tau^{(k)}) = \text{similarity}(o^{(i)}, o^{(k)}) + \text{similarity}(a^{(i)}, a^{(k)}) \quad (5)$$

where the similarity of distributions can be calculated by some measures such as Kullback-Leibler (KL) divergence [29]. Then we define that two observations are neighbors if they are from two neighboring trajectories, and share the most similar observations, since the moment that sharing similar observations from the similar trajectories may represent close underlying state of humans [62, 63]. To calculate similarity between two observations, one can use distance measures such as Euclidean. We can always find one and only one neighboring observation on a given neighboring trajectory, by breaking ties appropriately – such as selecting the earliest observation. Algorithm 1 summarizes how to find the K -nearest neighbors of $o_t^{(i)}$. In our experiment, the proposed distance to find nearest neighbors is compared to random selection in Section 4.5, and the results show the superior performance of the proposed distance.

3.2 HOPE

Based on the proposed solutions regarding the challenges within human-centric environment, we propose our *human-centric OPE (HOPE)* by incorporating estimated immediate rewards with weighted importance sampling (WIS). Let $w_i := \prod_{t=1}^T \pi(a_t^{(i)} | o_t^{(i)}) / \beta(a_t^{(i)} | o_t^{(i)})$ be an *importance weight*, which is the probability of the first T steps of $\tau^{(i)}$ under the *target policy*, π , divided by its probability under the behavior policy, β [44]. The estimated return of HOPE for an

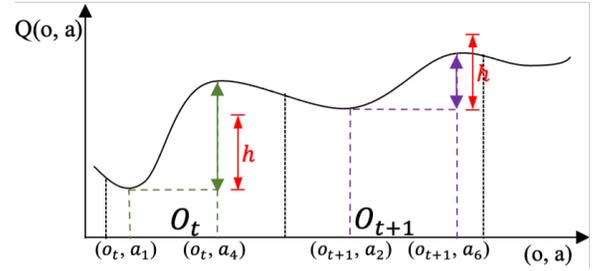


Figure 2: Identify critical observations based on the maximum difference of Q values. o_t is a critical observation with difference of Q greater than a threshold h .

Algorithm 1: Find K Nearest Neighbors of $o_t^{(i)}$.

Input: Historical data \mathcal{D} , hyper-parameter K .

Begin:

- 1: Initialize an empty set Z_t^i .
 - 2: **for** each trajectory $\tau^{(j)} \in \mathcal{D}$ **do**
 - 3: Calculate the distance $d(\tau^{(i)}, \tau^{(j)})$ following (5).
 - 4: **end for**
 - 5: Arrange $d(\tau^{(i)}, \tau^{(j)})$, $j \in N$ in non-decreasing order, take the first K elements and find the corresponding trajectories as *neighboring trajectories of $\tau^{(i)}$* .
 - 6: **for** each neighboring trajectory $\tau^{(k)}$ of $\tau^{(i)}$ **do**
 - 7: **for** each observation $o_{t'}^{(k)} \in \tau^{(k)}$ **do**
 - 8: Calculate the distance $l_{t'}^{ik}$ between $o_{t'}^{(k)}$ and $o_t^{(i)}$.
 - 9: **end for**
 - 10: Arrange $l_{t'}^{ik}$, $t' \in T$ in non-decreasing order, take the first element and find the corresponding observation $o_{t'}^{(k)}$ as *neighboring observation of $o_t^{(i)}$* .
 - 11: $Z_t^i \leftarrow Z_t^i \cup o_{t'}^{(k)}$
 - 12: **end for**
 - 13: Return the set Z_t^i .
-

target policy π is

$$V^\pi(\text{HOPE}) = \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i \sum_{t=1}^T \gamma^{t-1} \hat{r}_t^{(i)}, \quad (6)$$

We choose to use WIS here since it is the most straightforward approach. Moreover, this can help isolate the source of improvements brought in by immediate rewards reconstruction framework we propose. Prior work note that the lower variance of WIS may produce a larger reduction in expected square error than the additional error incurred due to the bias compared to some unbiased OPE such as importance sampling (IS) in practice [54]. Our real-world experimental results in Section 4 further support this.

In general, the importance weight assumes that the support of the evaluation policy π is a subset of the behavior policy β , which is enforced by Assumption 1 [53]:

ASSUMPTION 1. *If $\pi(a|o) \neq 0$, then $\beta(a|o) \neq 0$, where $a \in \mathcal{A}$, $o \in \mathcal{O}$.*

Upper and Lower Bounds of HOPE. We define δ^{ik} as the event where $\tau^{(k)}$ is a neighbor of $\tau^{(i)}$. Also, we define the counts of neighboring events as $K = \sum_{k=1}^N \mathbb{1}(\delta^{ik})$, where $\mathbb{1}(\cdot)$ is the indicator function. We then construct the nearest-neighbors matrix: $M^{ik} = \frac{\mathbb{1}(\delta^{ik})}{K}$. As the $N \times N$ matrix M^{ik} can be computed from the data and be used to compute the estimated immediate rewards for all observation-action pairs using the following proposition.

PROPOSITION 1. *For all transitions in the data, the estimated immediate rewards for corresponding observation-action pairs are given by*

$$\hat{r}_t^{(i)} = \sum_k M^{ik} \equiv [\mathbf{M}\mathbf{u}]_i \quad (7)$$

PROOF. A given reward on trajectory i at timestamp t , that averaging over all rewards on trajectory j at timestamp t' such that $\delta_{tt'}^{ij}$ holds, can be written as $\frac{1}{K} \sum_{(j,t'):\delta_{tt'}^{ij}} \delta_{tt'}^{ij} = \sum_{j,t'} \frac{\mathbb{1}(\delta_{tt'}^{ij})}{K} = \sum_{j,t'} M_{tt'}^{ij}$. Therefore, assume $u(o, a)$ is a function over the observation-action space and \mathbf{u} is the vector containing the quantity $u_i = u(o^{(i)}, a^{(i)})$ for every $(o^{(i)}, a^{(i)})$, the nearest-neighbor estimation of $u(o^{(i)}, a^{(i)})$ is given by $[\mathbf{M}\mathbf{u}]_i$ [19]. \square

In the case of $\tilde{r} \in [\tilde{r}_{lb}, \tilde{r}_{ub}]$, $\hat{r} \in [\hat{r}_{lb}, \hat{r}_{ub}]$ according to Proposition 1. Denote the returns of trajectory $\tau^{(i)}$ as $G(\tau^{(i)})$. Following [53], we also write $G(\tau^{(i)}) \in [0, 1]$ s.t. $G(\tau^{(i)}) := \frac{(\sum_{t=1}^T \gamma^{t-1} \tilde{r}_t^{(i)}) - G_{lb}}{G_{ub} - G_{lb}}$ to denote quantification of how good a trajectory $\tau^{(i)}$ is. Then the HOPE estimation is written by $\frac{1}{\sum_{n=1}^N w_n} \sum_{i=1}^N w_i G(\tau^{(i)}) = \frac{\sum_{i=1}^N w_i G(\tau^{(i)})}{\sum_{i=1}^N w_i}$. Then the upper and lower bound of HOPE estimation can be calculated using the following lemma, the proof of which is given in [53].

LEMMA 1. *Let π and β be any policy such that Assumption 1 holds, then for any constant integer $m \geq 1$, $\mathbb{E}[\prod_{t=1}^m \frac{\pi(a_t|o_t)}{\beta(a_t|o_t)} | \tau \sim \beta] = 1$.*

From Lemma 1, with the number of samples increasing, the denominator of HOPE tends towards n . HOPE estimator is bounded within $[0, 1]$, and so $\rho_{lb}^{HOPE}(\pi, \beta) = 0$ and $\rho_{ub}^{HOPE}(\pi, \beta) = 1$ for all π and β , where $\rho(\pi) = \mathbb{E}[G(\tau) | \tau \sim \beta]$.

Consistency of HOPE. We also show that HOPE is a consistent estimator of $\rho(\pi)$ if there is a single behavior policy (*Theorem 1*) or if there are multiple behavior policies that satisfy a technical requirement (*Theorem 2*), following work by [53].

THEOREM 1. *If Assumption 1 holds and there is only one behavior policy, then HOPE is a consistent estimator of $\rho(\pi)$.*

PROOF. When there is only one behavior policy, HOPE estimation can be rewrote as $\frac{\frac{1}{N} \sum_{i=1}^N w_i G(\tau^{(i)})}{\frac{1}{N} \sum_{i=1}^N w_i}$ by multiplying both its numerator and denominator by $\frac{1}{N}$. Then the numerator is equal to IS^π , which is a consistent estimator of $\rho(\pi)$ as proved in prior work [44, 53]), thus the numerator converges almost surely to $\rho(\pi)$. For the denominator, by Lemma 1, we have that

$$\mathbb{E}\left[\prod_{t=1}^T \frac{\pi(a_t^{(i)}|o_t^{(i)})}{\beta(a_t^{(i)}|o_t^{(i)})}\right] = 1, \text{ for all } i \in [1, \dots, N]. \quad (8)$$

Each term $\prod_{t=1}^T \frac{\pi(a_t^{(i)}|o_t^{(i)})}{\beta(a_t^{(i)}|o_t^{(i)})}$ is identically distributed for each $i \in [1, \dots, N]$, as there is only one behavior policy. By the Khintchine strong law of large numbers [47], we have that

$$\frac{1}{N} \sum_{i=1}^N \prod_{t=1}^T \frac{\pi(a_t^{(i)}|o_t^{(i)})}{\beta(a_t^{(i)}|o_t^{(i)})} \xrightarrow{\text{a.s.}} 1. \quad (9)$$

By the property of almost sure convergence [23], HOPE converges almost surely to $\rho(\pi)$, and so HOPE is a consistent estimator of $\rho(\pi)$. \square

We also provide the proof for the consistency of HOPE when there are multiple behavior policies.

THEOREM 2. *If Assumption 1 holds and there exists a constant $\epsilon > 0$ such that $\beta_i(a|o) \geq \epsilon$ for all $i \in \{1, \dots, N\}$ and $(a|o)$ where $\pi(a|o) \neq 0$, then HOPE is a consistent estimator of $\rho(\pi)$ if there are multiple behavior policies.*

PROOF. When there are multiple behavior policies, similar to the proof of Theorem 1, the numerator is equal to IS^π . The numerator converges almost surely to $\rho(\pi)$. For the denominator, by Lemma 1 we have that Equation 8 holds. Each term $\prod_{t=1}^T \frac{\pi(a_t^{(i)}|o_t^{(i)})}{\beta_i(a_t^{(i)}|o_t^{(i)})} \in [0, \frac{1}{\epsilon^T}]$ and therefore has bounded variance. By the Kolmogorov strong law of large numbers [47], we have that almost surely convergence 9 holds. Therefore, HOPE is a consistent estimator of $\rho(\pi)$ if there are multiple behavior policies. \square

4 EXPERIMENTS

We conduct experiments on two real-world human-centric tasks, sepsis treatment, and intelligent tutoring, to validate our proposed approach. Since our focus is the human-centric environments, we don't conduct experiments on control tasks such as D4RL [11].

4.1 Benchmarks

We use *nine* state-of-the-art benchmarks, which cover a variety of approaches that have been explored for OPE: *Four IPS methods*, including IS [44], WIS [44], Per-Decision IS (PDIS) [44], and Per-Horizon WIS (PHWIS) [10]. For PHWIS, we follow the PHWIS-Behavior as in [10], as we assume the lengths of the trajectories do not depend on the policy that is used to generate them. *Two direct methods*, Fitted Q Evaluation (FQE) [30] and Dual stationary

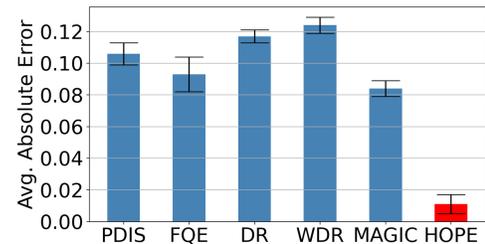


Figure 3: Average absolute error and standard deviation from synthetic sepsis environment ($\gamma=0.99$). IS, WIS, PHWIS, DualDICE are unable to select the best policy.

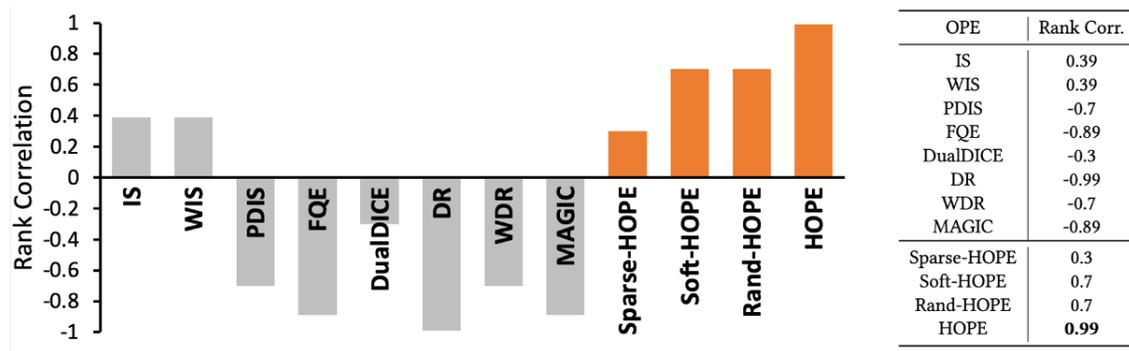


Figure 4: Spearman’s rank correlation coefficient across policies from real-world medical system ($\gamma=0.99$). PHWIS is unable to produce a meaningful result probably due to the highly varied lengths of trajectories.

Distribution Correction Estimation (DualDICE) [37]. For FQE, as in [30], we train a neural network to estimate the value of the evaluation policy π_e by bootstrapping from $Q(o', a')$. For DualDICE, we use the open-sourced code in its original paper. *Three hybrid methods*, including Doubly Robust (DR), Weighted DR (WDR) [52], and MAGIC [52], which finds an optimal linear combination among a set that varies the switch point between WDR and direct methods. For MAGIC, we use the implementation of [60].

4.2 Validating OPE Performance

In this work, we use two types of procedures to validate the performance of OPE methods. We use standard metrics including absolute error, $\text{regret}@1$, and Spearman’s rank correlation coefficient [49] that are commonly used in prior OPE approaches. Moreover, we use the human-centric significance test to measure *the statistical significance* between the OPE-estimated returns across different policies. For each evaluation policy, we use bootstrapping by episodes as introduced in [21]. For e-learning, we also empirically evaluated the induced policies. As for many human-centric tasks, one key measurement for the RL-induced policy is whether they significantly outperform the current expert policy [65]. Therefore, we conduct a t-test over OPE estimations obtained from bootstrapping. It measures whether there is a significant difference between the mean value of OPE estimations on one policy against another.

4.3 Sepsis Treatment

Sepsis is the leading cause of mortality and the most expensive condition associated with in-hospital stay [32]. Sepsis treatment is a highly challenging problem and has raised tremendous investigation [13, 14, 39].

4.3.1 Synthetic Sepsis Environment. We use a sepsis model [41] in the management of sepsis in ICU patients. Following the settings from [38], the discrete observation space consists of a binary indicator for diabetes, and four vital signs (heart rate, blood pressure, oxygen concentration, glucose level) that take values in a subset of {very_high, high, normal, low, very_low}. The simulated environment contains a total of 1440 observations, and 8 actions characterized by assigning a binary value (0 or 1) toward each option in {antibiotics, vasopressors, mechanical_ventilation}.

The simulation continues either until at most $T = 5$ (horizon) time steps (0 rewards), death (-1 reward), or discharge (+1 reward). Patients are discharged when all vital signs are in the normal range without treatment. Patients die if at least three vitals are out of the normal range. We use the behavior policy and three evaluation policies following [38]. Figure 3 shows that HOPE performs the best in terms of average absolute error across evaluation policies.

4.3.2 Real-World Medical System. In our experiment, we use Electronic Health Records (EHRs) collected from a large hospital in the United States with overall 221,700 visits patients over two years. The observation space consists of 15 continuous sepsis-related clinical attributes, including seven vital signs {HeartRate, RespiratoryRate, PulseOx, SystolicBP, DiastolicBP, MAP, Temperature} and eight lab analytes {Bands, BUN, Lactate, Platelet, Creatinine, BiliRubin, WBC, FIO2}. The size of action space is 4 with two binary treatment options {antibiotic_administration, oxygen_assistance}. Four stages of sepsis are defined by the clinicians, and the rewards are set for each stage: infection (± 5), inflammation (± 10), organ failure (± 20), and septic shock (± 50). The designated negative rewards are given when a patient enters the corresponding stage and positive rewards are given back when the patient recovers from the stage. The collected trajectories’ lengths range from 1 to 1160. We assume that the clinical care team is well-trained with medical knowledge and follows standard protocols in sepsis treatments, thus we learn the expert policy as introduced in [5]. We train policies using Deep Q Network (DQN) [36] with varied hyperparameters and select five as evaluation policies. As prior work in sepsis research [5, 28, 45] identifies septic shock rate as an important criterion for learning policies, we calculate Spearman’s rank correlation coefficient between the policies’ ranks using estimated values given by OPE and the actual policies’ ranks in terms of septic shock rates.

Figure 4 shows the results of HOPE and benchmarks. The grey-shaded columns represent the benchmark results, and the orange-shaded columns represent the results from HOPE and its variations. Overall, HOPE performs the best in terms of rank correlation. Interestingly, we notice that IS and WIS outperform other benchmarks, while they can be suffering from long-horizon in prior theoretical work [31]. A possible reason is that both methods benefit more from the reduction in expected error than the variance incurred

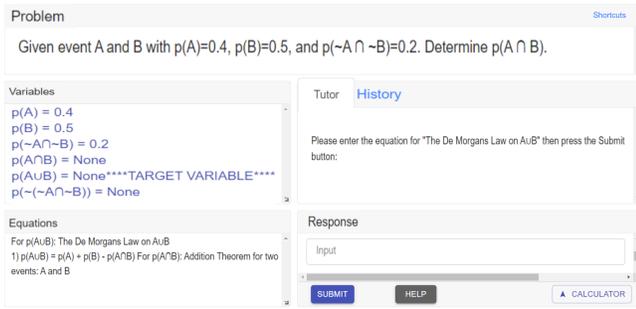


Figure 5: Our ITS GUI. The problem statement window (top) presents the statement of the problem. The dialog window (middle right) shows the message the tutor provides to the students. Responses, e.g., writing an equation, are entered in the response window (bottom right). Any variables and equations generated through this process are shown on the variable window (middle left) and equation window (bottom left).

due to horizon under our real-world settings. Similar findings are reported in some long-horizon environments [12].

4.4 Real-World Intelligent Tutor

Intelligent Tutoring Systems (ITSs) are computer systems that mimic aspects of human tutors and have also been shown to be successful [1, 2, 59]. They aim to provide instruction or feedback to support students’ learning, which is an important application of RL to improve students’ engagement and learning outcomes. We use a web-based ITS which teaches computer science students probability knowledge, covering ten major principles such as the complement theorem. Students’ interaction logs are collected over seven semesters of classroom studies (including 1,307 students) in an undergraduate computer science course at a large public university in the United States. Figure 5 shows the GUI of the tutor.

During tutoring, there are many factors that might determine or indicate students’ learning state, but many of them are not well understood by educators. Thus, to be conservative, we extract varieties of attributes that might determine or indicate student learning observations from student-system interaction logs. In sum, 142 attributes with both discrete and continuous values are extracted, which can be categorized into the following five groups: (i) **Autonomy (10 features)**: the amount of work done by the student, such as the number of times the student restarted a problem; (ii) **Temporal Situation (29 features)**: the time-related information about the work process, such as average time per step; (iii) **Problem-Solving (35 features)**: information about the current problem-solving context, such as problem difficulty; (iv) **Performance (57 features)**: information about the student’s performance during problem-solving, such as percentage of correct entries; (v) **Hints (11 features)**: information about the student’s hint usage, such as the total number of hints requested.

The agent will make 10 decisions: for each problem, the agent will decide whether the student should *solve* the next problem, *study* a solution provided by the tutor, or *work together* with the tutor to solve the problem. The rewards are obtained after all problems are

accomplished, which is defined as the students’ normalized learning gain calculated by the two test scores that students took before and after the experiments [9], respectively. A total of four policies, including three DQN-induced policies (denoted as π_1, π_2, π_3) and one expert policy (denoted as π_{expert}), are deployed to the ITS used by students. The log data from students in the prior six semesters are used to train policies and the following semester to test.

Figure 6 shows the results of comparing HOPE (the last row) with the nine original OPE benchmarks (the top section of the left table) using both standard validation methods and signed significance tests for OPE. Overall, HOPE (the last row) outperforms all nine benchmarks in terms of average absolute error (AAE, column 2), regret@1 (column 3), and rank correlation (column 4). There is no clear winner among the nine original OPE benchmarks.

Column 5 in Table 6 shows the best policy determined by each OPE. While all nine original OPE benchmarks select other sub-optimal policies as the best policy, HOPE is the only method that successfully identifies π_1 to be the best policy in the empirical study. More importantly, while the nine original OPE benchmarks predict that their selected best policy would significantly outperform the Expert policy (column 6), the empirical results (the 7th/last column) show no significant difference was found. The offline significance test using HOPE, however, perfectly aligns with the empirical result, that π_1 is significantly different from the Expert policy in both OPE t-test (column 6) and the empirical t-test (column 7).

4.5 Ablation Studies

For a better understanding of our proposed approach to tackle the partial observability and missing immediate rewards in real-world human-centric environments, we conduct three ablation studies:

(i) *Sparse-HOPE*. One variation of our proposed approach is assuming the preliminary rewards are sparse and calibrating immediate rewards via nearest neighbors, named Sparse-HOPE. Figure 4 and Figure 6 show that Sparse-HOPE can outperform all benchmarks, except IS, in terms of rank correlation. In real-world intelligent tutors, Sparse-HOPE outperforms all benchmarks in terms of average absolute error and regret@1. Those indicate that our proposed nearest-neighbors-based immediate rewards reconstruction is effective for estimating the return of a policy. On the other hand, Sparse-HOPE performs worse than HOPE, which could indicate the importance of considering rewards as aggregated and the effectiveness of our preliminary rewards reconstruction.

(ii) *Soft-HOPE*. We define another variation of HOPE, named Soft-HOPE, which assumes decisions made on any observation could contribute equally to the final outcomes, i.e.

$$V^\pi(\text{Soft-HOPE}) = \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i \sum_{t=1}^T \gamma^{t-1} \left(\frac{1}{K} \sum_k r_{i'}^{s(k)} \right). \quad (10)$$

Note that it performs neighbors-based estimation on all observations, as opposed to (6) which only estimates neighbors-based immediate rewards on critical observations. Figure 4 shows that Soft-HOPE outperforms all benchmarks from real-world medical systems in terms of rank correlation, and from intelligent tutoring in terms of average absolute error. However, it performs worse than most benchmarks in terms of regret@1 and rank correlation from intelligent tutoring. A possible reason is using averaged rewards

OPE	Standard metrics			Human-centric significance test		
	AAE	Regret@1	Rank	π^{Best}	OPE t-test (p)	Empirical t-test (p)
IS	6.42	0.46	0.8	π_2	7.071 (0.000)	0.612 (0.543)
WIS	0.22	0.46	0.2	π_2	17.081 (0.000)	0.612 (0.543)
PDIS	6.42	0.46	0.8	π_2	7.071 (0.000)	0.612 (0.543)
PHWIS	0.22	0.46	0.2	π_2	17.081 (0.000)	0.612 (0.543)
FQE	0.10	0.46	-0.8	π_2	12.546 (0.000)	0.612 (0.543)
DualDICE	0.33	N/A	N/A	N/A	N/A	N/A
DR	18.5	0.60	0.8	π_3	4.893 (0.000)	0.203 (0.840)
WDR	1.63	0.60	-0.8	π_3	8.344 (0.000)	0.203 (0.840)
MAGIC	17.3	0.60	0.6	π_3	20.750 (0.000)	0.203 (0.840)
Sparse-HOPE	0.06	0.00	0.4	π_1^*	10.518 (0.000)	2.011 (0.040)
Soft-HOPE	0.07	0.00	0.8	π_1^*	1.332 (0.186)	2.011 (0.040)
Rand-HOPE	0.10	0.00	0.8	π_1^*	6.902 (0.000)	2.011 (0.040)
HOPE	0.03	0.00	0.8	π_1^*	11.670 (0.000)	2.011 (0.040)

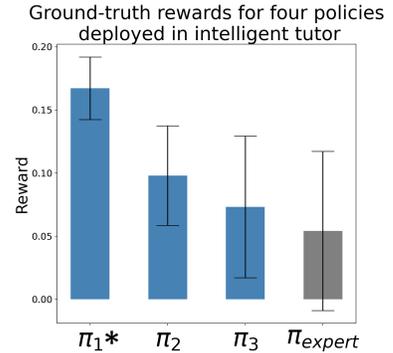


Figure 6: Validating OPE performance (left) and empirical results (right) from real-world intelligent tutoring ($\gamma = 0.9$). Left: Standard metrics include average absolute error (AAE), regret@1, and Spearman’s rank correlation coefficient (Rank). The best results on each metric are bolded. Human-centric significance test results include the best policy selected by OPE, offline significance test with bootstrapping on the best policy selected by OPE and expert policy, and empirical significance test on these two policies, at the level of $p < 0.05$. The OPE significance test that aligns with the empirical test is bolded. DualDICE estimates the performance of all policies equally, thus its rank-related results are unavailable. Right: Ground-truth rewards for four policies. π_1 receives the highest average reward and is the only policy that differs significantly from π_{expert} ($p = 0.04$), as indicated by the asterisk.

on all observations could introduce noise to OPE and weaken its estimation of ranking. Moreover, it performs worse than HOPE from both environments, which indicates that our defined critical observations can help extract meaningful information for OPE.

(iii) *Rand-HOPE*. The third variation of our proposed approach is randomly selecting neighbors instead of using our defined distance for immediate rewards reconstruction, which we call Rand-HOPE. We repeat Rank-HOPE 100 times and report average results. Figure 4 and Figure 6 show that Rand-HOPE outperforms all benchmarks in both real-world environments. A possible reason is that inferring preliminary immediate rewards can provide much more useful information than sparse rewards, thus even randomly averaged rewards would perform better than using sparse rewards. Rand-HOPE performs worse than HOPE, which indicates that our defined distance is more accurate to reconstruct immediate rewards for OPE in e-learning and healthcare.

Moreover, in real-world intelligent tutoring, HOPE and its three variations, are the only methods that successfully select the best policy. Table 1 further shows the mean and standard error on policies π_1 and π_{expert} estimated by HOPE-related methods. HOPE achieves the best estimation that is closest to ground truth.

5 CONCLUSION & SOCIAL IMPACT

In this work, we proposed an approach, *HOPE*, for OPE in real-world human-centric environments with partial observability and aggregated rewards. It first inferred preliminary immediate rewards from historical observations, then used nearest neighbor methods to fully reconstruct immediate rewards. We also introduced critical observations, that can impact final outcomes of a trajectory over others,

	π_1 Mean \pm se	π_{expert} Mean \pm se
Sparse-HOPE	0.094 \pm 0.01	0.005 \pm 0.00
Soft-HOPE	0.021 \pm 0.00	0.004 \pm 0.00
Rand-HOPE	0.140 \pm 0.02	-0.023 \pm 0.02
HOPE	0.176\pm0.01	0.008\pm0.00
Empirical result	0.167 \pm 0.02	0.054 \pm 0.06

Table 1: Mean and standard error with bootstrapping on policies π_1 and π_{expert} from real-world intelligent tutoring. HOPE achieves the best estimation on both policies.

to enrich provided information for OPE. We conducted extensive real-world experiments with two challenging tasks for OPE, sepsis treatment and intelligent tutoring, using both standard validations and human-centric significance tests to validate OPE. The results showed that HOPE outperformed state-of-the-art benchmarks in both applications. A part of our methodology leverages WIS, which may introduce variance to estimations. We kept it straightforward such that the performance can be easily isolated. In the future, WIS can be replaced with DR or DICE for reduced variance.

All real-world data were obtained anonymously through an exempt IRB-approved protocol and were scored using established rubrics. No demographic data or class grades were collected. All data were shared within the research group under IRB, and were de-identified and automatically processed for labeling. This research seeks to remove societal harms that come from lower engagement and retention of students who need more personalized interventions and developing more robust medical interventions for patients.

ACKNOWLEDGMENTS

This research was supported by the NSF Grants: Integrated Data-driven Technologies for Individualized Instruction in STEM Learning Environments (1726550), CAREER: Improving Adaptive Decision Making in Interactive Learning Environments (1651909), and Generalizing Data-Driven Technologies to Improve Individualized STEM Instruction by Intelligent Tutors (2013502).

REFERENCES

- [1] Mark Abdelshihed, Mehak Maniktala, Song Ju, Ayush Jain, Tiffany Barnes, and Min Chi. 2021. Preparing Unprepared Students For Future Learning. In *Proceedings of the 43rd annual conference of the cognitive science society*, Vol. 43.
- [2] Mark Abdelshihed, Guojing Zhou, Mehak Maniktala, Tiffany Barnes, and Min Chi. 2020. Metacognition and Motivation: The Role of Time-Awareness in Preparation for Future Learning. In *Proceedings of the 42nd annual conference of the cognitive science society*, Vol. 42.
- [3] Kenshi Abe and Yusuke Kaneko. 2021. Off-Policy Exploitability-Evaluation in Two-Player Zero-Sum Markov Games. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 78–87.
- [4] Markel Sanz Ausin, Hamoon Azizsoltani, Song Ju, Yeo Jin Kim, and Min Chi. 2021. Infernet for delayed reinforcement tasks: Addressing the temporal credit assignment problem. In *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 1337–1348.
- [5] Hamoon Azizsoltani and Yeo Jin. 2019. Unobserved Is Not Equal to Non-existent: Using Gaussian Processes to Infer Immediate Rewards Across Contexts. In *In Proceedings of the 28th International Joint Conference on Artificial Intelligence*.
- [6] Andrew Bennett, Nathan Kallus, Lihong Li, and Ali Mousavi. 2021. Off-policy evaluation in infinite-horizon reinforcement learning with latent confounders. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1999–2007.
- [7] Anthony R Cassandra. 1998. A survey of POMDP applications. In *Working notes of AAAI 1998 fall symposium on planning with partially observable Markov decision processes*, Vol. 1724.
- [8] Yash Chandak, Scott Niekum, Bruno da Silva, Erik Learned-Miller, Emma Brunskill, and Philip S Thomas. 2021. Universal off-policy evaluation. *Advances in Neural Information Processing Systems* 34 (2021).
- [9] Min Chi, Kurt VanLehn, Diane Litman, and Pamela Jordan. 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 (2011), 137–180.
- [10] Shayan Doroudi, Philip S Thomas, and Emma Brunskill. 2017. Importance Sampling for Fair Policy Selection. *Grantee Submission* (2017).
- [11] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219* (2020).
- [12] Justin Fu, Mohammad Norouzi, Ofir Nachum, George Tucker, Ziyu Wang, Alexander Novikov, Mengjiao Yang, Michael R Zhang, Yutian Chen, Aviral Kumar, et al. 2021. Benchmarks for deep off-policy evaluation. *arXiv preprint arXiv:2103.16596* (2021).
- [13] Ge Gao, Qitong Gao, Xi Yang, Miroslav Pajic, and Min Chi. 2022. A Reinforcement Learning-Informed Pattern Mining Framework for Multivariate Time Series Classification. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [14] Ge Gao, Farzaneh Khoshnevisan, and Min Chi. 2022. Reconstructing Missing EHRs Using Time-Aware Within-and Cross-Visit Information for Septic Shock Early Prediction. In *2022 IEEE 10th International Conference on Healthcare Informatics (ICHI)*. IEEE Computer Society, 151–162.
- [15] Ge Gao, Samiha Marwan, and Thomas W Price. 2021. Early performance prediction using interpretable patterns in programming process data. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 342–348.
- [16] Qitong Gao, Ge Gao, Min Chi, and Miroslav Pajic. 2023. Variational Latent Branching Model for Off-Policy Evaluation. In *The Eleventh International Conference on Learning Representations*.
- [17] Qitong Gao, Stephen L Schmidt, Karthik Kamaravelu, Dennis A Turner, Warren M Grill, and Miroslav Pajic. 2022. Offline Policy Evaluation for Learning-based Deep Brain Stimulation Controllers. In *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCP)*. IEEE, 80–91.
- [18] Qitong Gao, Dong Wang, Joshua D Amason, Siyang Yuan, Chenyang Tao, Ricardo Henao, Majda Hadziahmetovic, Lawrence Carin, and Miroslav Pajic. 2022. Gradient Importance Learning for Incomplete Observations. *International Conference on Learning Representations* (2022).
- [19] Omer Gottesman, Joseph Futoma, Yao Liu, Sonali Parbhoo, Leo Celi, Emma Brunskill, and Finale Doshi-Velez. 2020. Interpretable off-policy evaluation in reinforcement learning by highlighting influential transitions. In *International Conference on Machine Learning*. PMLR, 3658–3667.
- [20] Joy Paul Guilford. 1950. *Fundamental statistics in psychology and education*. (1950).
- [21] Botao Hao, Xiang Ji, Yaqi Duan, Hao Lu, Csaba Szepesvári, and Mengdi Wang. 2021. Bootstrapping statistical inference for off-policy evaluation. *arXiv preprint arXiv:2102.03607* (2021).
- [22] Alexander Irpan, Kanishka Rao, Konstantinos Bousmalis, Chris Harris, Julian Ibarz, and Sergey Levine. 2019. Off-policy evaluation via off-policy classification. *Advances in Neural Information Processing Systems* 32 (2019).
- [23] Jiming Jiang. 2010. *Large sample techniques for statistics*. Springer.
- [24] Nan Jiang and Lihong Li. 2016. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*. PMLR, 652–661.
- [25] Song Ju, Shitian Shen, Hamoon Azizsoltani, Tiffany Barnes, and Min Chi. 2019. Importance Sampling to Identify Empirically Valid Policies and their Critical Decisions.. In *EDM (Workshops)*. 69–78.
- [26] Ramtin Keramati, Omer Gottesman, Leo Anthony Celi, Finale Doshi-Velez, and Emma Brunskill. 2022. Identification of Subgroups With Similar Benefits in Off-Policy Policy Evaluation. In *Conference on Health, Inference, and Learning*. PMLR, 397–410.
- [27] Farzaneh Khoshnevisan and Min Chi. 2021. Unifying domain adaptation and domain generalization for robust prediction across minority racial groups. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 521–537.
- [28] Matthieu Komorowski, Leo A Celi, Omar Badawi, Anthony C Gordon, and A Aldo Faisal. 2018. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature medicine* 24, 11 (2018), 1716–1720.
- [29] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.
- [30] Hoang Le, Cameron Voloshin, and Yisong Yue. 2019. Batch policy learning under constraints. In *International Conference on Machine Learning*. PMLR, 3703–3712.
- [31] Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. 2018. Breaking the curse of horizon: Infinite-horizon off-policy estimation. *Advances in Neural Information Processing Systems* 31 (2018).
- [32] Vincent Liu, Gabriel J Escobar, John D Greene, Jay Soule, Alan Whippy, Derek C Angus, and Theodore J Iwashyna. 2014. Hospital deaths in patients with sepsis from 2 independent cohorts. *Jama* 312, 1 (2014), 90–92.
- [33] Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. 2014. Offline policy evaluation across representations with applications to educational games. In *AAMAS*, Vol. 1077.
- [34] Samiha Marwan, Ge Gao, Susan Fisk, Thomas W Price, and Tiffany Barnes. 2020. Adaptive immediate feedback can improve novice programming engagement and intention to persist in computer science. In *Proceedings of the 2020 ACM conference on international computing education research*. 194–203.
- [35] Maja J Mataric. 1994. Reward functions for accelerated learning. In *Machine learning proceedings 1994*. Elsevier, 181–189.
- [36] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [37] Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. 2019. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. *Advances in Neural Information Processing Systems* 32 (2019).
- [38] Hongseok Namkoong, Ramtin Keramati, Steve Yadlowsky, and Emma Brunskill. 2020. Off-policy policy evaluation for sequential decisions under unobserved confounding. *Advances in Neural Information Processing Systems* 33 (2020), 18819–18831.
- [39] Allen Nie, Yannis Flet-Berliac, Deon Richmond Jordan, William Steenberg, and Emma Brunskill. 2022. Data-Efficient Pipeline for Offline Reinforcement Learning with Limited Data. In *Advances in Neural Information Processing Systems*.
- [40] Xinkun Nie, Emma Brunskill, and Stefan Wager. 2021. Learning when-to-treat policies. *J. Amer. Statist. Assoc.* 116, 533 (2021), 392–409.
- [41] Michael Oberst and David Sontag. 2019. Counterfactual off-policy evaluation with gumbel-max structural causal models. In *International Conference on Machine Learning*. PMLR, 4881–4890.
- [42] Emily Miller Payne, Russ Hodges, and Elda Patricia Hernandez. 2017. Changing Demographics and Needs Assessment for Learning Centers in the 21st Century. *Learning Assistance Review* 22, 1 (2017), 21–36.
- [43] Michael JD Powell and J Swann. 1966. Weighted uniform sampling—a Monte Carlo technique for reducing variance. *IMA Journal of Applied Mathematics* 2, 3 (1966), 228–236.
- [44] Doina Precup. 2000. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series* (2000), 80.
- [45] Aniruddh Raghu, Matthieu Komorowski, Imran Ahmed, Leo Celi, Peter Szolovits, and Marzyeh Ghassemi. 2017. Deep reinforcement learning for sepsis treatment. *arXiv preprint arXiv:1711.09602* (2017).
- [46] Desik Rengarajan, Gargi Vaidya, Akshay Sarvesh, Dileep Kalathil, and Srinivas Shakkottai. 2022. Reinforcement Learning with Sparse Rewards using Guidance from Offline Demonstration. *arXiv preprint arXiv:2202.04628* (2022).

- [47] Pranab K Sen and Julio M Singer. 1994. *Large sample methods in statistics: an introduction with applications*. Vol. 25. CRC press.
- [48] Shitian Shen and Min Chi. 2016. Reinforcement learning: the sooner the better, or the later the better?. In *Proceedings of the 2016 conference on user modeling adaptation and personalization*. 37–44.
- [49] Charles Spearman. 1987. The proof and measurement of association between two things. *The American journal of psychology* 100, 3/4 (1987), 441–471.
- [50] Guy Tennenholtz, Uri Shalit, and Shie Mannor. 2020. Off-policy evaluation in partially observable environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 10276–10283.
- [51] Joel Tetreault and Diane Litman. 2006. Comparing the utility of state features in spoken dialogue using reinforcement learning. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*. 272–279.
- [52] Philip Thomas and Emma Brunskill. 2016. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2139–2148.
- [53] Philip S Thomas. 2015. Safe reinforcement learning. (2015).
- [54] Philip S Thomas, Scott Niekum, Georgios Theodorou, and George Konidaris. 2015. Policy evaluation using the Ω -return. *Advances in Neural Information Processing Systems* 28 (2015).
- [55] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 5026–5033.
- [56] Lisa Torrey and Matthew Taylor. 2013. Teaching on a budget: Agents advising agents in reinforcement learning. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. 1053–1060.
- [57] Masatoshi Uehara, Jiawei Huang, and Nan Jiang. 2020. Minimax weight and q-function learning for off-policy evaluation. In *International Conference on Machine Learning*. PMLR, 9659–9668.
- [58] Masatoshi Uehara, Haruka Kiyohara, Andrew Bennett, Victor Chernozhukov, Nan Jiang, Nathan Kallus, Chengchun Shi, and Wen Sun. 2022. Future-Dependent Value-Based Off-Policy Evaluation in POMDPs. *arXiv preprint arXiv:2207.13081* (2022).
- [59] Kurt VanLehn. 2006. The Behavior of Tutoring Systems. *International Journal of Artificial Intelligence in Education* 16, 3 (2006), 227–265.
- [60] Cameron Voloshin, Hoang M Le, Nan Jiang, and Yisong Yue. 2019. Empirical study of off-policy policy evaluation for reinforcement learning. *arXiv preprint arXiv:1911.06854* (2019).
- [61] Mengjiao Yang, Bo Dai, Ofir Nachum, George Tucker, and Dale Schuurmans. 2022. Offline policy selection under uncertainty. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 4376–4396.
- [62] Xi Yang, Guojing Zhou, Michelle Taub, Roger Azevedo, and Min Chi. 2020. Student Subtyping via EM-Inverse Reinforcement Learning. *International Educational Data Mining Society* (2020).
- [63] Changchang Yin, Ruoqi Liu, Dongdong Zhang, and Ping Zhang. 2020. Identifying sepsis subphenotypes via time-aware multi-modal auto-encoder. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 862–872.
- [64] Guojing Zhou, Hamoon Azizoltani, Markel Sanz Ausin, Tiffany Barnes, and Min Chi. 2022. Leveraging granularity: Hierarchical reinforcement learning for pedagogical policy induction. *International journal of artificial intelligence in education* 32, 2 (2022), 454–500.
- [65] Guojing Zhou, Xi Yang, Hamoon Azizoltani, Tiffany Barnes, and Min Chi. 2020. 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*. 284–292.