Value Driven Representation for Human-in-the-Loop Reinforcement Learning

Ramtin Keramati

Institute of Computational and Mathematical Engineering
Stanford, California
keramati@cs.stanford.edu

Emma Brunskill

Department of Computer Science
Stanford, California
ebrun@cs.stanford.edu

ABSTRACT

Interactive adaptive systems powered by Reinforcement Learning (RL) have many potential applications, such as intelligent tutoring systems. In such systems there is typically an external human system designer that is creating, monitoring and modifying the interactive adaptive system, trying to improve its performance on the target outcomes. In this paper we focus on algorithmic foundation of how to help the system designer choose the set of sensors or features to define the observation space used by reinforcement learning agent. We present an algorithm, value driven representation (VDR), that can iteratively and adaptively augment the observation space of a reinforcement learning agent so that is sufficient to capture a (near) optimal policy. To do so we introduce a new method to optimistically estimate the value of a policy using offline simulated Monte Carlo rollouts. We evaluate the performance of our approach on standard RL benchmarks with simulated humans and demonstrate significant improvement over prior baselines.

CCS CONCEPTS

• **Human-centered computing** → *Human computer interaction (HCI)*.

KEYWORDS

Reinforcement Learning, Human-in-the-Loop

ACM Reference Format:

Ramtin Keramati and Emma Brunskill. 2019. Value Driven Representation for Human-in-the-Loop Reinforcement Learning. In *27th Conference on User Modeling, Adaptation and Personalization (UMAP '19)*, June 9–12, 2019, Larnaca, Cyprus. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3320435.3320471>

1 INTRODUCTION

Interactive adaptive systems powered by reinforcement learning can improve over time and have many potential applications. These include intelligent tutoring systems that improve their teaching as they instruct more students, smart home devices that adjust temperature settings in response to weather and human preferences,

Figure 1: We consider interactive adaptive systems which may interact with people (person on the right) and which have a system designer (person on the left) that is trying to optimize the adaptive system. We look at a case where the adaptive system proactively provides input to the designer (outgoing arrow) to propose potential modifications to its sensors that may help improve performance.

and mobile wellness applications that support chronic care management. In such cases the system is learning a good decision policy—a mapping from a current observation (of a student, of a home, of a person and their context) to action (what pedagogical activity to propose, how to set the temperature, what health nudge to suggest) in order to maximize overall outcomes over time (how much student has learned, minimize energy, maximize total activity).

In such systems there is typically an external system designer that is creating, monitoring and modifying the interactive adaptive system, trying to improve its performance on the target outcomes (Figure 1). A key question is how interactive adaptive systems can better provide input back to their system designers to help support the designer in improving the performance. In this paper we focus on one particular direction for this, how to help the system designer choose the set of sensors or features to define the observation space used by the interactive system to make decisions. For example, does the tutoring system only log student responses to problems, or is a webcam also used to detect frustration levels?

The observation space specification implicitly constrains the decision policy class the interactive system can optimize over. Here we consider the case when the interactive system itself can monitor its own performance, and propose to the system designer some potential modifications to the observation space that it expects might yield improved performance. The human designer can then choose to augment the system with additional features and/or sensors that the adaptive system can then use going forward. For example, the system may recognize that there are number of observations which are currently identical (for example, same performance on a set of problems) where selecting the action (next problem) does not seem to yield the best outcomes, and ask the human if there might be an additional feature/ sensor that could be used to distinguish such currently aliased observations. This set of observations can provide

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP '19, June 9–12, 2019, Larnaca, Cyprus

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6021-0/19/06...\$15.00
<https://doi.org/10.1145/3320435.3320471>

information to the human designer about what types of features/sensors might be useful to add to the system, like a webcam and emotion classifier to detect frustration.

In this short paper we focus on the algorithmic foundations of this idea, providing a proof of concept in simulated domains. Precisely, we present an algorithm, value driven representation (VDR), that can iteratively augment the observation space of a reinforcement learning agent (such as an interactive adaptive system), if the algorithm estimates that the resulting augmented observation space could yield an improved policy performance in the real environment. VDR can be applied both to situations where additional features could be added later in a demand driven way (from a human system designer), or when the full set of features is known in advance but there are computational, performance and interpretability benefits to starting with a more compact representation.

Our approach starts with a coarse observation space (a small minimal set of features). Note that this small subset may only be known because these are the initial features seem to be relevant by a human system designer or due to cost or other constraints. We assume the algorithm is part of an adaptive interactive system that is acting in a Markov decision process (MDP), but the set of initial features provided may be a small subset of the features needed to satisfy the Markov assumption. We can view the coarse observation space as a state abstraction. A goal of our algorithm is to be able to augment the observation space in order to reach the minimal set of features sufficient to make the same optimal decisions as would be possible with the full (unknown) set of features (formally known as a optimal π^* -irrelevance abstraction [3]).

Our algorithm proceeds by proposing splits of existing observations that look identical under the current set of features. The key contribution of our work is to estimate the potential value of policies with new augmented observation space using old data without making the Markov assumption. Our algorithm for doing so is inspired by Upper Confidence Trees [2], a popular Monte Carlo Tree Search method, but adapted to focus on decision policy evaluation and, more importantly, does not require an MDP dynamics and reward so that it can be run using old data.

We evaluate the performance of our approach on several simulation domains where the true Markovian state space is known and demonstrate significant improvement over prior baselines. While a key next step is to try this out with a human designer in the loop, this is an encouraging step of the potential benefit of adaptively adjusting the feature representation.

2 PROBLEM SETUP

We consider human-in-the-loop reinforcement learning, where a human system designer can modify the observation space definition used by a reinforcement learning agent, such as when a designer can modify the observation space of a RL intelligent tutoring system interacting students. More precisely, we assume the RL agent is acting in an episodic Markov decision process $M = \langle S, \mathcal{A}, T, R, \gamma \rangle$, where S is a finite set of states (such as a student's current state of learning), $\mathcal{A}$ is a finite set of actions, and T is a dynamics model that specifies $p(s'|s, a)$ – the probability of transitioning to state s' after taking action a in state s (for example, the probability the student will not understand 1 digit addition, do a problem on addition,

and transition to a new state s' in which the student understands addition.) R specifies the reward $r(s, a)$ received by taking action a in state s : e.g., high reward when a student takes a test and passes it. γ is a discount factor that weighs immediate vs future reward.

We assume the state space S (such as the true internal state of the student) is only indirectly observable by the agent through sensors that provide the observation space $\mathcal{O}$ (e.g., the agent can observe if the student got a problem correct). There is a many-to-one deterministic mapping from states to observations, and therefore the observation space can be viewed as an aggregated state space. We denote the aggregated states s under observation o as S_o .

A decision policy π for the RL agent is a stochastic mapping from states to actions. The state-action value of a policy $Q^\pi(s, a)$ is the expected discounted sum of rewards the RL agent would obtain by taking action a from state s and then following the policy. In RL the dynamics and reward model are unknown.

The agent is provided with an initial observation space $\mathcal{O}_0$ that can be modified by the human system designer. The goal is for the interactive reinforcement learning agent to, together with the designer, find the smallest observation space that yields the maximal expected reward policy such that the resulting policy matches the performance of the best policy under the (unknown) MDP.

3 ALGORITHM

We present a novel human-in-the-loop RL algorithm, Value Driven Representation (VDR). VDR involves two key components. First, VDR performs optimistic reinforcement learning given the current observation space specification (e.g. it tries to optimize a decision policy for teaching a student given the current available features that distinguish student learning states) (Section 3.1). Second, VDR evaluates potential augmentations of the existing observation space that might enable a better decision policy and proposes to a human designer to split the observation that is evaluated to be most beneficial (e.g. for the human to provide another feature, like frustration, that can be used to refine an observation from "solved problem 1" to "solved problem 1 with frustration" and "solved problem 1 without frustration") (Section 3.2). Note in this initial work we only simulate input human experts and leave a human user study to later work.

Algorithm 1 VDR

```

1:  $\mathcal{O} \leftarrow \mathcal{O}_0, \mathcal{D} \leftarrow []$  // data
2: for each episode do
3:    $(\pi, V^\pi) \leftarrow$  optimistic Off-Policy Policy Optimization ( $\mathcal{D}$ )
4:   Execute  $\pi$  for 1 episode and collect trajectory  $\tau, \mathcal{D} \leftarrow \mathcal{D} \cup \tau$ 
5:   for all observations  $o_i \in \mathcal{O}$  do
6:     Propose augmentation of  $o_i \rightarrow (o_i^1, o_i^2)$  // section 3.2
7:      $\mathcal{D}' \leftarrow$  relabel  $o_i$  in the data  $\mathcal{D}$  with  $o_i^1, o_i^2$ 
8:      $(\pi_{o_i}^*, V^{\pi_{o_i}^*}) \leftarrow$  Off-Policy Policy Optimization ( $\mathcal{D}'$ ) //section 3.1
9:   end for
10:   $(score_{best}, o_{best}) \leftarrow \max_{o_i} D_{KL}(\pi_{o_i}^* || \pi)(V^{\pi_{o_i}^*} - V^\pi)$ 
11:  if  $score_{best} \geq c$  then
12:     $\mathcal{O} \leftarrow \mathcal{O} \cup \{o_{best}^1, o_{best}^2\} \setminus o_{best}$ 
13:  end if
14: end for
```

Algorithm 2 $\text{OPTE}(\pi, \text{graph}, N, \text{opt})$

Input: π : evaluation policy, graph : trajectory tree, N : number of simulations, opt : optimism flag

```

1: for  $i = 1 : N$  do
2:    $\text{node} \leftarrow \text{graph.initialNode}()$ 
3:    $o \leftarrow \text{node.observation}()$ ;  $R(i) \leftarrow []$ ;  $\text{done} \leftarrow \text{False}$ 
4:   while not  $\text{done}$  do
5:      $\text{node}, r, \text{done} \leftarrow \text{next}(\text{node}, \pi(o), \text{opt})$ 
6:      $R(i).append(r)$ ;  $o \leftarrow \text{node.observation}()$ 
7:   end while
8: end for
9: return  $\frac{1}{N} \sum_i \sum_t \gamma^t R_t(i)$ 

```

3.1 Off-Policy Policy Optimization

We first consider how an RL agent should act given the current experience in order to quickly learn a good decision policy for the current observation space. A key consideration here is that the current observation space $\mathcal{O}$ is generally not Markovian: therefore standard techniques that rely on the system description being a Markov process (like estimating a Markov dynamics and reward model and planning) can (and often will) fail [4, 5].

Instead, we propose an Off-Policy Tree Evaluation (OPTE) approach, similar to Monte Carlo Tree Search (MCTS) planning [2] and evaluation methods. The key distinction is that *we do not assume access to a domain model of the decision process* which standard MCTS methods rely on to perform simulations. OPTE first uses prior data to construct a tree representing observed trajectory sequences, storing counts of the number of times each node has been previously visited at each node. We call this a trajectory tree $\mathcal{T}$.

Similar to MCTS, OPTE uses the tree structure to simulate potential sequences of observations, actions and rewards (called roll-outs). A single rollout involves starting at the root node (o_{t_0}) of $\mathcal{T}$ and sampling transitions using a maximum likelihood model of transition probabilities and rewards given the data associated with this node. It is possible the observed outcomes have not included all feasible next observations. To handle this case, our algorithm maintains a pseudo-count of $C = 1$ over an additional next unseen observation transition, for each action in the tree. If this outcome is sampled, or if an action has not been tried from a particular node previously (e.g. at step d), we terminate the simulated roll-out and use a model free Monte Carlo estimate of the current observation-action pair Q -value $Q(o_d, a_d)$ at the leaf node, computed by averaging over all returns obtained after observing (o, a) tuple. Our complete Off Policy Tree Evaluation (OPTE) is shown in algorithm 2 and 3, which takes the average of N simulations.

$$V_O^\pi = \frac{1}{N} \sum_{i=1}^N \left(\sum_{t=1}^{d_i} \gamma^{t-1} r_t + \gamma^{d_i} Q(o_{d_i}, a_{d_i}) \right)$$

Additionally, in order to encourage strategic exploration, we propose a method that is inspired by optimism under uncertainty approaches. Precisely, we take the OPTE algorithm (Algorithm 2) described above and add in a reward bonus to the $Q(o, a)$ used at the tree leaves. Similar to upper confidence bound RL algorithms [9] we use $Q(o, a) + \frac{r_{\max}}{1-\gamma} \sqrt{\log(n(o))/n(o, a)}$, where $r_{\max}$ is the maximum reward and $n(o)$, $n(o, a)$ are visitation counts of an observation and an observation-action pair. This approach is a minor modification of OPTE, which can be computed by setting the input opt to *True* in algorithm 3. Off-Policy Tree Optimization (OPTO) can be done using

Algorithm 3 $\text{next}(\text{node}, a, \text{opt})$

Input: node : node in trajectory graph, a : action, opt : optimism flag

```

1:  $C \leftarrow 1$ 
2:  $\text{counts} \leftarrow \text{node.getTransitionCounts}(a)$ 
3:  $\text{counts}.append(C)$  // Unobserved outcome
4:  $\text{model} \sim \text{MLE}(\text{counts})$ 
5:  $\text{nextNode} \sim \text{model}(a)$ 
6: if  $\text{nextNode}$  is -1 then
7:   // Sample a new potential outcome
8:    $o \leftarrow \text{node.OBSERVATION}()$ ;  $\text{done} \leftarrow \text{True}$ 
9:   if  $\text{opt} == \text{True}$  then
10:     $r \leftarrow \text{graph}.Q[o, a] + \frac{r_{\max}}{1-\gamma} \sqrt{\frac{\log(n[o])}{n[o, a]}}$ 
11:   else
12:     $r \leftarrow \text{graph}.Q[o, a]$ 
13:   end if
14: else
15:    $r, \text{done} \leftarrow \text{nextNode.Info}()$ 
16: end if
17: return  $(\text{nextNode}, r, \text{done})$ 

```

any policy optimization method combined with using simulated roll outs on the trajectory tree $\mathcal{T}$.

3.2 Observation Augmentation

The initial input observation space may be insufficient to achieve a high performance policy: for example, it may be crucial to change the policy depending on whether a student is frustrated after completing problem 1 correctly, yet initially this distinction may be lacking in the observation space. We now propose how the RL system can itself try to identify which observation refinements might yield an improved performance (policy value) if a human designer could provide a feature that distinguished between observations that are currently aliased.

To do this, at every episode the algorithm creates and scores $\{O_i\}$ new potential observation spaces. Each observation space O_i is derived from taking observation o_i in the current O_{current} observation space, splitting it into two new observations o_i^1 and o_i^2 , and adding these two new observations to all of the other non-split observations $o_{k \neq i} \in O_{\text{current}}$ ($\{o_i^1, o_i^2, \forall_{k \neq i} : o_k \in O_{\text{current}}\}$).

Splitting a particular observation o_i into two is performed by executing Expectation Maximization (EM) [8] on the existing collected trajectories to hypothesize 2 potential latent observations with different dynamics and/or rewards models. Given the EM learned parameters for a observation split (o_i into observations o_i^1 and o_i^2), the Viterbi algorithm [1] can be used to relabel prior trajectories, turning all instances of observation o_i into o_i^1 or o_i^2 . Using the relabeled trajectories $\mathcal{D}'$, we can build a new trajectory tree $\mathcal{T}'$ and perform off-policy policy optimization using OPTO to estimate the value of the best policy $V^{\pi_{o_i^*}}$ for the modified space.

The objective is to present an augmentation to the system designer if 1) splitting an observation yields an optimal policy with a higher value than the existing best policy for the observation space, and 2) if the best policy for the augmented observation (evaluated using OPTO) is sufficiently different than the previous best policy, measured by the KL-divergence of two policies. Precisely, define $\text{score}(i) = D_{KL}(\pi_{old}^* || \pi_i)(V^{\pi_i} - V^{\pi_{old}})$ where D_{KL}

Figure 2: Cheese Maze, (a) Environment, agent has 4 actions: move up, down, left, right and receives a reward of -10 if it hits the wall, -1 for moving to an empty state and 10 for getting to the cheese state which also teleports the agent to the initial state (bottom left). Numbers in each cell indicate the initial observation space, and split of observation 1 into (1a, 1b). (b) Comparison of our algorithm to MC-AIXI-CTW

is the KL-divergence, π_i is the optimal policy after splitting observation o_i . Our agent proposes an observation augmentation $i = \operatorname{argmax}_i \operatorname{score}(i)$ to the system designer if $\max_i \operatorname{score}(i) \geq c$ exceeds a threshold c .

4 SIMULATION RESULTS

This human-in-the-loop RL system is designed ultimately to be used for helping humans and RL agents best work together to achieve a good representation that can be used to quickly identify good policies. However, as an initial proof of concept, we first conduct simulated experiments where a simulated human designer will agree a proposed observation refinement is beneficial if $c = 0.25$ and where the revised associated observation is split as follows. Assuming MDP states $s \in \mathcal{S}_o$ was clustered under observation o , we assign state s to observation o^1 (or o^2) if more than 50 percent of the time s was assigned to o^1 (or o^2) by the Viterbi algorithm.

We simulate our VDR algorithm in two existing RL tasks. One is a navigation task CheeseMaze [11] inspired by a robot needing to learn how to reach a destination: the robot requires particular features in order to be able to learn and represent the optimal policy. The second is mountain car [10], where a car on a hill must reach the goal position up the right side of the hill (figure 3a).

The goal is to see if our method will propose the necessary augmentation to the simulated human to learn the optimal policy. Additionally, we are interested to evaluate how fast our algorithm can learn starting from a coarse observation space, when representing the optimal policy does not require full Markovian state space.

4.1 Cheese Maze

Cheese Maze was used as a benchmark environment in [11]; for details refer to figure 2a. We set the maximum length of each episode to 20, and consider an augmentation every 5 episodes. We compare to MC-AIXI-CTW that outperformed other history based and feature RL method including UTree [6, 11].

As shown in Figure 2b VDR outperforms MC-AIXI-CTW and finds the optimal policy in fewer number of episodes. VDR finds the optimal policy by splitting observation 1 (see figure 2a) into two (1a, 1b) that requires different action for representing the optimal policy in center and left side of the maze.

Figure 3: Mountain Car, (a) Environment. (b) Comparison of VDR with Q-learning, tile coding and DQN. (c) Observation space augmentation: each color represents distinct observation and the observation space is shown superimposed on the underlying 8x8 grid.

4.2 Mountain Car

We considered mountain car where the agent always starts at the same location and velocity $x_0 = -0.5, v_0 = 0.03$. In all simulations the initial starting position and velocity is fixed for that entire process. We set the maximum episode length to 500 and consider an observation augmentation every 20 episodes. We treat the underlying true state space as a discrete 8x8 grid, though the true space is best modeled continuously. We compared our algorithm with Q-learning with ϵ -greedy exploration on 20x20 grid (location and velocity). Additionally we compared to DQN [7] (with two hidden layers of size 64) and tile coding with 2 tilings.

Figure 3b shows the result for $v_0 = 0.03$, where VDR can learn the optimal policy in 100 episodes with only 3 augmentations: this is enough to represent the optimal policy. Our experiments with other initial velocity ($v_0 = 0.05$) showed the same results. Figure 3c shows the observation augmentation. Starting with only two observation shown in figure 3c superimposed on the underlying 8x8 grid, VDR find the optimal policy after three splits.

5 FUTURE WORK AND CONCLUSION

Adaptive interactive reinforcement learning systems often benefit from a human-in-the-loop system designer that can modify the state or action space of the RL system to improve performance. Given the potential set of modifications, guidance from the RL system to the human designer could be helpful. In this short paper we propose a way for a RL system to proactively propose augmentations to its observation representation that may enable improved performance. Our simulations suggest the potential benefit of this approach in small domains. Exploring the scalability of this approach and testing it with real humans in the loop are clear interesting next steps.

ACKNOWLEDGEMENT

The research reported here was supported by NSF BIGDATA award and ONR Young Investigator award.

REFERENCES

- [1] G David Forney. 1973. The viterbi algorithm. *Proc. IEEE* 61, 3 (1973), 268–278.
- [2] Levente Kocsis and Csaba Szepesvári. 2006. Bandit based monte-carlo planning. In *European conference on machine learning*. Springer, 282–293.
- [3] Lihong Li, Thomas J Walsh, and Michael L Littman. 2006. Towards a unified theory of state abstraction for MDPs.. In *ISAIM*.
- [4] Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. 2014. Offline policy evaluation across representations with applications to educational games. In *International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1077–1084.
- [5] Andrew McCallum. 1995. *Reinforcement Learning with Selective Perception and Hidden State*. Ph.D. Dissertation.
- [6] Andrew Kachites McCallum et al. 1996. Learning to use selective attention and short-term memory in sequential tasks. In *From animals to animats 4: proceedings of the fourth international conference on simulation of adaptive behavior*, Vol. 4. MIT Press, 315.
- [7] 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.
- [8] Todd K Moon. 1996. The expectation-maximization algorithm. *IEEE Signal processing magazine* 13, 6 (1996), 47–60.
- [9] Ian Osband and Benjamin Van Roy. 2014. Near-optimal reinforcement learning in factored MDPs. In *Advances in Neural Information Processing Systems*. 604–612.
- [10] Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*. Vol. 1. MIT press Cambridge.
- [11] Joel Veness, Kee Siong Ng, Marcus Hutter, William Uther, and David Silver. 2011. A monte-carlo aixo approximation. *Journal of Artificial Intelligence Research* 40, 1 (2011), 95–142.