

A Q-learning Approach for Adherence-Aware Recommendations

Ioannis Faros¹, *Student Member, IEEE*, Aditya Dave², *Member, IEEE*,
Andreas A. Malikopoulos^{1,2}, *Senior Member, IEEE*

Abstract—In many real-world scenarios involving high-stakes and safety implications, a human decision-maker (HDM) may receive recommendations from an artificial intelligence while holding the ultimate responsibility of making decisions. In this letter, we develop an “adherence-aware Q-learning” algorithm to address this problem. The algorithm learns the “adherence level” that captures the frequency with which an HDM follows the recommended actions and derives the best recommendation policy in real time. We prove the convergence of the proposed Q-learning algorithm to the optimal value and evaluate its performance across various scenarios.

Index Terms—Q-learning, Markov Decision Processes, Recommender systems, Reinforcement learning.

I. INTRODUCTION

Decisions driven by artificial intelligence have recently found applications in complex cyber-physical systems [1] such as transportation [2], [3], finance [4], and healthcare [5]. However, decisions involving high-stakes or safety-critical applications [6] are often ultimately taken by human decision-makers (HDMs) under advice from an artificial intelligence algorithm. Since this is at the discretion of the HDM, the algorithm’s recommendations may not be followed at every instance of time [7]. The phenomenon of unexpected decisions influencing the performance of such “expert-in-the-loop” systems has garnered increasing interest in recent years, primarily in the fields of operations [8], human trust on machines [9], finance [4], and healthcare services [5].

Many research efforts have focused on the factors influencing the adherence of HDMs to recommendations. In [10], it was established that HDMs usually prefer following recommendations that match their comfort and expertise and may ignore recommendations that contradict their opinions. The authors in [9] provide evidence for the hypothesis that HDMs trust their knowledge rather than an algorithm. A similar phenomenon for partial adherence is also observed when humans recommend actions to other humans, e.g., medical advice [11]. Conversely, an algorithm may not be able to account for real-life limitations faced by humans when implementing an action [12]. Thus, an HDM may

This research was supported by NSF under Grants CNS-2149520 and CMMI-2219761.

¹Systems Engineering, Cornell University, Ithaca, NY 14850, USA.

²School of Civil and Environmental Engineering, Cornell University, Ithaca, NY 14850, USA. email: {if74, adidave, amaliko}@cornell.edu

trust their judgment to navigate such situations and disregard algorithmic recommendations [13]. As a consequence of these factors, algorithmic recommendations to HDMs can perform significantly worse than anticipated [11].

Two main approaches have been proposed in the literature to address performance degradation: (1) *Increasing adoption*: Many research efforts have focused on increasing the adoption of recommendations among HDMs [14]. In [15], reinforcement learning techniques were utilized to increase HDM adherence, assuming they are likelier to follow recommendations close to their baseline strategy. In [16], a Lazy Markov decision processes (MDPs) formulation was proposed to improve the alignment of recommendations with the baseline strategy. (2) *Incorporating adherence awareness*: A more recent approach is to improve recommendations by considering partial adherence within the algorithm’s problem. In [17], an adherence-aware MDP was analyzed, and a value iteration algorithm was proposed to improve the performance of recommendations. However, this result relies upon prior knowledge of the HDM’s adherence level and system dynamics. To bring this framework closer to real-world implementation, we require further insights into adherence-aware recommendations that can be obtained with incomplete prior knowledge of dynamics and adherence levels.

In this letter, we address this challenge by incorporating partial adherence into reinforcement learning [18], [19] by proposing an “adherence-aware Q-learning algorithm.” Specifically, we consider an MDP comprising an HDM that implements actions to influence the evolution of an unknown environment. The HDM’s strategy for generating actions combines an algorithmic recommendation and a baseline strategy. The adherence level of the HDM is not known to the algorithm a priori. The algorithm and the HDM share the same objective, and thus, our formulation aims to compute in real time the best-recommended actions to maximize an expected total discounted reward.

Our main contributions in this letter are the (1) introduction of an adherence-aware Q-learning algorithm (Algorithm (7)) to compute an optimal control law, (2) convergence for the algorithm to the optimal Q-function (Theorem 1), and (3) establishment of the advantages of our approach against baseline policies and classical Q-learning using numerical examples (Section IV).

The remainder of the letter proceeds as follows. In Section II, we present our formulation, the definitions of the adherence-aware Q-learning function, and the updating rule of

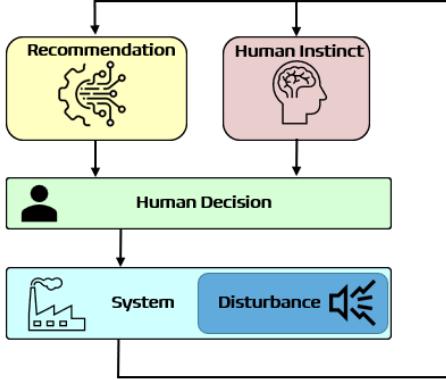


Fig. 1: Process of the actual law implementation.

the adherence level. In Section III, we prove the convergence of the proposed algorithm. In Section IV, we demonstrate our result in an inventory control problem, and in Section V, we draw concluding remarks.

II. MODELING FRAMEWORK

A. Problem Formulation

We consider a system comprising an artificial intelligence or algorithm recommending actions to an HDM. In turn, the HDM implements actions to influence the evolution of a dynamic environment. At each instance of time, the HDM can select to either follow the recommendation provided by the algorithm or select an action using a baseline law. The evolution of the state of the environment is modeled as an infinite horizon discounted MDP $\mathcal{S} = (\mathcal{X}, \mathcal{U}, P, R, \lambda)$, where \mathcal{X} is a finite set of states and \mathcal{U} is a finite set of actions. At any time $t \in \mathbb{N} = \{0, 1, \dots\}$, the state of the system is denoted by the random variable $X_t \in \mathcal{X}$ and the action input to the MDP is the random variable $U_t \in \mathcal{U}$. The function $P : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \rightarrow [0, 1]$ yields the transition probability for all t as $P(x_{t+1} | x_t, u_t) = p(X_{t+1} = x_{t+1} | X_t = x_t, U_t = u_t)$ from any realized state $x_t \in \mathcal{X}$ and realized action $u_t \in \mathcal{U}$ to the next realized state $x_{t+1} \in \mathcal{X}$. The function $R : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ yields the reward $R(X_t, U_t)$ for all $t \in \mathbb{N}$, and $\lambda \in (0, 1]$ is a discount factor applied to future rewards when measuring performance.

In the MDP \mathcal{S} , the actions are implemented by an HDM with access to a *baseline* law $g^b : \mathcal{X} \rightarrow \mathcal{U}$ that belongs to a set G of stationary Markovian laws. The HDM also receives an action recommendation at each instance of time from an algorithm using a *recommendation* law $g^r : \mathcal{X} \rightarrow \mathcal{U}$ from the set of feasible laws G . Subsequently, we consider that the HDM implements the recommended action with a probability $\theta \in [0, 1]$ and the baseline action with a probability $1 - \theta$. Thus, the HDM follows a mixed *actual* control law given by a convex combination of the recommended and baseline laws as follows:

$$g^a = \theta \cdot g^r + (1 - \theta) \cdot g^b. \quad (1)$$

The process of an HDM taking an action using the baseline and the recommendation law is illustrated in Fig. 1. The probability θ is a representation of the *adherence* level of

the HDM to a recommended action, i.e., it captures the fact that an HDM does not systematically implement the latter. The objective of an HDM is to maximize the expected total discounted reward

$$J(g^a; x) = \lim_{T \rightarrow \infty} \mathbb{E}^{g^a} \left[\sum_{t=1}^T \lambda^{t+1} \cdot R(X_t, U_t) \mid X_1 = x \right], \quad (2)$$

where $\mathbb{E}^{g^a}[\cdot]$ is the expectation with respect to the distributions of all random variables generated using the actual control law.

Problem 1. The objective is to derive the optimal control law g^{a*} such that $g^{a*} = \theta g^r + (1 - \theta) g^b$ satisfies

$$J(g^{a*}) \geq J(g^a), \quad \forall g^a \in G^a(g^b), \quad (3)$$

where $G^a(g^b)$ is the set of all possible control laws for a given baseline law g^b and θ .

Remark 1. The term U_t refers to the action implemented by the HDM at any $t \in \mathbb{N}$. Additionally, we denote the recommended action as $U_t^r = g^r(X_t)$ and the baseline action as $U_t^b = g^b(X_t)$ for all t .

In our modeling framework, we impose the following assumptions :

Assumption 1. The stationary adherence level θ and transition matrix P may be unknown. However, the baseline law g^b is known a priori.

The baseline law g^b can be determined by observing the system over a finite number of instances since the sets of states and actions are finite. Thus, we consider the law to be known in our problem formulation.

Assumption 2. The baseline law g^b is deterministic.

The baseline law must be deterministic to recognize when the HDM implements the baseline action as opposed to when they implement the recommended action. This assumption is satisfied by the optimal solutions to regular MDPs. An extension of our results for mixed baseline laws is a direction of future research.

B. Preliminary Results

In Problem 1, an optimal value function is given by the fixed point of the following Bellman-like equation for all possible states $x \in \mathcal{X}$ [17]:

$$V(x) = \max_{u^r} \theta \cdot \sum_{x' \in \mathcal{X}} \left(R(x, u^r) + P(x' | x, u^r) \cdot V(x') \right) + (1 - \theta) \cdot \sum_{x' \in \mathcal{X}} \left(R(x, g^b(x)) + P(x' | x, g^b(x)) \cdot V(x') \right), \quad (4)$$

where x' denotes a realization of the next state, u^r the recommended action, $u^b = g^b(x)$ the base action for a given state $x \in \mathcal{X}$. Here, we use a time-invariant notation because (4) holds for all $t \in \mathbb{N}$ and all possible states in \mathcal{X} . When the system dynamics and adherence level θ are known, we can use value iteration with (4) to derive an optimal recommendation law.

In our formulation, we consider that these functions and quantities may be unknown. Next, we provide a Q-learning algorithm to compute the optimal recommendation law in Problem 1.

III. Q-LEARNING APPROACH

In this section, we propose an *adherence-aware Q-learning* algorithm that can learn a Q-function for Problem 1 and subsequently use this to derive an optimal recommendation law. We also prove the convergence of the proposed algorithm to the optimal Q-function. We begin by constructing an unbiased estimate for the unknown θ using a point estimator of θ that considers the action implemented by the HDM. The tuple $(\theta, 1 - \theta)$ can be interpreted as the probability distribution of a sequence of identical and independently distributed random variables $\{Y_t\}_{t \in \mathbb{N}}$, where each $Y_t \in \{0, 1\}$ indicates whether the HDM implements the recommended action at time t , i.e., $Y_t = 1$ when $U_t = g^r(X_t)$ and $Y_t = 0$ when $U_t = g^b(X_t)$. To construct this estimator for any sample size $n \in \mathbb{N}$, we first collect a sample s^t comprising observations of the random variables Y_0, \dots, Y_n . Then, we utilize the point estimator to approximate θ as $\theta^t = s^t/n$. To carry out this procedure online during Q-learning, we can write the update rule for the $n+1$ -th estimate θ_{t+1} of θ from the n -th estimate θ^t as

$$\theta_{t+1} = \frac{\theta_t \cdot n + \mathbb{1}[Y_{t+1} = 1]}{n+1}, \quad (5)$$

where $\mathbb{1}[\cdot]$ is the indicator function. Using standard arguments, we can show that this update rule will converge to the true θ value [20] in (1), i.e.,

$$\mathbb{E}[\theta_{t+1}] = \theta. \quad (6)$$

Using the update rule for θ , we propose our adherence-aware Q-learning algorithm to learn the value function in (4) as

$$\begin{aligned} Q(x, u) \leftarrow Q(x, u) + \alpha(x, u) \cdot & \left\{ \theta_{t+1} \cdot \left[r(x, u) \right. \right. \\ & + \lambda \cdot \max_{g^r} Q(x', g^r(x')) \left. \right] + (1 - \theta_{t+1}) \cdot \left[r(x, u) \right. \\ & \left. \left. + \lambda \cdot Q(x', g^b(x')) \right] - Q(x, u) \right\}, \end{aligned} \quad (7)$$

where $x' \in \mathcal{X}$ is the next state, $x \in \mathcal{X}$ is the current state, $u \in \mathcal{U}$ is the current action, and $\alpha(x, u)$ is the learning rate. We use time-invariant notation as this algorithm holds for all $t \in \mathbb{N}$ and all states in \mathcal{X} and actions in \mathcal{U} . The term $Q(x, u)$ in the RHS is the current value for a given state x and action u . The term $\max_{g^r} Q(x', g^r(x'))$ is the estimate of optimal future value, and $Q(x', g^b(x'))$ is the future value given the baseline law. Next, we show the adherence-aware Q-learning algorithm in procedural form.

A. Proof of convergence

In this subsection, we prove the convergence of the adherence-aware Q-learning algorithm to the optimal value function. Before we prove convergence, we first define an adherence-aware operator for our Bellman-like recursion, and we prove that this is a contraction mapping.

Algorithm 1 Adherence-aware Q-learning algorithm

Algorithm parameters: step size $\alpha \in (0, 1]$, small $\epsilon > 0$, $\lambda \in [0, 1]$. Consider a baseline law $g^r(x) = u^r$.
Initialize: $Q(x, u)$, for all $x \in \mathcal{X}, u \in \mathcal{U}$, arbitrarily.

```

for each episode do
  Initialize  $x$ 
  for each iteration do
    Choose  $u$  using a law derived from  $Q$  (e.g.,  $\epsilon$ -greedy)
    Take action  $u$ , compute  $r, x'$ 
    Update  $Q$  using Equation (7)
     $x \leftarrow x'$ 
  end for
end for

```

Definition 1. The adherence-aware operator $\mathcal{J} : [\mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}] \rightarrow [\mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}]$ for any bounded $q : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is

$$\begin{aligned} (\mathcal{J}q)(x_t, u_t) = \theta_{t+1} \max_{u_t^r} \sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^r) \cdot & \left[R(x_t, u_t^r) \right. \\ & + \lambda \cdot q(x_{t+1}, u_t^r) \left. \right] + (1 - \theta_{t+1}) \cdot \sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^b) \\ & \cdot \left[R(x_t, u_t^b) + \lambda q(x_{t+1}, u_t^b) \right], \quad \forall t \in \mathbb{N}, \end{aligned} \quad (8)$$

for all possible realization $x_t \in \mathcal{X}$ and $u_t \in \mathcal{U}$, where $u^r = g^r(x_t)$ is the recommended action and $u^b = g^b(x_t)$ is the baseline action.

Lemma 1. The operator \mathcal{J} is a contraction mapping.

Proof. We prove the result using the definition of a contraction mapping, i.e., we prove that \mathcal{J} satisfies

$$\|\mathcal{J}q_1 - \mathcal{J}q_2\|_\infty \leq \lambda \|q_1 - q_2\|_\infty, \quad (9)$$

for any bounded $q : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ and for any possible realization $x_t \in \mathcal{X}$ and $u_t \in \mathcal{U}$.

Substituting (8), using the definition of the ∞ -norm, in (9), we obtain:

$$\begin{aligned} \|\mathcal{J}q_1 - \mathcal{J}q_2\|_\infty = \max_{x, u^r} & \left| \theta_{t+1} \cdot \lambda \cdot \sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^r) \right. \\ & \cdot \left[\max_{u^r} q_1(x_{t+1}, u_t^r) - \max_{u^r} q_2(x_{t+1}, u_t^r) \right] + (1 - \theta_{t+1}) \cdot \\ & \lambda \cdot \sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^b) \cdot \left[q_1(x_{t+1}, u_t^b) - q_2(x_{t+1}, u_t^b) \right] \left. \right|. \end{aligned} \quad (10)$$

From the triangle inequality, we obtain:

$$\begin{aligned} \|\mathcal{J}q_1 - \mathcal{J}q_2\|_\infty \leq \max_{x, u^r} & \theta_{t+1} \cdot \lambda \cdot \sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^r) \\ & \cdot \left| \max_{s_{t+1}, U_t^r} q_1(s_{t+1}, u_t^r) - \max_{s_{t+1}, U_t^r} q_2(s_{t+1}, u_t^r) \right| + (1 - \theta_{t+1}) \cdot \\ & \lambda \cdot \sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^b) \cdot \left| q_1(x_{t+1}, u_t^b) - q_2(x_{t+1}, u_t^b) \right|. \end{aligned} \quad (11)$$

Since $\sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^r) = 1$ and $\sum_{x_{t+1} \in \mathcal{X}} P(x_{t+1} | x_t, u_t^b) = 1$ in the RHS of (11), $\|\mathcal{J}q_1 - \mathcal{J}q_2\|_\infty \leq \max_{x, u^r} \theta_{t+1} \cdot \lambda \cdot |q_1 - q_2| + (1 - \theta_{t+1}) \cdot \lambda \cdot |q_1 - q_2|$.

We complete the proof using the definition for the ∞ -norm to state that $\|\mathcal{J}q_1 - \mathcal{J}q_2\|_\infty \leq \theta_{t+1} \cdot \lambda \cdot (\|q_1 - q_2\|_\infty) + (1 - \theta_{t+1}) \cdot \lambda \cdot (\|q_1 - q_2\|_\infty) = \lambda \cdot \|q_1 - q_2\|_\infty$. \square

Using the Banach fixed point theorem and the result of Lemma 1, the equation $Q = \mathcal{J}Q$ admits a unique solution $Q^\infty = \mathcal{J}Q^\infty$. Furthermore, starting at $Q^0(x_t, u_t) = 0$, the fixed point iteration around \mathcal{J} generates a sequence of functions $Q^{k+1}(x_t, u_t) = \mathcal{J}Q^k(x_t, u_t) = \mathcal{J}^k Q^0(x_t, u_t)$ for all $k = 1, \dots$, such that $\lim_{k \rightarrow \infty} \mathcal{J}^k Q^0 = Q^\infty$. Next, we prove that the operator is monotone and then use this to establish that the fixed point Q^∞ is the optimal value.

Lemma 2. *Consider the maximum value over all control laws in Problem 1 given by*

$$Q^* = \max_{g^a} \mathbb{E}^{g^a} \left[\sum_{t=1}^{\infty} \lambda^{t-1} R(X_t, U_t) \mid X_1 = x \right]. \quad (12)$$

Then, the fixed point solution $Q^\infty = \mathcal{J}Q^\infty$ also satisfies $Q^\infty = Q^*$.

Proof. Suppose that $R(x_t, u_t)$ is bounded by $C \in \mathbb{R}$ for all possible states and actions, and consider a finite truncation as $Q_T^* = \max_{g^a} \mathbb{E}^{g^a} \left[\sum_{t=1}^T \lambda^{t-1} R(X_t, U_t) \mid X_1 = x \right]$. Case (1), by construction, the truncate Q_T^* is sub-optimal w.r.t. Q^* , i.e., $Q^* \geq Q_T^*$. However, Q_T^* satisfies $Q_T^* = \mathcal{J}^k Q^0$, and hence as $k \rightarrow \infty$ we can write $Q^* \geq Q_T^* = \mathcal{J}^k Q^0 = Q^\infty$. In Case (2), we use the assumption that the reward is bounded, and hence we can write $Q^* \leq Q_T^* + \sum_{t=T}^{\infty} \lambda^T C$ for all T . Taking the $\lim_{T \rightarrow \infty}$, we get $Q^* \leq Q_T^* = Q^\infty$. From Cases (1) and (2), we conclude that $Q^\infty = Q^*$. \square

Next, we prove the convergence of Algorithm 1 in the tabular setting with finite-valued random variables.

Theorem 1. *Given a finite MDP $(\mathcal{X}, \mathcal{U}, P, R, \lambda)$, the adherence-aware Q-learning algorithm, defined by the update rule (7) converges with probability 1 to the optimal Q^* in (12) under the following conditions for all $(x, u) \in \mathcal{X} \times \mathcal{U}$:*

$$1) \quad \sum_t \alpha_t(x, u) = \infty, \quad (13)$$

$$2) \quad \sum_t \alpha_t(x, u)^2 < \infty, \quad (14)$$

$$3) \quad R(x, u^r) \text{ and } R(x, u^b) \text{ are bounded.} \quad (15)$$

Proof. For any $t \in \mathbb{N}$, we write the $t+1$ -th iteration of the Q -update rule (7) as:

$$\begin{aligned} Q_{t+1}(x, u) &= (1 - \alpha_t(x, u))Q_t(x, u) \\ &+ \alpha_t(x, u) \left\{ \theta_{t+1} \left[\max_{u^r} (R(x, u^r) + \lambda \cdot Q_t(x', u^r)) \right] \right. \\ &\quad \left. + (1 - \theta_{t+1}) [R(x, u^b) + \lambda \cdot Q_t(x', u^b)] \right\}, \end{aligned} \quad (16)$$

where recall that $u^b = g^b(x)$. By subtracting $Q^*(x, u)$ from

both sides, we get:

$$\begin{aligned} \Delta_{t+1}(x, u) &= (1 - \alpha_t(x, u)) \cdot \Delta_t(x, u) + \alpha_t(x, u) \\ &\cdot \left\{ \theta_{t+1} \cdot \left[\max_{u^r} (R(x, u^r) + \lambda \cdot Q_t(x', u^r)) \right] + (1 - \theta_{t+1}) \right. \\ &\quad \left. \cdot [R(x, u^b) + \lambda \cdot Q_t(x', u^b)] - Q^*(x, u) \right\}, \end{aligned} \quad (17)$$

where $\Delta_t(x, u) = Q_t(x, u) - Q^*(x, u)$. Let

$$\begin{aligned} F_t(x, u, x') &= \theta_{t+1} \cdot \left[\max_{u^r} (R(x, u^r) + \lambda \cdot Q_t(x', u^r)) \right] \\ &+ (1 - \theta_{t+1}) \cdot [R(x, u^b) + \lambda \cdot Q_t(x', u^b)] - Q^*(x, u). \end{aligned} \quad (18)$$

Next, we prove that the expectation and the variance of $F_t(x, u, x')$ are bounded above. By taking the expectation in both sides in (18) given the history $\mathcal{F}_t = \{\Delta_t, \Delta_{t-1}, \dots, F_{t-1}, \dots\}$, we obtain

$$\begin{aligned} \mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] &= \mathbb{E} \left\{ \theta_{t+1} \cdot \left[\max_{u^r} (R(X, U^r) \right. \right. \\ &+ \lambda \cdot Q_t(X', U^r)) \mid \mathcal{F}_t, x, u \left. \right] \left. \right\} + \mathbb{E} \left\{ (1 - \theta_{t+1}) \cdot \left[R(X, U^b) \right. \right. \\ &+ \lambda \cdot Q_t(X', U^b) \mid \mathcal{F}_t, x, u \left. \right] \left. \right\} - Q^*(x, u). \end{aligned} \quad (19)$$

From Lemma 1, the update rule for θ_{t+1} is independent of state and action so that we can write

$$\begin{aligned} \mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] &= \mathbb{E}[\theta_{t+1} \mid \mathcal{F}_t] \\ &\cdot \mathbb{E} \left[\max_{u^r} (R(X, U^r) + \lambda \cdot Q^n(X', U^r)) \mid \mathcal{F}_t, x, u \right] \\ &+ \mathbb{E}[(1 - \theta_{t+1}) \mid \mathcal{F}_t] \cdot \mathbb{E} \left[R(X, U^b) \right. \\ &\quad \left. + \lambda \cdot Q_t(X', U^b) \mid \mathcal{F}_t, x, u \right] - Q^*(x, u). \end{aligned} \quad (20)$$

From (6), the last equation becomes

$$\begin{aligned} \mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] &= \theta \cdot \mathbb{E} \left[\max_{u^r} (R(X, U^r) \right. \\ &+ \lambda \cdot Q_t(X', U^r)) \mid \mathcal{F}_t, x, u \left. \right] + (1 - \theta) \cdot \mathbb{E} \left[R(X, U^b) \right. \\ &\quad \left. + \lambda \cdot Q_t(X', U^b) \mid \mathcal{F}_t, x, u \right] - Q^*(x, u). \end{aligned} \quad (21)$$

By expanding the expectations we have

$$\begin{aligned} \mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] &= \theta \cdot \max_{u^r} \sum_{x' \in \mathcal{X}} P(x' \mid x, u^r) \cdot \\ &\left[(R(x, u^r) + \lambda Q_t(x', u^r)) \right] + (1 - \theta) \cdot \sum_{x' \in \mathcal{X}} P(x' \mid x, u^b) \cdot \\ &\left[R(x, u^b) + \lambda \cdot Q_t(x', u^b) \right] - Q^*(x, u). \end{aligned} \quad (22)$$

The first two terms in the RHS of the last equation form the contraction operator \mathcal{J} , hence from Lemma 2,

$$\mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] = (\mathcal{J}Q)(xu) - (\mathcal{J}Q^*)(x, u). \quad (23)$$

To conclude the result, we use the ∞ -norm on both sides of (23) and the result of Lemma 1, and we obtain

$$\begin{aligned} \|\mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t]\|_\infty &= \|(\mathcal{J}Q)(x, u) - (\mathcal{J}Q^*)(x, u)\|_\infty \\ &\leq \lambda \cdot \|Q_t - Q^*\|_\infty = \lambda \cdot \|\Delta_t\|_\infty. \end{aligned} \quad (24)$$

Next, we show that the variance is also bounded above. By using the definition of the variance of $F_t(X, U, X')$ we write:

$$\begin{aligned} \text{var}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] &= \mathbb{E}\left\{\left(\theta_{t+1} \cdot \left(\max_{u^r} (R(X, U^r) + \lambda \cdot Q_t(X', U^r))\right) + (1 - \theta_{t+1}) \cdot (R(X, U^b) + \lambda \cdot Q_t(X, U^b))\right) - Q^*(X, U) \right. \\ &\quad \left. - \left[\underbrace{(\mathcal{J}Q)(X, U) - (\mathcal{J}Q^*)(X, U)}_{\mathbb{E}[F_t(X, U, X') \mid \mathcal{F}_t, x, u]}\right]^2 \mid \mathcal{F}_t, x, u\right\}. \end{aligned} \quad (25)$$

From Lemma 2, we can simplify the last term in the RHD of (25) as

$$\begin{aligned} \text{var}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] &= \mathbb{E}\left\{\left(\theta_{t+1} \cdot \left(\max_{u^r} (R(X, U^r) + \lambda \cdot Q_t(X', U^r))\right) + (1 - \theta_{t+1}) \cdot (R(X, U^b) + \lambda \cdot Q_t(X', U^b))\right) - (\mathcal{J}Q)(X, U)\right\}^2 \mid \mathcal{F}_t, x, u, \end{aligned} \quad (26)$$

which is the definition of the variance with variable $\theta_{t+1} \cdot (\max_{u^r} (R(X, U^r) + \lambda \cdot Q_t(X', U^r))) + (1 - \theta_{t+1}) \cdot (R(X, U) + \lambda \cdot Q_t(X', U^b))$. By using the Assumption (3), i.e., the reward. Hence its variance is bounded, i.e.,

$$\text{var}[F_t(X, U, X') \mid \mathcal{F}_t, x, u] \leq C \cdot (1 + \|\Delta_t\|_\infty)^2, \quad (27)$$

and the proof is complete. \square

IV. NUMERICAL EXAMPLES

A. Inventory Example

We consider an HDM in **inventory control** as a shop owner, where the shop provides multiple items (K) for sale. The goal of the HDM is to maximize its cumulative revenue at each time step by deciding the number of items that need to be ordered based on the stochastic demand of the future. Let x_k be the state of commodity k , where $x_k \in \{0, 1, \dots, 100\}$ and $k = 1, 2, \dots, K$. The action $u_k \in \{0, 1, \dots, 100\}$ corresponds to the amount of the order, and $d_k \in \{0, 1, \dots, 100\}$ corresponds to the amount of the demand that is stochastic and follows the uniform distribution with $B(0, 100)$. We assume that the baseline law of the HDM is according to (s, S) method [21], where s represents the threshold of which the HDM should order the amount of S . The state evolution for the inventory is according to $x_{k+1} = \max(0, x_k + u_k - D_k)$, and the reward is defined as the incomes of sales minus the holding cost H , and the ordering cost C . The holding cost is given by $H(x_k, u_k) = h \cdot \max(0, x_k - u_k)$, where $h > 0$ while the ordering cost is given by $C(u_k) = c \cdot u_k$, where $c > 0$.

B. Machine replacement

We consider the **Machine replacement** MDP problem with ten states. The set of states is $\{1, 2, 3, 4, 5, 6, 7, 8, S, L\}$ and the set of actions is $\{\text{repair}, \text{wait}\}$. Each state represents the condition of the machine, where in-state 8 the machine is broken, and states S, L represent the short and the long repair, respectively. To model this problem, we adopt the same rewards and transitions as in [17]. In particular, we show the

transition probabilities for both actions in Figure 2. Also, we define a reward of 18 in state S , 20 in state L , and 0 in state 8. All other states have a reward of 20. We assume the baseline law is always wait when the machine is not broken or in the long repair L . In any other case, the action is repair.

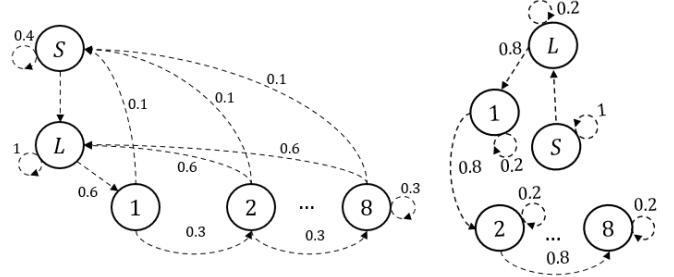


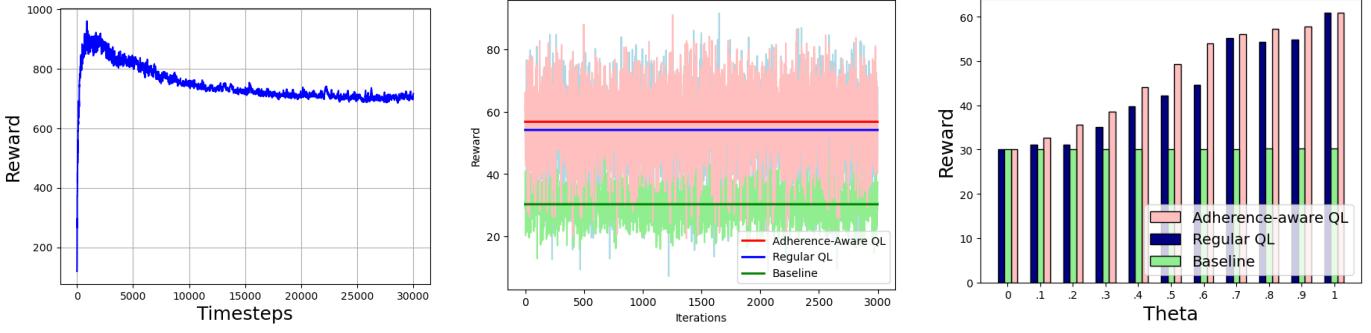
Fig. 2: Transition probabilities for repair (left) and wait (right).

C. Simulation results

We considered the learning rate α and the discount factor λ equal to 0.9 for the simulation results. We also considered that the update rule of θ will converge to 0.7. First, we show convergence for the two different demand distributions by tracking the initial state value for 10,000 and 100 time steps, respectively. In both numerical examples, the algorithm converges (Figs. 3a and 4a). For the inventory case, though, the algorithm needs more time to converge to optimal value due to the large state space. Furthermore, in both scenarios, we compared our adherence-aware Q-learning algorithm with the regular Q-learning and the baseline law, as it was the only one the HDM would implement. Figures 3b and 4b illustrate the result in which the average actual reward using the adherence-aware Q-learning algorithm is better than the other two. Finally, we investigated the algorithm's performance for different values of θ . Figures 3c and 4c illustrate the result and show that when $\theta \in [0, 0.5]$ our approach is slightly better than the baseline law, while for $\theta \in [0.5, 1]$ our approach outperforms the other two.

V. CONCLUDING REMARKS

In this letter, we proposed an “adherence-aware Q-learning” designed to derive optimal recommendation actions for HDMs. Our approach considers the complexity of the problem, where both the dynamics of the environment and the level of adherence to recommendations remain unknown. The structure of our algorithm is based on a combination of the HDM’s baseline law and an update rule for estimating the adherence level to the recommendations. We proved the convergence of the adherence-aware Q-learning algorithm to the optimal value function, and we applied this algorithm to two numerical examples, illustrating its ability to converge to the optimal value and outperform alternative methods in various system scenarios. Future work should consider situations where the baseline law is unknown and can be learned over time or the HDM has partial observability of the system state. Extending the results in a team of human-driven vehicles with a decentralized information structure [22] should be also a potential direction of future research.

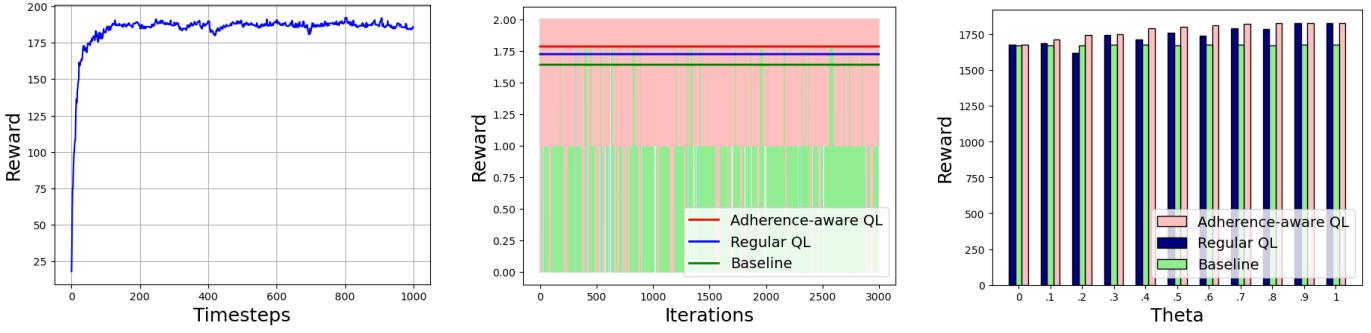


(a) Convergence of the adherence-aware algorithm.

(b) Actual reward for the three different approaches.

(c) Influence of the parameter θ .

Fig. 3: Inventory control results.



(a) Convergence of the adherence-aware algorithm.

(b) Actual reward for the three different approaches.

(c) Influence of the parameter θ .

Fig. 4: Machine replacement results.

REFERENCES

- [1] A. A. Malikopoulos, “Separation of learning and control for cyber-physical systems,” *Automatica*, vol. 151, no. 110912, 2023.
- [2] N. Venkatesh, V.-A. Le, A. Dave, and A. A. Malikopoulos, “Connected and automated vehicles in mixed-traffic: Learning human driver behavior for effective on-ramp merging,” in *Proceedings of the 62nd IEEE Conference on Decision and Control (CDC)*, 2023 (to appear, arXiv:2304.00397).
- [3] A. A. Malikopoulos, L. E. Beaver, and I. V. Chremos, “Optimal time trajectory and coordination for connected and automated vehicles,” *Automatica*, vol. 125, no. 109469, 2021.
- [4] F. D’Acunto, N. Prabhala, and A. G. Rossi, “The promises and pitfalls of robo-advising,” *The Review of Financial Studies*, vol. 32, no. 5, pp. 1983–2020, 2019.
- [5] W. Lin, S.-H. Kim, and J. Tong, “Does algorithm aversion exist in the field?” *An empirical analysis of algorithm use determinants in diabetes self-management. An Empirical Analysis of Algorithm Use Determinants in Diabetes Self-Management (July 23, 2021)*, 2021.
- [6] A. Dave, I. Faros, N. Venkatesh, and A. A. Malikopoulos, “Worst-case control and learning using partial observations over an infinite time-horizon,” 2023.
- [7] M. Balakrishnan, K. Ferreira, and J. Tong, “Improving human-algorithm collaboration: Causes and mitigation of over-and under-adherence,” *Available at SSRN 4298669*, 2022.
- [8] S. Kesavan and T. Kushwaha, “Field experiment on the profit implications of merchants’ discretionary power to override data-driven decision-making tools,” *Management Science*, vol. 66, no. 11, pp. 5182–5190, 2020.
- [9] J. M. Logg, J. A. Minson, and D. A. Moore, “Algorithm appreciation: People prefer algorithmic to human judgment,” *Organizational Behavior and Human Decision Processes*, vol. 151, pp. 90–103, 2019.
- [10] B. J. Dietvorst, J. P. Simmons, and C. Massey, “Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them,” *Management science*, vol. 64, no. 3, pp. 1155–1170, 2018.
- [11] E. Sabaté, *Adherence to long-term therapies: evidence for action*. World Health Organization, 2003.
- [12] K. H. Van Donselaar, V. Gaur, T. Van Woensel, R. A. Broekmeulen, and J. C. Fransoo, “Ordering behavior in retail stores and implications for automated replenishment,” *Management Science*, vol. 56, no. 5, pp. 766–784, 2010.
- [13] J. Sun, D. J. Zhang, H. Hu, and J. A. Van Mieghem, “Predicting human discretion to adjust algorithmic prescription: A large-scale field experiment in warehouse operations,” *Management Science*, vol. 68, no. 2, pp. 846–865, 2022.
- [14] B. J. Dietvorst, J. P. Simmons, and C. Massey, “Algorithm aversion: people erroneously avoid algorithms after seeing them err,” *Journal of Experimental Psychology: General*, vol. 144, no. 1, p. 114, 2015.
- [15] H. Bastani, O. Bastani, and W. P. Sinchaisri, “Improving human decision-making with machine learning,” *arXiv preprint arXiv:2108.08454*, 2021.
- [16] A. Jacq, J. Ferret, O. Pietquin, and M. Geist, “Lazy-mdps: Towards interpretable rl by learning when to act,” in *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 2022, pp. 669–677.
- [17] J. Grand-Clément and J. Pauphilet, “The best decisions are not the best advice: Making adherence-aware recommendations,” *arXiv preprint arXiv:2209.01874*, 2022.
- [18] S.-M. Hung and S. N. Givigi, “A q-learning approach to flocking with uavs in a stochastic environment,” *IEEE transactions on cybernetics*, vol. 47, no. 1, pp. 186–197, 2016.
- [19] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, “Safe learning in robotics: From learning-based control to safe reinforcement learning,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.
- [20] E. L. Lehmann and G. Casella, *Theory of point estimation*. Springer Science & Business Media, 2006.
- [21] A. F. Veinott Jr and H. M. Wagner, “Computing optimal (s, s) inventory policies,” *Management Science*, vol. 11, no. 5, pp. 525–552, 1965.
- [22] A. A. Malikopoulos, “On team decision problems with nonclassical information structures,” *IEEE Transactions on Automatic Control*, 2023.