

Reinforcement Learning with Partial Parametric Model Knowledge

Shuyuan Wang* Philip D. Loewen** Nathan P. Lawrence**
Michael G. Forbes*** R. Bhushan Gopaluni*

* *Department of Chemical and Biological Engineering, University of British Columbia, Vancouver, BC V6T 1Z3, Canada*
(e-mail: antergravity@gmail.com, bhushan.gopaluni@ubc.ca)

** *Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z2, Canada*
(e-mail: lawrence@math.ubc.ca, loew@math.ubc.ca)

*** *Honeywell Process Solutions, North Vancouver, BC V7J 3S4, Canada*
(e-mail: michael.forbes@honeywell.com)

Abstract: We adapt reinforcement learning methods for continuous control to bridge the gap between complete ignorance and perfect knowledge of the environment. Our method, Partial Knowledge Least Squares Policy Iteration (PLSPI), takes inspiration from both model-free RL and model-based control. It uses incomplete information from a partial model and retains RL’s data-driven adaption towards optimal performance. The linear quadratic regulator provides a case study; numerical experiments demonstrate the effectiveness and resulting benefits of the proposed method.

Keywords: Reinforcement Learning, partial parametric model information, LQR, LSPI, sample efficiency.

1. INTRODUCTION

Reinforcement learning (RL) is an online process that enables an agent to autonomously adapt to its environment. It has recently achieved great success on complex tasks such as playing Atari games (Mnih et al., 2015), winning at Go (Silver et al., 2016), and OpenAI Five (Berner et al., 2019). However, RL is famously data-hungry (Ye et al., 2021; Yu, 2018). Poor sample efficiency is a secondary concern in the purely virtual tasks mentioned above, but it is a critical limitation in real-world systems like robotics, healthcare, and industrial processes, where every observation costs time and money.

Researchers have made a lot of progress in sample efficient RL in recent years, including exploration (Plappert et al., 2017), environment modeling (Tamar et al., 2016), abstraction (Bacon et al., 2017), meta-RL (McClement et al., 2022) and more. In most of these cases, however, the RL formulation assumes the learner knows nothing of the system. In many practical problems, we have some prior information about the system, such as its order, structure, or even some of the model parameters. If RL could utilize such information instead of learning everything about its environment from scratch, we can reasonably expect an improvement in sample efficiency.

Of course, when a complete and accurate model for the system is available, classic model-based control techniques,

such as optimal control, will be preferred. Sampling and data are not required at all. One even has tidy analytical solutions in certain well-studied cases, such as the linear quadratic regulator (LQR). These classic control methods do not adapt to changes in the environment, however. Furthermore, complete and accurate models are vanishingly rare in practice (Spielberg et al., 2019). In this research, we aim to build a bridge from the theoretical world of LQR to the practical world of partial knowledge and inevitable uncertainty.

The recent advances in model-based RL (MBRL) share the motivations above (Deisenroth and Rasmussen, 2011). MBRL was also developed to improve sample efficiency, however, sample efficiency is achieved by trading off performance. In spite of the term “model-based” in its name, MBRL works in a different way from control. It first uses sampled data to fit a model, then uses the synthetic model to generate more data. These artificial data and real data are merged into a single batch of synthetic data to train a policy. Hence, the fundamental idea behind MBRL is still feeding the learning agent with more data. In contrast, this research aims to develop a learning framework using *sample-free* control techniques while retaining the ability to explore the uncertain part of the environment in pursuit of optimal performance.

Within the limited research on this topic, (Tamar et al., 2011; Shelton, 2002) define and use partial model knowledge in an RL framework, but their methods are based on discrete Markov decision process (MDP) models with finite

* The authors gratefully acknowledge the financial support of Natural Sciences and Engineering Research Council of Canada and Honeywell Process Solutions.

action and state spaces. This does not fit with continuous control.

In the initial explorations described here, we focus on LQR, but we expect nonlinear applications to follow. In practice, LQR models are often used to approximate and control more general systems. Further, even though LQR is the most basic and important optimal control problem with unbounded, continuous state and action spaces, the problem of incorporating *a priori* information has not been thoroughly investigated.

To this end, we propose a novel off-policy algorithm named Partial Knowledge Least Squares Policy Iteration (PLSPI) for learning-based LQR with partially known parametric model information. Specifically, our method combines classic LSPI (Lagoudakis and Parr, 2003), which is a model-free method utilizing a linear structure of the value function, with optimal control techniques.

The method proposed here offers two key benefits:

- (1) PLSPI gives RL the ability to use partial parametric system information, so that the sample efficiency can be improved.
- (2) PLSPI provides a conceptual link between model-free learning and model-based control.

2. PRELIMINARIES

2.1 Notation

For any fixed vector q in \mathbb{R}^n , the quadratic form $q^T Q q$ is a linear function of its symmetric $n \times n$ matrix Q . To make this explicit, define the “overbar” function on vectors and matrices. Let \bar{q} denote the column vector of length $\frac{1}{2}n(n+1)$ whose elements list all possible products of two elements from q , i.e.,

$$\bar{q} = [q_1^2, q_1 q_2, \dots, q_1 q_n, q_2^2, q_2 q_3, \dots, q_2 q_n, q_3^2, \dots, q_n^2]^T.$$

Then let \bar{Q} denote the vector of length $\frac{1}{2}n(n+1)$ whose elements list the upper triangular entries in Q , ordering the elements of \bar{q} and \bar{Q} so that the usual quadratic form involving q and Q turns into an inner product:

$$q^T Q q = \bar{q}^T \bar{Q}.$$

Use $\mathbf{Tr}(\cdot)$ to represent the trace of a matrix; and $\mathbb{E}(\cdot)$ for the mathematical expectation; let $[\cdot]^T$ and $[\cdot]'$ equivalently represent transpose of a vector or a matrix.

2.2 The linear quadratic regulator (LQR)

LQR aims to minimize the accumulated quadratic cost for a linear dynamical system. This paper focuses on the infinite horizon time-invariant discrete-time LQR, in which the cost function is

$$r(x_t, u_t) = x_t^T Q x_t + u_t^T R u_t$$

for prescribed matrices $Q = Q^T \geq 0$ and $R = R^T > 0$. In the general formulation, a discount factor $\gamma \in (0, 1]$ is given and the problem is

$$\begin{aligned} \min \quad & J = \mathbb{E} \left(\sum_{t=0}^{\infty} \gamma^t r(x_t, u_t) \right) \\ \text{s.t.} \quad & x_{t+1} = A x_t + B u_t + \xi_t \end{aligned} \quad (1)$$

Here the ξ_t are independent and identically distributed Gaussian random vectors with $\mathbb{E}(\xi_t) = 0$ and $\mathbb{E}(\xi_t \xi_t^T) = W$. The deterministic case arises when $W = 0$; in this situation the expectation is superfluous, and we typically consider $\gamma = 1$. In the stochastic case, the covariance $W = W^T > 0$ is known and we choose $\gamma \in (0, 1)$.

Both stochastic and deterministic LQR problems can be solved using dynamic programming, where the value function is first evaluated, and then an optimal controller with linear form is obtained. An advantage of the LQR formulation is that the cost-to-go function is a quadratic function of state x . The value function is obtained by solving a Discrete Algebraic Riccati Equation (DARE)

$$P = Q + \gamma A^T P A - \gamma A^T P B (R + B^T P B)^{-1} B^T P A \quad (2)$$

where P induces a positive definite quadratic form which can be interpreted as a value function for the LQR problem. For deterministic LQR, the value function is of the form

$$V(x_t) = x_t^T P x_t, \quad (3)$$

while for stochastic LQR, the value function includes a term independent of the state

$$V(x_t) = x_t^T P x_t + \frac{\gamma}{1-\gamma} \mathbf{Tr}(W P). \quad (4)$$

Once the matrix P has been determined, the unique optimal controller is given by

$$u_t^* = -K x_t = -\gamma (R + B^T P B)^{-1} B^T P A x_t. \quad (5)$$

A companion to the state value function V , that takes the state x as input, is the state-action value function Q , that takes state and action pairs (x, u) as inputs. It can also be written as a quadratic form. For deterministic LQR, the derivation is as follows (Bradtke et al., 1994):

$$\begin{aligned} Q_K(x, u) &= r(x, u) + \gamma V_K(f(x, u)) \\ &= x^T Q x + u^T R u + \gamma (A x + B u)^T P_K (A x + B u) \\ &= [x, u]^T \begin{bmatrix} Q + \gamma A^T P_K A & \gamma A^T P_K B \\ \gamma B^T P_K A & R + \gamma B^T P_K B \end{bmatrix} [x, u] \\ &= [x, u]^T \begin{bmatrix} H_{K(11)} & H_{K(12)} \\ H_{K(21)} & H_{K(22)} \end{bmatrix} [x, u] \\ &= [x, u]^T H_K [x, u]. \end{aligned} \quad (6)$$

Under our hypotheses, the matrix H_K is positive definite. For stochastic LQR, there is an additional constant term in it

$$Q_K(x, u) = [x, u]^T H_K [x, u] + \frac{\gamma}{1-\gamma} \mathbf{Tr}(W P_K) \quad (7)$$

where H_K has the exact same form as that of deterministic case.

2.3 Policy iteration for LQR

Policy iteration is the most important and popular method for model-free control. Policy iteration involves two repeated steps: policy evaluation, and policy improvement. For LQR, both steps have analytical solutions. As shown in (6), the cost-to-go function has been parameterized as a quadratic function with respect to state-action pairs. Hence, with the help of the ‘overbar’ function $\bar{[\cdot]}$, the quadratic form can be written as a linear form, and least squares estimation (LSE) can be applied to the policy evaluation step.

Bradtke et al. (1994) first followed this straightforward idea and developed the policy iteration method for deterministic LQR. Given a tuple of data (x_i, u_i, r_i, x_{i+1}) belonging to dataset \mathcal{D} , the following recursion can be written

$$\begin{aligned} Q_K(x_i, u_i) &= r(x_i, u_i) + \gamma V_K(x_{i+1}) \\ &= r(x_i, u_i) + \gamma Q_K(x_{i+1}, Kx_{i+1}). \end{aligned} \quad (8)$$

With its quadratic representation, the equation for identifying the state-action value function can be constructed

$$\begin{aligned} [x_i, u_i]' H_K [x_i, u_i] &= r(x_i, u_i) \\ &\quad - \gamma [x_{i+1}, Kx_{i+1}]' H_K [x_{i+1}, Kx_{i+1}] \end{aligned} \quad (9)$$

Using the ‘overbar’ function $\overline{[\cdot]}$ for vectors and matrices, equation (9) can be written as

$$\left(\overline{[x_i, u_i]}' - \gamma \overline{[x_{i+1}, Kx_{i+1}]}' \right) \overline{H}_K = r(x_i, u_i). \quad (10)$$

For the policy improvement step, the updated controller is obtained through

$$K = -H_{K(22)}^{-1} H_{K(21)}. \quad (11)$$

Equation (10) is often referred to as *Bellman Residual Minimizing Approximation*. While also utilizing the quadratic structure of the value function, a more advanced approach is proposed in Least Squares Temporal Differences (LSTD) (Bradtke and Barto, 1996) and Least Squares Policy Iteration (LSPI) (Lagoudakis and Parr, 2003) for the policy evaluation step. The detailed derivation is out of scope for this paper, but the final estimation equation turns out to require only a slight modification of (10):

$$\begin{aligned} \overline{[x_i, u_i]} (\overline{[x_i, u_i]}' - \gamma \overline{[x_{i+1}, Kx_{i+1}]}') \overline{H}_K \\ = r(x_i, u_i) \overline{[x_i, u_i]}, \end{aligned} \quad (12)$$

where $\overline{[x_i, u_i]} (\overline{[x_i, u_i]}' - \gamma \overline{[x_{i+1}, Kx_{i+1}]}')$ is a square matrix. Line (12) is often referred as *Least-Squares Fixed-Point Approximation*. (Lagoudakis and Parr, 2003) state that this way of constructing the learning equation requires fewer samples and can obtain a superior policy. As shown above for stochastic LQR, there is a constant term in the value function. In this paper, we still use $\overline{[x, u]}' \overline{H}_K$ to approximate Q_K . This won’t affect too much on the results under noise with low variance, which is a reasonable assumption in practice.

In contrast to (Bradtke et al. (1994)), LSPI estimates the variable \overline{H}_K by summing all equations (12) with respect to all the tuples in the dataset \mathcal{D} and then solving the final equation

$$\begin{aligned} \sum_{i \in \mathcal{D}} \overline{[x_i, u_i]} (\overline{[x_i, u_i]}' - \gamma \overline{[x_{i+1}, Kx_{i+1}]}') \overline{H}_K \\ = \sum_{i \in \mathcal{D}} r(x_i, u_i) \overline{[x_i, u_i]}. \end{aligned} \quad (13)$$

Furthermore, the framework of LSPI contains an outer loop and an inner loop: the inner loop evaluates and improves the policy by iterating with *the same dataset*; the outer loop interacts with the environment and updates the dataset with the newest learned policy.

3. PROBLEM STATEMENT

This paper focuses on the stochastic LQR formulated above, with $W = \sigma_\omega^2 I$. Our method can also be adapted to deterministic settings.

We assume that we have, prior to running the algorithm, some information about system dynamics A and B . We are trying to learn an optimal solution, under the assumption that some of the elements of A and B are known, and others are unknown.

4. RL WITH PARTIAL KNOWLEDGE

In this section, we describe our method of endowing a model-free RL algorithm with partial model knowledge. Our method aims to reduce the sample data consumed by RL, and this is achieved by constructing a better estimation for the value function in the policy evaluation step. Specifically, we utilize optimal control results and elegantly fuse them into the LSPI scheme to enhance learning based LQR. A novel method named Partial Knowledge Least Square Policy Iteration (PLSPI) is developed to consider partial model information and improve sample efficiency in RL. The method is outlined as follows.

4.1 Representing Partial Model Information

Our method of considering partial model information is based on decomposing the system dynamics A and B into two parts

$$\begin{aligned} A &= A_1 + A_2 \\ B &= B_1 + B_2, \end{aligned} \quad (14)$$

where A_1 and B_1 contain all the known parameters, and A_2 and B_2 contain all the unknown parameters. Specifically, A_1 and B_1 set the parameters located in the unknown place as 0; A_2 and B_2 set all the known part as 0, and keeps the unknown part there. However, it should be noted that one can have multiple choices of value to plug-in. If some inaccurate estimation of the unknown parts exists, plug-in the estimated value would be a better choice.

Take a second order scalar system for instance. Knowing the order of the system, the decomposition can be constructed as

$$\begin{aligned} x_{n+1} &= \begin{bmatrix} 0 & 1 \\ ? & ? \end{bmatrix} x_n + B u_n \\ &\quad \downarrow \\ x_{n+1} &= \left(\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} ? & 0 \\ ? & ? \end{bmatrix} \right) x_n + B u_n \end{aligned} \quad (15)$$

where the left matrix is A_1 , the right matrix is A_2 . It will be illustrated later that without identifying A_2 , the LQR can still be solved with online data and A_1 .

With A_1 and B_1 , and a given controller K_1 , a sub-model containing known model information can be constructed as

$$\tilde{x}_1 = A_1 x_0 - B_1 K_1 x_0 \quad (16)$$

where \tilde{x}_1 represents the consequent state executed from the sub-model, and x_0 is the sampled data input to the sub-model. The sub-model shares the same cost matrix Q and R with the original model. The choice of K_1 will be discussed in subsection 4.3.

With the sub-model, a DARE related to A_1 and B_1 can be obtained as

$$P_{K_1} = Q + K_1^T R K_1 + \gamma(A_1 - B_1 K_1)^T P_{K_1} (A_1 - B_1 K_1). \quad (17)$$

With (17), P_{K_1} can be solved offline. After having P_{K_1} , the corresponding state-action value function H_{K_1} can also be calculated offline, as follows:

$$H_{K_1} = \begin{bmatrix} R + \gamma A_1^T P_{K_1} A_1 & \gamma A_1^T P_{K_1} B_1 \\ \gamma B_1^T P_{K_1} A_1 & Q + \gamma B_1^T P_{K_1} B_1 \end{bmatrix} \quad (18)$$

The result will be introduced into reinforcement learning process in the next subsection.

4.2 Introducing Partial Model Information into RL

Having H_{K_1} , the partial knowledge can be transferred into the RL process. With the property of value function, that the cost-to-go from current state-action can be expanded as one stage cost plus the cost-to-go from the consequent state-action, the following equation can be constructed

$$([\overline{x_i, u_i}]' - \gamma[\overline{\tilde{x}_{i+1}, K_1 \tilde{x}_{i+1}}]') \overline{H}_{K_1} = r(x_i, u_i), \quad (19)$$

where (x_i, u_i) denotes a (state,action) pair sampled from the real system. Unlike the state x_{i+1} sampled online, \tilde{x}_{i+1} represents the virtual next state obtained by stepping the sub-model from the current state x_i :

$$\tilde{x}_{i+1} = A_1 x_i - B_1 K_1 x_i \quad (20)$$

The core of integrating partial model information into RL is achieved by forming the difference between (9) and (19):

$$\begin{aligned} &([\overline{x_i, u_i}]' - \gamma[\overline{x_{i+1}, K x_{i+1}}]') \overline{H}_K - ([\overline{x_i, u_i}]' \\ &- \gamma[\overline{\tilde{x}_{i+1}, K_1 \tilde{x}_{i+1}}]') \overline{H}_{K_1} = r(x_i, u_i) - r(x_i, u_i) \\ \Rightarrow &([\overline{x_i, u_i}]' - \gamma[\overline{x_{i+1}, K x_{i+1}}]') (\overline{H}_K - \overline{H}_{K_1}) \\ &= \gamma([\overline{x_{i+1}, K x_{i+1}}]' - [\overline{\tilde{x}_{i+1}, K_1 \tilde{x}_{i+1}}]') \overline{H}_{K_1} \end{aligned} \quad (21)$$

Define $\overline{H}_K - \overline{H}_{K_1}$ as \overline{H}_{K_2} . We will have

$$\begin{aligned} &([\overline{x_i, u_i}]' - \gamma[\overline{x_{i+1}, K x_{i+1}}]') \overline{H}_{K_2} \\ &= \gamma([\overline{x_{i+1}, K x_{i+1}}]' - [\overline{\tilde{x}_{i+1}, K_1 \tilde{x}_{i+1}}]') \overline{H}_{K_1}. \end{aligned} \quad (22)$$

We hypothesize that this will lead to an improvement on sample efficiency. The idea behind this operation is that instead of identifying H_K from scratch, we can instead identify the gap between what we have known (or a initial guess/baseline of it) and the actual value of it, and hence save data. Note that the right-hand side of the equation can be viewed as a reformulated reward after introducing partial model information. Plus, since the Q function takes state and action as input, $r(x_i, u_i)$ can always be cancelled out no matter what value K_1 is fed into the sub-model (20).

Inspired by LSPI, a more advanced way to construct the LSE equation is used. This can be simply done by slightly adjusting (22)

$$\begin{aligned} &([\overline{x_i, u_i}]([\overline{x_i, u_i}]' - \gamma[\overline{x_{i+1}, K x_{i+1}}]') \overline{H}_{K_2} \\ &= \gamma[\overline{x_i, u_i}]([\overline{x_{i+1}, K x_{i+1}}]' - [\overline{\tilde{x}_{i+1}, K_1 \tilde{x}_{i+1}}]') \overline{H}_{K_1}. \end{aligned} \quad (23)$$

Given a batch of data \mathcal{D} , the overall policy evaluation equation is given by

$$\begin{aligned} &\sum_{i \in \mathcal{D}} [\overline{x_i, u_i}]([\overline{x_i, u_i}]' - \gamma[\overline{x_{i+1}, K x_{i+1}}]') \overline{H}_{K_2} \\ &= \gamma \sum_{i \in \mathcal{D}} [\overline{x_i, u_i}]([\overline{x_{i+1}, K x_{i+1}}]' - [\overline{\tilde{x}_{i+1}, K_1 \tilde{x}_{i+1}}]') \overline{H}_{K_1}. \end{aligned} \quad (24)$$

Now (24) is the improved policy evaluation equation, that takes both online data and prior model into account. The sample efficiency can benefit from the improved equation.

4.3 Overall Algorithm

In (23), the choice of K_1 is still pending. To facilitate a connection between the sub-model and the real model, the preferred choice of K_1 is the same controller K that is currently evaluated. However, sometimes K may destabilize the sub-model. When this happens, K_1 will be chosen as the optimal controller of the sub-model, which can be obtained by first solving (17) and then plugin it into (5). The optimal controller will be noted as K_1^* .

Another way to choose K_1 is to fix it as K_1^* through the whole learning process. We tested this method, and it has similar performance but the hybrid method above works slightly better.

Following LSPI, PLSPI also contains two loops. It is summarized as Algorithm 1.

Algorithm 1 Partial Knowledge LSPI (PLSPI)

- 1: **Input:** Rollout numbers N ; Time horizon T ; Exploration noise ϵ_t ; Partial model A_1 and B_1 ; Maximum iterations I_1 and I_2 .
 - 2: **Initialization:** Initialize controller K .
 - 3: **for** $i = 0, \dots, I_1$ **do**
 - 4: $\mathcal{D} \leftarrow \emptyset$
 - 5: Execute $u_t = -K x_t + \epsilon_t$ for N episodes. Store $N \times T$ pairs of data for \mathcal{D} .
 - 6: **for** $j = 0, \dots, I_2$ **do**
 - 7: **if** $A_1 - B_1 K$ is stable **then**
 - 8: $K_1 = K$
 - 9: **else**
 - 10: $K_1 = K_1^*$
 - 11: Solve (17)(18) to obtain H_{K_1} .
 - 12: Identify H_{K_2} with (24), and obtain H_K with $H_K = H_{K_1} + H_{K_2}$.
 - 13: Improve policy with $K = -H_{K(22)}^{-1} H_{K(21)}$.
 - 14: **return** K
-

Remark 1. Our method can serve as a bridge between control and RL. Consider the extreme cases. When the prior information is the full model, and the initial controller is initialized with the full model, the estimation actually needs no data: it reduces to solving a fully-specified optimal control problem. When prior information is completely absent, then the scheme reduces to classic LSPI, which finds the optimal policy totally by data.

5. NUMERICAL EXAMPLES

We present several simulation examples to illustrate our algorithm. The first example tests the effectiveness of

LSPI and the proposed method on LQR, specifically on a deterministic undiscounted setting. The second example is a more complex and unstable model with practical origins (cooling system) which is used for comparison purposes. In all experiments, the exploration noise follows *i.i.d.* Gaussian random vectors with $\mathbb{E}(\epsilon_t) = 0$ and $\mathbb{E}(\epsilon_t \epsilon_t^T) = \sigma_\eta^2 I$, with $\sigma_\eta^2 = 0.1$. The controller K is initialized as 0.

5.1 Example 1

Consider the classic discrete-time double integrator with dynamic coefficients

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

and quadratic cost coefficients

$$Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad R = 1.$$

The system is deterministic so that the noise $\xi_t = 0$. The discount factor γ is set to 1. The eigenvalues of the system both equal 1, meaning that the system is on the boundary of stability.

For our PLSPI method, the partial information is set as

$$A_1 = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, \quad B_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The problem is simple, so both LSPI and PLSPI perform well. They converge to the optimal controller within just 1 or 2 iterations. Each iteration contains 30 rollouts, with one rollout being a run with a horizon of 20 steps.

5.2 Example 2

Next we consider a simplified model of a 3-level cooling system, which is a popular test scenario for RL applications (Tu and Recht, 2018; Krauth et al., 2019; Dean et al., 2020). The detail is as follows,

$$A = \begin{bmatrix} 1.01 & 0.01 & 0 \\ 0.01 & 1.01 & 0.01 \\ 0 & 0.01 & 1.01 \end{bmatrix},$$

$B = I$, $Q = I$, $R = 1000I$, $\sigma_\omega^2 = 0.01$, and $\gamma = 0.98$. The open-loop system is unstable, which increases the difficulty for the learning algorithm. To make the problem harder, we use the scale factor 1000 to penalize control action through R much more heavily than state offsets through Q . This puts the maximum eigenvalue of the optimal closed-loop around 0.98, making the system challenging to tune. We assume knowledge of only the diagonal entries in the dynamic matrices, setting the partial information as

$$A_1 = \begin{bmatrix} 1.01 & 0 & 0 \\ 0 & 1.01 & 0 \\ 0 & 0 & 1.01 \end{bmatrix}, \quad B_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The performance of PLSPI and LSPI is compared in terms of the average trend and the variation of the evolution from multiple simulations. Specifically, we run LSPI and our PLSPI for 10 times each, and obtain Fig. 1.

Particularly, two versions of LSPI are tested here. LSPI-v1 follows the original setting, where the inner loop iterates for multiple times (5 in our case) given a fixed batch of data; LSPI-v2 adopts the setting in (Krauth et al., 2019),

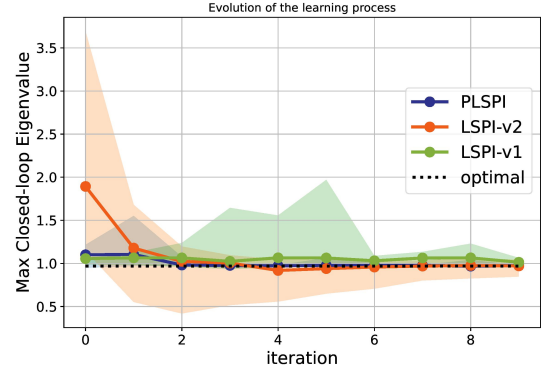


Fig. 1. The comparison between different methods in terms of the evolution of closed-loop system with respect to iterations. Each iteration contains 30 rollouts, with each rollout running 20 steps. The shaded area represents the variance of each method, with percentile 0-75%. The solid line represents the median of each method.

where there is no inner loop, meaning each batch of data is only used for once.

Fig. 1 shows PLSPI converging faster than both versions of LSPI. More importantly, the shaded region showing observed variance is also smaller for PLSPI, which indicates another benefit: given limited data, PLSPI learns the optimal controller with better accuracy. These benefits contribute to the better sample efficiency of our method.

The stochastic system response is also tested, with the results shown in Fig.2. Trajectories in Fig.2 show that the learned controller can regulate the states under a stochastic system point of view.

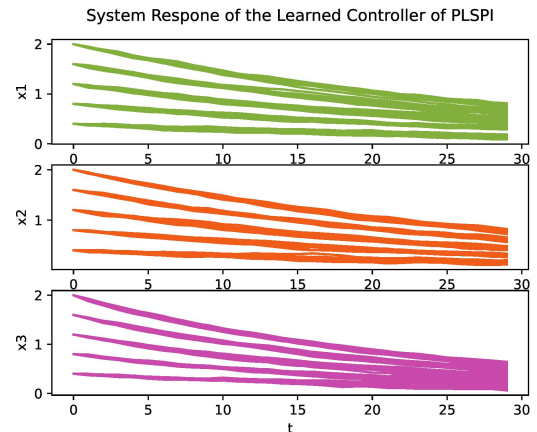


Fig. 2. Stochastic trajectories generated from the learned controller. The X -axis represents simulation steps.

We further apply the algorithm to a completely known system model (knowledge on noise term is not required) to test an extreme case. Fig.3 shows the result.

In Fig. 3, the learner gets to the optimal solution with only 1 iteration, and evolves without any variance. The initial variance is related to the random initialization of the controller. If the controller is initialized with the given

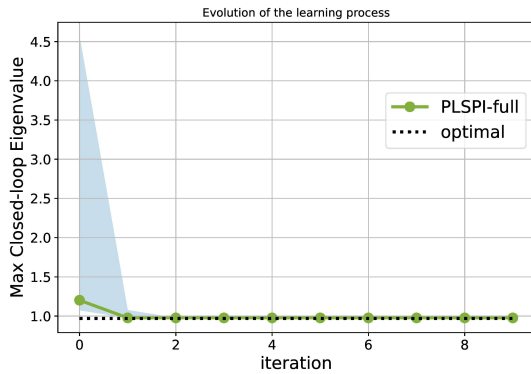


Fig. 3. Results of PLSPI given a completely known system model, with percentile 0-100%.

model, the variance can also be eliminated. This result follows the purpose of the algorithm design: given complete information about the system, the algorithm will follow a fully optimal control solution process, without consuming any data.

6. CONCLUSION

In control tasks, some partial parametric model information is often known, but seemingly under-utilized in model-free RL. Taking it as a starting point, we have adopted complementary benefits in model-free RL and model-based control to develop a framework that aims to bridge both sides. This paper is a proof of concept and there are many avenues to explore. These include the exact learning for value function under stochastic setting; theoretical guarantee on when partial information is useful; and extensions to nonlinear systems. We believe this is a promising area to investigate further as RL gains traction in process systems engineering.

REFERENCES

- Bacon, P.L., Harb, J., and Precup, D. (2017). The option-critic architecture. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1). doi: 10.1609/aaai.v31i1.10916.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al. (2019). Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- Bradtke, S., Ydstie, B., and Barto, A. (1994). Adaptive linear quadratic control using policy iteration. In *Proceedings of 1994 American Control Conference - ACC '94*, volume 3, 3475–3479.
- Bradtke, S.J. and Barto, A.G. (1996). Linear least-squares algorithms for temporal difference learning. *Machine learning*, 22(1), 33–57.
- Dean, S., Mania, H., Matni, N., Recht, B., and Tu, S. (2020). On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4), 633–679.
- Deisenroth, M. and Rasmussen, C.E. (2011). PILCO: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, 465–472.
- Krauth, K., Tu, S., and Recht, B. (2019). Finite-time analysis of approximate policy iteration for the linear quadratic regulator. *Advances in Neural Information Processing Systems*, 32.
- Lagoudakis, M.G. and Parr, R. (2003). Least-squares policy iteration. *The Journal of Machine Learning Research*, 4, 1107–1149.
- McClement, D.G., Lawrence, N.P., Backström, J.U., Loewen, P.D., Forbes, M.G., and Gopaluni, R.B. (2022). Meta-reinforcement learning for the tuning of PI controllers: An offline approach. *Journal of Process Control*, 118, 139–152.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540), 529–533.
- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R.Y., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. (2017). Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*.
- Shelton, C.R. (2002). Reinforcement learning with partially known world dynamics. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence (UAI)*, 8, 461–468. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587), 484–489.
- Spielberg, S., Tulsyan, A., Lawrence, N.P., Loewen, P.D., and Bhushan Gopaluni, R. (2019). Toward self-driving processes: A deep reinforcement learning approach to control. *AIChE journal*, 65(10), e16689.
- Tamar, A., Di Castro, D., and Meir, R. (2011). Integrating partial model knowledge in model free rl algorithms. In *ICML*.
- Tamar, A., Wu, Y., Thomas, G., Levine, S., and Abbeel, P. (2016). Value iteration networks. *Advances in neural information processing systems*, 29.
- Tu, S. and Recht, B. (2018). Least-squares temporal difference learning for the linear quadratic regulator. In *International Conference on Machine Learning*, 5005–5014. PMLR.
- Ye, W., Liu, S., Kurutach, T., Abbeel, P., and Gao, Y. (2021). Mastering Atari games with limited data. *Advances in Neural Information Processing Systems*, 34, 25476–25488.
- Yu, Y. (2018). Towards sample efficient reinforcement learning. In *IJCAI*, 5739–5743.