

A New Algorithm for the LQR Problem with Partially Unknown Dynamics^{*}

Agnese Pacifico^(✉)[0000–0002–5198–5984], Andrea Pesare^[0000–0002–0131–0268],
and Maurizio Falcone^[0000–0002–4277–2370]

Dipartimento di Matematica, Sapienza Università di Roma, Rome, Italy
pacifico.1699761@studenti.uniroma1.it, {pesare,falcone}@mat.uniroma1.it

Abstract. We consider an LQR optimal control problem with partially unknown dynamics. We propose a new model-based online algorithm to obtain an approximation of the dynamics *and* the control at the same time during a single simulation. The iterative algorithm is based on a mixture of Reinforcement Learning and optimal control techniques. In particular, we use Gaussian distributions to represent model uncertainty and the probabilistic model is updated at each iteration using Bayesian regression formulas. On the other hand, the control is obtained in feedback form via a Riccati differential equation. We present some numerical tests showing that the algorithm can efficiently bring the system towards the origin.

Keywords: Reinforcement Learning · LQR problem · Numerical methods

1 Introduction

The Linear-Quadratic Regulator (LQR) optimal control problem [1,3] is a classical problem in control theory with a wide range of applications. When the dynamics is known, the optimal control is obtained in feedback form solving a backward Riccati differential equation.

Some Reinforcement Learning (RL) problems can be seen as LQR problems where the dynamics is partially or completely unknown. Some RL algorithms try to learn a model for the dynamics and use this model to find the optimal policy. These are called *model-based algorithms*. Others recover the optimal control directly, without reconstructing the dynamics, and these are called *model-free algorithms*. For a broad overview on RL, we recommend [14].

The connection between RL and optimal control theory was already identified in the past [15]. Recently, Palladino and co-authors tried to clarify this relationship via some rigorous proofs, identifying some RL tasks as real optimal control problems with unknown dynamics [8,11,10]. In this context, we propose a *new model-based algorithm for LQR problems* where the dynamics is partially unknown, which takes contributions from both fields. In fact, our algorithm can be considered as a case of Bayesian RL, a class of model-based algorithms where

^{*} This research has been partially supported by the INdAM Research group GNCS.

the controller builds a stochastic model of the dynamics and updates it according to Bayesian statistics [6]. On the other hand, we borrow the LQR solution from optimal control theory, to get the synthesis of a suitable control [1].

In particular, our algorithm is similar to PILCO [2], from which it takes the use of Gaussian distributions and the whole process of Bayesian update. However, we propose here some novelties. The first is that in PILCO the optimal control is chosen through a gradient descent algorithm in a class of controls, whereas we solve a Riccati differential equation to identify the minimizer. The second is that PILCO needs several trials to reconstruct the dynamics and stabilise the system. Our algorithm is designed to approximate the dynamics *and* to find a suitable control in a single run but can be applied only to linear dynamical systems. Finally, let us mention that other works have already dealt with LQR problems with unknown dynamics (see i.e. [4,5,7,9] and references therein), but they all need several trials to converge, whereas our method works with just one.

The paper is structured as follows. In Section 2 we recall the LQR problem. In Section 3 we present our algorithm and discuss some implementation details. Finally, in Section 4 we show and discuss some numerical tests.

2 The classical LQR problem

The Linear-Quadratic Regulator (LQR) problem [1,3] is an optimal control problem with linear dynamics and quadratic cost. In the finite horizon case, the state of the system $x(t) \in \mathbb{R}^n$ evolves according to the following controlled dynamics

$$\begin{cases} \dot{x}(t) = \hat{A}x(t) + Bu(t), & t \in [0, T] \\ x(0) = x_0. \end{cases} \quad (1)$$

We will denote by $\mathbb{R}^{m \times n}$ the space of matrices with m rows and n columns; \mathbb{I}_n and $\mathbf{0}_n$ will be respectively the identity and zero matrices in $\mathbb{R}^{n \times n}$. Here $\hat{A} \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$. The control function $u(t) \in \mathbb{R}^m$ must be chosen among the admissible controls $\mathcal{U} := \{u : [0, T] \rightarrow \mathbb{R}^m \text{ Lebesgue measurable}\}$ to minimize the quadratic cost functional

$$J_{x_0}[u] := \frac{1}{2} \left(\int_0^T (x(t)^T Q x(t) + u(t)^T R u(t)) dt + x(T)^T Q_f x(T) \right). \quad (2)$$

The *main assumptions* on the cost matrices are the following:

- $Q, Q_f \in \mathbb{R}^{n \times n}$ are symmetric and positive semi-definite;
- $R \in \mathbb{R}^{m \times m}$ is symmetric and positive definite.

For any $x_0 \in \mathbb{R}^n$, we will call \bar{u} an optimal control starting from x_0 if and only if

$$J_{x_0}[\bar{u}] \leq J_{x_0}[u] \quad \forall u \in \mathcal{U}. \quad (3)$$

When the dynamics is fully known, the optimal control can be obtained in feedback form [1]. Indeed, if $P(t)$ is the unique symmetric solution of the Riccati

differential equation

$$\begin{cases} -\dot{P}(t) = \widehat{A}^T P(t) + P(t)\widehat{A} - P(t)BR^{-1}B^T P(t) + Q, & t \in [0, T] \\ P(T) = Q_f, \end{cases} \quad (4)$$

the optimal control is given by $\bar{u}(t) = -R^{-1}B^T P(t)x(t)$.

We want to investigate: what happens if the dynamics is partially unknown? How can one find a suitable control? We considered the following framework:

Setting. We assume that the matrices $B \in \mathbb{R}^{n \times m}$, $Q, Q_f \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$ are given, whereas the state matrix $\widehat{A} \in \mathbb{R}^{n \times n}$ is unknown.

3 An Online Algorithm for the LQR Problem

In this section, we describe our model-based algorithm able to solve the LQR problem without knowing the matrix \widehat{A} . The algorithm is online, meaning that it doesn't need any previous simulation or computation. Our goal is twofold: first, we look for a good estimate for the unknown dynamics matrix \widehat{A} ; and secondly, we want to choose a control that can steer the trajectory towards the origin. Furthermore, we assume that we can run the experiment just once, so the system must be controlled while the dynamics is still uncertain. To this end, we use a technique to get an estimate of the matrix \widehat{A} and to update the control at the same time.

We divide the interval $[0, T]$ into equal time steps of length Δt , globally we will have N time steps and we group them into rounds of S steps each. The i -th round will be denoted by $[t_{i-1}, t_i]$ and a superscript index will indicate a single time step:

$$t_i^j = t_{i-1} + j\Delta t, \quad j = 0, \dots, S.$$

The current knowledge of the dynamics matrix is represented as a probability distribution over matrices. For each round, two major operations are carried out:

1. At the beginning of each round, a probability π_{i-1} is given from the previous round. We use the mean \bar{A}_{i-1} of this distribution and compute a feedback control for the round solving a Riccati equation;
2. At the end of each round, the current probability distribution is updated using Bayesian formulas, according to the trajectory observed during the round. The output is a new probability distribution π_i .

The whole algorithm is summarised below as Algorithm 1. In the following, we will give more technical details.

Prior distribution. The algorithm requires the choice of a prior distribution π_0 . We fix some $m_0 \in \mathbb{R}^n$ and $\Sigma_0 \in \mathbb{R}^{n \times n}$ and consider a random matrix A such that each of its rows r_k is distributed as an independent gaussian vector with

Algorithm 1 An online algorithm for the LQR problem

Divide $[0, T]$ into M intervals of length Δt , with $\Delta t = \frac{T}{M}$;
 Group the intervals in *rounds*, each containing S intervals;
 Choose a prior distribution π_0 over matrices;
for i from 1 to the *rounds* number **do**
 Find a feedback control u_i^* solving a Riccati eq. with the mean matrix \bar{A}_{i-1} ;
 Use u_i^* as control for all the steps in the i -th *round*;
 Observe the actual trajectory;
 Update π_i according to the data from the observed trajectory
end for

mean m_0 and covariance matrix Σ_0

$$A = \begin{bmatrix} r_1^T \\ \vdots \\ r_n^T \end{bmatrix} \quad r_k \sim \mathcal{N}(m_0, \Sigma_0) \quad \forall k = 1, \dots, n$$

A typical choice for the prior distribution, when no information about the true matrix \hat{A} is available, is $m_0 = (0, \dots, 0)^T$ and $\Sigma_0 = n m \mathbb{I}_n$.

Feedback control. At the beginning of the round $[t_{i-1}, t_i]$ our knowledge of the matrix \hat{A} is described by the distribution π_{i-1} . In order to find the control to apply, we solve the evolutive *Riccati* equation associated with the matrix \bar{A}_{i-1} , where \bar{A}_{i-1} is the mean of the distribution π_{i-1} . The Riccati equation reads

$$\begin{cases} -\dot{P}(t) = \bar{A}_{i-1}^T P(t) + P(t) \bar{A}_{i-1} - P(t) B R^{-1} B^T P(t) + Q & t \in [t_{i-1}, T] \\ P(T) = Q_f. \end{cases} \quad (5)$$

If we denote by $P_i(t)$ the solution of (5), our control will be the feedback control given by

$$u_i^*(t_i^j) = -R^{-1} B^T P_i(t_i^j) x_i^j \quad j = 0, \dots, S-1, \quad (6)$$

where $x_i^j = x(t_i^j)$ is the observed state of the system. Since the control must be defined for all instants $t \in [t_{i-1}, t_i]$, we choose a piecewise constant control: $u_i^*(t) = u_i^*(t_i^j)$ for $t \in [t_i^j, t_i^{j+1}]$. Note that we need real-time observations of the system state to compute u_i^* , since it depends on x .

Distribution update. At the end of each round we can update the probability distribution using *Bayesian regression formulas*. More precisely, we know that during the time step $[t_i^j, t_i^{j+1}]$ the system evolves according to (1) where we plug-in the chosen piecewise constant control. We can easily get an approximation of the state derivative with a finite difference scheme using the state observations $x_i^j = x(t_i^j)$. By a first-order scheme, rearranging the terms, we get

$$\hat{A} x_i^j \simeq \frac{x_i^{j+1} - x_i^j}{\Delta t} - B u^*(t_i^j) \quad j = 0, \dots, S-1. \quad (7)$$

We interpret these data as observations of the dynamics with a Gaussian noise due to the error in the derivative approximation and the measurements. We treat each row \widehat{r}_k of \widehat{A} separately. Denoting by y^j the right hand side in (7), we can write its k -th component as

$$y_{(k)}^j = \widehat{r}_k x_i^j + \varepsilon \quad \text{with } \varepsilon \sim \mathcal{N}(0, \sigma^2),$$

where σ is a parameter chosen according to the order of the derivative approximation.

For each row k we have a prior distribution $r_k \sim \mathcal{N}(m_0, \Sigma_0)$ given by π_{i-1} . We define X as the matrix whose columns are x_i^0, \dots, x_i^{S-1} and y as the column vector $y_{(k)}^0, \dots, y_{(k)}^{S-1}$. Therefore we obtain a posterior distribution $r_k|y, X \sim \mathcal{N}(m, \Sigma)$ where $\Sigma^{-1} = \frac{1}{\sigma^2} X X^T + \Sigma_0^{-1}$ and $m = \Sigma(Xy/\sigma^2 + \Sigma_0^{-1}m_0)$. For more details about Bayesian Linear Regression see for instance the extensive monograph by Rasmussen and Williams [12].

Higher-order schemes. The derivative in (7) can be approximated with higher order finite difference schemes [13]. Note that trajectory regularity is required in the interval containing the nodes used in the approximation. While first-order approximation uses only two nodes, higher-order approximations use more nodes, thus the control cannot jump at each time step and we have to keep it constant for more steps.

Remark 1 (Heuristic argument for convergence). The algorithm cannot find the optimal control for the problem, since at the beginning the matrix \widehat{A} is unknown and it needs at least some steps to have a good estimate for \widehat{A} . However, from Bayesian Regression theory (see [12]) we know that the more data we observe, the more precise our distribution π_i becomes, eventually converging to the Dirac delta $\delta_{\widehat{A}}$.

Furthermore, let us note that since π_i is converging to $\delta_{\widehat{A}}$, then also the mean matrix \bar{A}_i of π_i is converging to \widehat{A} . Thus, after few rounds, the feedback control computed by the algorithm (which is using \bar{A}_i) in the interval $[t_i, T]$ should be close to the optimal control of a trajectory of the real dynamics which starts from the same point x_{i-1}^S .

Finally, when $\Delta t \rightarrow 0$, the algorithm reaches earlier a good estimate of the dynamics matrix. This means that also the computed control is closer to the optimal one. All these heuristics are confirmed by the numerical simulations in the next section.

4 Some Numerical Tests

The following numerical tests for the algorithm described above were performed in MATLAB and took few seconds to run.

Test 1. We first consider a dynamical system where the state lies in \mathbb{R}^2 and the control is 1-dimensional, i.e. $n = 2$ and $m = 1$. The LQR problem is defined

by the following matrices:

$$\hat{A} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \quad B = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad Q = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad R = 0.1 \quad Q_f = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

The time horizon is set to $T = 5$ and the starting point is $x_0 = (0, 1)^T$. We assume that the matrix \hat{A} is unknown to the algorithm, though we use it to simulate the dynamics. We choose the prior distribution as recommended in Section 3, using $m_0 = (0, 0)^T$ and $\Sigma_0 = 2I_2$; for all the tests we set $\sigma = \sqrt{10\Delta t^p}$ and $S = 2p$, where p is the order of the scheme used in the derivative approximation. Fig. 1 shows the *piecewise constant* controls chosen by the algorithm with $p = 1$

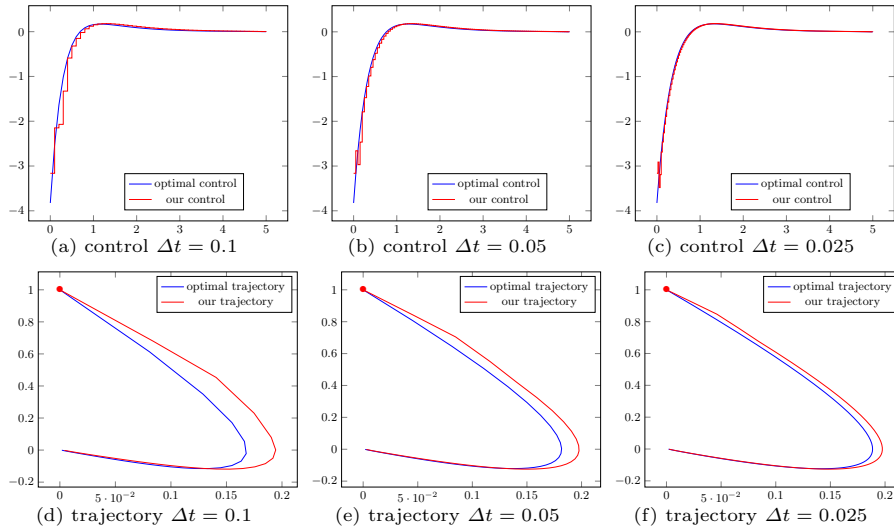


Fig. 1: Simulations of Test 1 for different values of Δt . The first row shows the control chosen by the algorithm as a function of time (in red) compared with the optimal control (in blue), computed knowing the matrix \hat{A} ; the second row shows the trajectories in \mathbb{R}^2 . The red dot is the trajectory starting point.

for different values of Δt and the corresponding trajectories. Recall that the matrix \hat{A} is completely unknown at the beginning, so the control we apply in the first steps depends only on the prior distribution we have chosen and clearly is not accurate. This causes the trajectory to deviate from the optimal solution. However, after few steps the matrix \hat{A} is well approximated and the algorithm manages to steer the state towards the origin anyway. In Table 1a we have reported the cost of the solution found by the algorithm for different choices of Δt . When Δt is smaller, the algorithm recovers the matrix \hat{A} quickly and thus the deviation from the optimal solution is smaller. This confirms the heuristics of Remark 1. We can also observe a numerical order of convergence equal to 1.

We tried different finite difference schemes for the approximation of the state derivatives (see "*Higher-order schemes*" in Section 3). Table 1b shows the error

Table 1: Numerical results for Test 1. (a): Cost of the trajectory for different values of Δt . The cost is compared with the optimal cost computed knowing \hat{A} (C^* last row). (b): Error for \hat{A} after the simulation using different Δt and different finite difference schemes for the derivatives approximation; p is the scheme order.

(a)				(b)					
Δt	Cost	Error	Order	$p = 1$		$p = 2$		$p = 4$	
Δt	Error	Order	Error	Order	Error	Order	Error	Order	
0.1	0.3897	0.0063	-	0.167	-	9.74e-3	-	1.63e-4	-
0.05	0.3866	0.0032	0.98	0.087	0.94	1.91e-3	2.3	3.00e-6	5.8
0.025	0.3849	0.0015	1.09	0.045	0.95	4.14e-4	2.2	1.03e-7	4.8
C^*	0.3834	-	-	0.018	1.00	6.30e-5	2.1	1.70e-9	4.5

in the approximation of \hat{A} at the end of the simulation, when using schemes of order $p = 1, 2, 4$ and for different values of Δt . As expected, when we consider more accurate approximations of the gradient, we get better estimations of the matrix \hat{A} . Unfortunately, the same does not hold for the solution costs, which are not significantly improved if compared with the ones found by the first-order approximation.

Test 2. For the second test we choose $n = 4, m = 3$ and $T = 10$. Our matrices are

$$\hat{A} = \begin{pmatrix} -0.0215 & -0.7776 & -0.1922 & 0.9123 \\ -0.3246 & 0.5605 & -0.8071 & 0.1504 \\ 0.8001 & -0.2205 & -0.7360 & -0.8804 \\ -0.2615 & -0.5166 & 0.8841 & -0.5304 \end{pmatrix}, B = \begin{pmatrix} -0.2937 & -0.6620 & -0.0982 \\ 0.6424 & 0.2982 & 0.0940 \\ -0.9692 & 0.4634 & -0.4074 \\ -0.9140 & 0.2955 & 0.4894 \end{pmatrix},$$

$Q = \frac{1}{4}\mathbb{I}_4, R = \frac{1}{3}\mathbb{I}_3$ and $Q_f = \mathbb{I}_4$, and the starting point is $x_0 = (1, 1, 1, 1)^T$. We set $\Delta t = 0.025, S = 4$ and use a second order approximation for the derivatives. Fig. 2 shows the control found by the algorithm and Fig. 3 the corresponding trajectory. The behaviour observed in Test 1 is even more visible here: in the first steps the control is not accurate since we do not know \hat{A} , but after few steps the algorithm learns more on the matrix and manages to bring the state to the origin. The cost of the control found by the algorithm is 1.111, whereas the optimal cost computed using \hat{A} is 1.056.




Fig. 2: The three components of the control in Test 2




Fig. 3: The trajectory of Test 2 in \mathbb{R}^4 , represented by projecting components in couples: (x_1, x_2) on the left and (x_3, x_4) on the right. The red dots indicate the starting point of the trajectory.

5 Conclusions

We proposed a new algorithm designed to deal with LQR problems when the dynamics is partially unknown. Numerical tests presented in Section 4 showed how it manages to approximate the dynamics *and* to find a suitable control that brings the system towards the origin in a single simulation. Future works include the convergence analysis of the algorithm and possible extensions to the nonlinear case.

References

1. Anderson, B.D., Moore, J.B.: Optimal control: linear quadratic methods. Courier Corporation (2007)
2. Deisenroth, M., Rasmussen, C.E.: PILCO: A model-based and data-efficient approach to policy search. In: Proceedings ICML-11. pp. 465–472 (2011)
3. Fleming, W.H., Rishel, R.W.: Deterministic and stochastic optimal control, vol. 1. Springer Science & Business Media (2012)
4. Fong, J., Tan, Y., Crocher, V., Oetomo, D., Mareels, I.: Dual-loop iterative optimal control for the finite horizon LQR problem with unknown dynamics. *Systems & Control Letters* **111**, 49–57 (2018)
5. Frueh, J.A., Phan, M.Q.: Linear quadratic optimal learning control (LQL). *International Journal of Control* **73**(10), 832–839 (2000)
6. Ghavamzadeh, M., Mannor, S., Pineau, J., Tamar, A., et al.: Bayesian Reinforcement Learning: A Survey. *Foundations and Trends® in Machine Learning* **8**(5-6), 359–483 (2015)
7. Li, N., Kolmanovsky, I., Girard, A.: LQ control of unknown discrete-time linear systems—A novel approach and a comparison study. *Optimal Control Applications and Methods* **40**(2), 265–291 (2019)
8. Murray, R., Palladino, M.: A model for system uncertainty in reinforcement learning. *Systems & Control Letters* **122**, 24–31 (2018)
9. Pang, B., Bian, T., Jiang, Z.P.: Adaptive dynamic programming for finite-horizon optimal control of linear time-varying discrete-time systems. *Control Theory and Technology* **17**(1), 73–84 (2019)

10. Pesare, A., Palladino, M., Falcone, M.: Convergence of the value function in optimal control problems with unknown dynamics. In: 2021 European Control Conference (ECC). IEEE (2021)
11. Pesare, A., Palladino, M., Falcone, M.: Convergence results for an averaged lqr problem with applications to reinforcement learning. *Mathematics of Control, Signals, and Systems* pp. 1–33 (2021)
12. Rasmussen, C., Williams, C.: *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA, USA (2006)
13. Strikwerda, J.C.: *Finite difference schemes and partial differential equations*. SIAM (2004)
14. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, second edn. (2018)
15. Sutton, R.S., Barto, A.G., Williams, R.J.: Reinforcement Learning is Direct Adaptive Optimal Control. *IEEE Control Systems* **12**(2), 19–22 (1992)