

On the Effectiveness of Iterative Learning Control

Anirudh Vemula

Robotics Institute, Carnegie Mellon University

VEMULA@CMU.EDU

Wen Sun

Department of Computer Science, Cornell University

WS455@CORNELL.EDU

Maxim Likhachev

Robotics Institute, Carnegie Mellon University

MAXIM@CS.CMU.EDU

J. Andrew Bagnell

Aurora Innovation

DBAGNELL@RI.CMU.EDU

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Iterative learning control (ILC) is a powerful technique for high performance tracking in the presence of modeling errors for optimal control applications. There is extensive prior work showing its empirical effectiveness in applications such as chemical reactors, industrial robots and quadcopters. However, there is little prior theoretical work that explains the effectiveness of ILC even in the presence of large modeling errors, where optimal control methods using the misspecified model (MM) often perform poorly. Our work presents such a theoretical study of the performance of both ILC and MM on Linear Quadratic Regulator (LQR) problems with unknown transition dynamics. We show that the suboptimality gap, as measured with respect to the optimal LQR controller, for ILC is lower than that for MM by higher order terms that become significant in the regime of high modeling errors. A key part of our analysis is the perturbation bounds for the discrete Riccati equation in the finite horizon setting, where the solution is not a fixed point and requires tracking the error using recursive bounds. We back our theoretical findings with empirical experiments on a toy linear dynamical system with an approximate model, a nonlinear inverted pendulum system with misspecified mass, and a nonlinear planar quadrotor system in the presence of wind. Experiments show that ILC outperforms MM significantly, in terms of the cost of computed trajectories, when modeling errors are high.¹

Keywords: Iterative Learning Control, Riccati Perturbation Bounds, Linear Quadratic Control

1. Introduction

Iterative learning control (ILC) has seen widespread adoption in a range of control applications where the dynamics of the system are subject to unknown disturbances or in instances where model parameters are misspecified [Moore et al. \(1992\)](#). While traditional feedback-based control methods have been successful at tackling non-repetitive noise, ILC has shown itself to be effective at adjusting to repetitive disturbance through feedforward control adjustment [Arimoto et al. \(1984\)](#). This was shown empirically in several robotic applications such as manipulation [Kuc et al. \(1991\)](#), and quadcopter trajectory tracking [Schoellig et al. \(2012\)](#); [Mueller et al. \(2012\)](#) among others. Prior work [An et al. \(1988\)](#) uses fixed point theory to analyze the conditions for convergence of ILC but

1. Full version of the paper, including the appendix, can be found at <https://arxiv.org/abs/2111.09434>

does not present performance bounds at convergence. Very recent work [Agarwal et al. \(2021\)](#) presented a ILC algorithm that is robust to model mismatch and uncertainty. However, they analyze the algorithm using planning regret, which measures regret with respect to the best open loop plan in hindsight, and do not study how the performance depends on modeling error. Our work contributes to understanding the effectiveness of ILC by studying its worst case performance, as a function of modeling error, in the linear quadratic regulator (LQR) setting with unknown transition dynamics and access to an approximate model of the dynamics.

A simple approach to the LQR problem with an approximate model of the dynamics is to do optimal control using the misspecified model (MM.) The resulting controller is similar to the certainty equivalent controller obtained by performing optimal control on estimated parameters of the regulator and ignoring the uncertainty of the estimates in adaptive control [Åström and Wittenmark \(2013\)](#). Despite the simplicity of MM, it is challenging to quantify its suboptimality, with respect to the optimal LQR controller, as a result of the modeling errors in the approximate model.

Our first contribution is proving worst case cost suboptimality bounds for MM in the finite horizon LQR setting in terms of the modeling error. This requires us to depart from the fixed point analysis used in prior work [Mania et al. \(2019\)](#); [Konstantinov et al. \(1993\)](#), as the solution to the discrete Riccati equation in the finite horizon is not a fixed point. A key part of our analysis is establishing perturbation bounds by carefully tracking the effect of modeling error through the horizon of the control task. This allows us to quantify the worst case suboptimality gap of MM in the finite horizon LQR setting.

The second contribution is to utilize the same proof techniques as we used for MM to analyze the suboptimality gap of ILC. This allows us to explicitly compare the worst case performance of ILC and MM for LQR problems, and understand why ILC works well in the regime of large modeling errors when MM often performs poorly. Our analysis highlights that the suboptimality gap for ILC is lower than that for MM by higher order terms that can become significant when modeling errors are high. We also show that ILC is capable of keeping the system stable and cost from blowing up even in the presence of large modeling errors, which MM is incapable of. By interpreting the worst case bounds, we identify several linear systems with key characteristics that enable ILC to be robust to large model misspecifications, unlike MM.

The final contribution of this work is to present simple empirical experiments involving optimal control tasks with linear and nonlinear dynamical systems that back the theoretical findings from our analysis. The experiment results reinforce our finding that in the regime of large modeling errors, ILC performs better than MM and synthesizes control inputs that result in smaller suboptimality gaps.

2. Problem Setup

We consider the finite horizon linear quadratic regulator (LQR) setting with a horizon H and a fixed initial state $x_0 \in \mathbb{R}^n$. The dynamics of the system are described by unknown matrices $A_t \in \mathbb{R}^{n \times n}$ and $B_t \in \mathbb{R}^{n \times d}$ for $t = 0, \dots, H-1$ as follows: $x_{t+1} = A_t x_t + B_t u_t$ where $u_t \in \mathbb{R}^d$ is the control input at time step t . Any sequence of control inputs (u_0, \dots, u_{H-1}) results in a state trajectory (x_0, \dots, x_H) . The cost function is defined using matrices $Q \in \mathbb{R}^{n \times n}$, $Q_f \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{d \times d}$ as follows:

$$V_0(x_0) = \sum_{t=0}^{H-1} x_t^T Q x_t + u_t^T R u_t + x_H^T Q_f x_H \quad (1)$$

From optimal control literature [Anderson and Moore \(2007\)](#), we know that the above cost is minimized by a linear time-varying state-feedback controller $K^* = (K_0^*, \dots, K_{H-1}^*)$ with control inputs $u_t = K_t^* x_t$ satisfying:

$$\begin{aligned} K_t^* &= -(R + B_t^T P_{t+1}^* B_t)^{-1} B_t^T P_{t+1}^* A_t \\ P_t^* &= Q + A_t^T P_{t+1}^* (I + B_t R^{-1} B_t^T P_{t+1}^*)^{-1} A_t \end{aligned}$$

where we initialize $P_H^* = Q_f$ and the matrices P_t^* define the optimal cost-to-go incurred using the optimal controller K^* from time step t as $V_t^*(x_t) = x_t^T P_t^* x_t$.

We are given access to an approximate model of the dynamics of the system specified by matrices $\hat{A}_t \in \mathbb{R}^{n \times n}$ and $\hat{B}_t \in \mathbb{R}^{n \times d}$ for $t = 0, \dots, H-1$ such that there exists some $\epsilon_A, \epsilon_B \geq 0$ (also referred to as the modeling error) satisfying $\|A_t - \hat{A}_t\| \leq \epsilon_A$ and $\|B_t - \hat{B}_t\| \leq \epsilon_B$. For the purposes of this paper, we use the notation $\|\cdot\|$ to refer to the matrix norm induced by the L2 vector norm. In this paper, we consider two control strategies: optimal control using the misspecified model (MM) and iterative learning control (ILC.)

2.1. Optimal Control using Misspecified Model

Optimal control using misspecified model uses the approximate model to synthesize a time-varying linear controller $K^{\text{MM}} = (K_0^{\text{MM}}, \dots, K_{H-1}^{\text{MM}})$ satisfying:

$$\begin{aligned} K_t^{\text{MM}} &= -(R + \hat{B}_t^T P_{t+1}^{\text{MM}} \hat{B}_t)^{-1} \hat{B}_t^T P_{t+1}^{\text{MM}} \hat{A}_t \\ P_t^{\text{MM}} &= Q + \hat{A}_t^T P_{t+1}^{\text{MM}} (I + \hat{B}_t R^{-1} \hat{B}_t^T P_{t+1}^{\text{MM}})^{-1} \hat{A}_t \end{aligned}$$

where we initialize $P_H^{\text{MM}} = Q_f$ and the control inputs are defined as $u_t^{\text{CE}} = K_t^{\text{MM}} x_t$. One can observe that the controller K^{MM} results in suboptimal cost when executed in the system as it is optimizing the cost under approximate dynamics rather than the true dynamics of the system. Thus, the suboptimality gap $V_0^{\text{MM}}(x_0) - V_0^*(x_0)$ depends on the approximate dynamics \hat{A}_t, \hat{B}_t , and how well they approximate the true dynamics.

2.2. Iterative Learning Control

Iterative learning control [Arimoto et al. \(1984\)](#); [Moore et al. \(1992\)](#) is a framework that is used to efficiently calculate the feedforward input signal adjustment by using information from previous trials to improve the performance in a small number of iterations. An example of an ILC algorithm is shown in [Algorithm 1](#). ILC assumes a rollout access to the system, i.e. we are allowed to conduct full rollouts of horizon H in the system to evaluate the cost and obtain the trajectory under true dynamics ([Line 4](#)). Note that this access is only restricted to rollouts, and the true dynamics A_t, B_t are unknown. ILC can be understood as an iterative shooting method where we synthesize control inputs by always evaluating in the true system while computing updates to the controls using the approximate model [Abbeel et al. \(2006\)](#); [Agarwal et al. \(2021\)](#). In [Algorithm 1](#), this is achieved by linearizing the dynamics and quadraticizing the cost around the observed trajectory ([Line 5](#)) resulting in an LQR problem with the objective:

$$J(\Delta x, \Delta u) = \sum_{t=0}^{H-1} (2x_t + \Delta x_t)^T Q \Delta x_t + (2u_t + \Delta u_t)^T R \Delta u_t + (2x_H + \Delta x_H)^T Q_f \Delta x_H \quad (2)$$

Algorithm 1 ILC Algorithm for Linear Dynamical System with Approximate Model

- 1: **Input:** Approximate model \hat{A}_t, \hat{B}_t , Initial state x_0 , Step size α , cost matrix Q, R, Q_f
 - 2: Initialize a control sequence $u_{0:H-1}$ using approximate model
 - 3: **while** not converged **do**
 - 4: Rollout $u_{0:H-1}$ on the true system to get trajectory $x_{0:H}$
 - 5: Compute LQR solution $\arg \min_{\Delta x, \Delta u} J(\Delta x, \Delta u)$ subject to $\hat{A}_t \Delta x_t + \hat{B}_t \Delta u_t = \Delta x_{t+1}$
 - 6: Update $u_{0:H-1} = u_{0:H-1} + \alpha \Delta u_{0:H-1}$
 - 7: **end while**
-

where $x_{0:H}$ is the observed trajectory on the true system when executing controls $u_{0:H-1}$, and for any $t = 0, \dots, H-1$ we have $\hat{A}_t \Delta x_t + \hat{B}_t \Delta u_t = \Delta x_{t+1}$. To characterize the control inputs ILC converges to, let us consider the gradient of the cost-to-go at time step t , $\nabla_u (x_t^T Q x_t + u_t^T R u_t + x_{t+1}^T P_{t+1}^{\text{ILC}} x_{t+1})$ where P_{t+1}^{ILC} captures the cost-to-go from time step $t+1$, and set it to zero:

$$\begin{aligned} 2Ru_t + \nabla_u x_{t+1} \nabla_x (x_{t+1}^T P_{t+1}^{\text{ILC}} x_{t+1}) &= 0 \\ 2Ru_t + 2\hat{B}_t^T P_{t+1}^{\text{ILC}} x_{t+1} &= 0 \\ 2Ru_t + 2\hat{B}_t^T P_{t+1}^{\text{ILC}} (A_t x_t + B_t u_t) &= 0 \\ u_t &= -(R + \hat{B}_t^T P_{t+1}^{\text{ILC}} B_t)^{-1} \hat{B}_t^T P_{t+1}^{\text{ILC}} A_t x_t \end{aligned}$$

where we use the approximate model to compute the derivative $\nabla_u x_{t+1} = \hat{B}_t^T$ and use the rollout trajectory to obtain $x_{t+1} = A_t x_t + B_t u_t$. It is important to note that converging to these control inputs require carefully choosing appropriate step sizes α at each iteration in the ILC Algorithm 1. Thus, we can see that the control inputs ILC converges to can be described using a time-varying state-feedback linear controller K^{ILC} defined as:

$$\begin{aligned} K_t^{\text{ILC}} &= -(R + \hat{B}_t^T P_{t+1}^{\text{ILC}} B_t)^{-1} \hat{B}_t^T P_{t+1}^{\text{ILC}} A_t \\ P_t^{\text{ILC}} &= Q + \hat{A}_t^T P_{t+1}^{\text{ILC}} (I + B_t R^{-1} \hat{B}_t^T P_{t+1}^{\text{ILC}})^{-1} A_t \end{aligned}$$

where we initialize $P_H^{\text{ILC}} = Q_f$ and the control inputs are defined as $u_t^{\text{ILC}} = K_t^{\text{ILC}} x_t$. We can observe that the ILC converges to control inputs that are different from the ones computed by the optimal controller K^* , and hence achieves suboptimal cost. In the next few sections, we will analyze the suboptimality bounds for both MM and ILC, and show how ILC converges to control sequence that achieves lower costs and is more robust to high modeling errors when compared to MM.

2.3. Assumptions

In this section, we will present all the assumptions used in our analysis. Our first assumption is on the cost matrices Q, Q_f and R , also used in [Mania et al. \(2019\)](#):

Assumption 1 *We assume that Q, Q_f , and R are positive-definite matrices. Note that simply scaling all of Q, Q_f , and R does not change the optimal controller K^* , so we can assume that the smallest singular value of R , $\sigma(R) \geq 1$.*

The above assumption allows us to ignore terms relating to singular values of R in the analysis, keeping it concise. The next assumption states that the true system is stable under the optimal controller K^* . Similar notions of stability have been considered in [Cohen et al. \(2018\)](#):

Assumption 2 We assume that the optimal controller K^* satisfies $\|A_t + B_t K^*\| \leq 1 - \delta$ for some $0 < \delta \leq 1$ and all $t = 0, \dots, H - 1$.

Assumption 3 We assume that the matrix $B_t R^{-1} \hat{B}_t^T$ has eigenvalues that have non-negative real parts for all $t = 0, \dots, H - 1$. A sufficient condition for this to hold is that the modeling error satisfy $\epsilon_B \leq \frac{\sigma(B_t^T R B_t)}{\|B_t^T R\|}$ for all $t = 0, \dots, H - 1$.

The above assumption ensures that $x^T B_t R^{-1} \hat{B}_t^T x \geq 0$ for any vector $x \in \mathbb{R}^n$ and all time steps t . Intuitively, if this is not true then ILC is not guaranteed to converge to a local minima.

3. Main Results

In this section, we will present the main results concerning the worst case performance bounds of MM and ILC in the LQR setting with an approximate model as described in Section 2. Our first theorem (proof in Appendix A) bounds the cost suboptimality of any time-varying linear controller \hat{K} in terms of the norm differences $\|K_t^* - \hat{K}_t\|$:

Theorem 1 Suppose $d \leq n$. Denote $\Gamma = 1 + \max_t \{\|A_t\|, \|B_t\|, \|P_t^*\|, \|K_t^*\|\}$ and \hat{V}_0 is the cost-to-go incurred by controller \hat{K} when executed on the system. Then under Assumption 2 and if $\|K_t^* - \hat{K}_t\| \leq \frac{\delta}{2\|B_t\|}$ for all $t = 0, \dots, H - 1$, we have

$$\hat{V}_0(x_0) - V_0^*(x_0) \leq d\Gamma^3 \|x_0\|^2 \sum_{t=0}^{H-1} e^{-\delta t} \|K_t^* - \hat{K}_t\|^2 \quad (3)$$

This theorem is central to our analysis as it states that as long as we can keep the norm differences $\|\hat{K}_t - K_t^*\|$ small, then the cost suboptimality scales with the norm difference squared at each time step and goes exponentially down with time step. We will now present results on how we can bound these norm differences for both MM and ILC.

Results for Optimal Control with Misspecified Model Our next lemma (proof in Appendix C) bounds the difference $\|K_t^{\text{MM}} - K_t^*\|$ in terms of $\|P_{t+1}^{\text{MM}} - P_{t+1}^*\|$ and modeling errors ϵ_A, ϵ_B :

Lemma 2 If $\|A_t - \hat{A}_t\| \leq \epsilon_A$ and $\|B_t - \hat{B}_t\| \leq \epsilon_B$ for $t = 0, \dots, H - 1$, and we have $\|P_{t+1}^* - P_{t+1}^{\text{MM}}\| \leq f_{t+1}^{\text{MM}}(\epsilon_A, \epsilon_B)$ for some function f_{t+1}^{MM} . Then we have under Assumption 1 for all $t = 0, \dots, H - 1$,

$$\|K_t^* - K_t^{\text{MM}}\| \leq 14\Gamma^3 \epsilon_t \quad (4)$$

where $\Gamma = 1 + \max_t \{\|A_t\|, \|B_t\|, \|P_t^*\|, \|K_t^*\|\}$ and $\epsilon_t = \max\{\epsilon_A, \epsilon_B, f_{t+1}^{\text{MM}}(\epsilon_A, \epsilon_B)\}$.

This result is very promising but there is a big piece still missing: how do we bound $f_{t+1}^{\text{MM}}(\epsilon_A, \epsilon_B)$. To do this, we need to establish perturbation bounds for the discrete ricatti equation in the finite horizon setting. Prior work [Konstantinov et al. \(1993\)](#); [Mania et al. \(2019\)](#) has only established such bounds in the infinite horizon setting using fixed point analysis. Our treatment is significantly different as the finite horizon solution is not a fixed point. Our final perturbation bounds are presented in the theorem (proof in Appendix C) below:

Theorem 3 *If the cost-to-go matrices for the optimal controller and MM controller are specified by $\{P_t^*\}$ and $\{P_t^{\text{MM}}\}$ such that $P_H^* = P_H^{\text{MM}} = Q_f$ then,*

$$\begin{aligned} \|P_t^* - P_t^{\text{MM}}\| &\leq \|A_t\|^2 \|P_{t+1}^*\|^2 (2\|B_t\| \|R^{-1}\| \epsilon_B + \|R^{-1}\| \epsilon_B^2) \\ &\quad + 2\|A_t\| \|P_{t+1}^*\| \epsilon_A + \|P_{t+1}^*\| \epsilon_A^2 \\ &\quad + c_{P_{t+1}^*} (\|A_t\| + \epsilon_A)^2 \|P_{t+1}^* - P_{t+1}^{\text{MM}}\| \end{aligned} \quad (5)$$

for $t = 0, \dots, H-1$ where $c_{P_{t+1}^*} \in \mathbb{R}^+$ is a constant that is dependent only on P_{t+1}^* if ϵ_A, ϵ_B are small enough such that $\|P_{t+1}^* - P_{t+1}^{\text{MM}}\| \leq \|P_{t+1}^*\|^{-1}$. Furthermore, the upper bound (5) is tight up to constants that only depend on the true dynamics A_t, B_t , cost matrix R , and P_{t+1}^* .

The above theorem gives us an upper bound for f_t^{MM} for $t = 0, \dots, H-1$ in Lemma 2 with $f_H^{\text{MM}} = 0$. The resulting upper bound on $\|K_t^* - K_t^{\text{MM}}\|$ from Lemma 2 combined with Theorem 1 gives us the cost suboptimality bound for MM. Notice that the bound on f_t^{MM} grows quickly as t decreases making f_{t+1}^{MM} in Lemma 2 the dominant error term that affects the cost suboptimality of MM.

Results for Iterative Learning Control Our final set of results establish similar worst case cost suboptimality bounds for ILC by first establishing a bound (proof in Appendix E) on the difference $\|K_t^{\text{ILC}} - K_t^*\|$ in terms of $\|P_{t+1}^{\text{ILC}} - P_{t+1}^*\|$ and modeling error ϵ_A, ϵ_B :

Lemma 4 *If $\|A_t - \hat{A}_t\| \leq \epsilon_A$ and $\|B_t - \hat{B}_t\| \leq \epsilon_B$ for $t = 0, \dots, H-1$, and we have $\|P_{t+1} - P_{t+1}^{\text{ILC}}\| \leq f_{t+1}^{\text{ILC}}(\epsilon_A, \epsilon_B)$ for some function f_{t+1}^{ILC} . Then we have under Assumption 1 for all $t = 0, \dots, H-1$,*

$$\|K_t^* - K_t^{\text{ILC}}\| \leq 6\Gamma^3 \epsilon_t \quad (6)$$

where $\Gamma = 1 + \max_t\{\|A_t\|, \|B_t\|, \|P_t^*\|, \|K_t^*\|\}$ and $\epsilon_t = \max\{\epsilon_A, \epsilon_B, f_{t+1}^{\text{ILC}}(\epsilon_A, \epsilon_B)\}$.

Similar to MM, we need to bound the crucial term $f_{t+1}^{\text{ILC}}(\epsilon_A, \epsilon_B)$ to bound the norm difference $\|K_t^* - K_t^{\text{ILC}}\|$ using Lemma 4. We will present perturbation bounds (proof in Appendix E) for the ILC recursion equation given in Section 2.2 in the finite horizon setting below:

Theorem 5 *If the cost-to-go matrices for the optimal controller and iterative learning control are specified by $\{P_t^*\}$ and $\{P_t^{\text{ILC}}\}$ such that $P_H^* = P_H^{\text{ILC}} = Q_f$ then we have under Assumption 3,*

$$\begin{aligned} \|P_t^* - P_t^{\text{ILC}}\| &\leq \|A_t\|^2 \|P_{t+1}^*\|^2 \|B_t\| \|R^{-1}\| \epsilon_B + \|A_t\| \|P_{t+1}^*\| \epsilon_A \\ &\quad + c_{P_{t+1}^*} \|A_t\| (\|A_t\| + \epsilon_A) \|P_{t+1}^* - P_{t+1}^{\text{ILC}}\| \end{aligned} \quad (7)$$

for $t = 0, \dots, H-1$ where $c_{P_{t+1}^*} \in \mathbb{R}^+$ is a constant that is dependent only on P_{t+1}^* if ϵ_A, ϵ_B are small enough that $\|P_{t+1}^* - P_{t+1}^{\text{ILC}}\| \leq \|P_{t+1}^*\|^{-1}$. Furthermore, the upper bound (7) is tight upto constants that depend only on the true dynamics A_t, B_t , cost matrix R , and P_{t+1}^* .

The above theorem gives us a bound on f_t^{ILC} for $t = 0, \dots, H-1$ in Lemma 4 with $f_H^{\text{ILC}} = 0$. The resulting upper bound on $\|K_t^* - K_t^{\text{ILC}}\|$ from Lemma 4 combined with Theorem 1 gives us the cost suboptimality bound for iterative learning control. Similar to MM, the dominant error term in Lemma 4 turns out to be f_{t+1}^{ILC} especially for smaller t as the upper bound (7) grows quickly as t decreases.

4. Interpreting the Worst Case Bounds

The recursive bounds presented in (5) and (7) make it difficult to compute a concise bound in Theorem 1. In this section, we will explicitly compare the cost suboptimality bounds for MM and ILC under different scenarios, where the bound can be simplified. Furthermore, it is important to note that in obtaining the controller K^{ILC} we have used data in the form of the rollouts performed on the system that MM did not have access to. An ideal comparison would be to compare the performance of ILC against MM that uses a model identified using the same data from rollouts, but such a study is out of scope for this work.

Small Modeling Errors In the regime of small modeling errors $\epsilon_A \ll 1$ and $\epsilon_B \ll 1$, we can ignore quadratic terms ϵ_A^2 and ϵ_B^2 in upper bound for MM (5) which results in an upper bound that matches that of ILC (7) upto a constant. This suggests that when the modeling errors are small, both ILC and MM have almost the same worst case performance, with ILC having better performance over MM by a constant factor. Intuitively, this makes sense as the approximate model is a very good approximation of the true dynamics, and despite using only the model, MM can synthesize a near-optimal controller.

Highly Damped Systems The second scenario we consider is that of a system that is highly damped which implies $\|A_t\| \ll 1$ for all $t = 0, \dots, H - 1$. In this regime, the upper bound for ILC (7) goes down to zero resulting in ILC achieving near-optimal cost despite having non-zero modeling errors ϵ_A, ϵ_B . The suboptimality in ILC (from Lemma 4) only arises from ϵ_A, ϵ_B and not from f_{t+1}^{ILC} which is 0. In contrast, the upper bound for MM (5) does not go down to zero and has terms that depend on ϵ_A^2 , which can be significant when ϵ_A is not small. Thus, for highly damped systems we have that the worst case performance of ILC can be significantly better than MM, especially when ϵ_A is large. Intuitively, this can be understood by observing that ILC removes the effect of modeling errors by always performing rollouts using true dynamics, while MM errors are exacerbated by using the approximate model for rollouts. Interestingly, we also notice that the modeling error ϵ_B does not affect the cost-suboptimality in upper bound for MM (5) when the system is highly damped.

Weakly Controlled Systems For systems with small $\|B_t\| \ll 1$, i.e. where the control inputs do not affect the dynamics of the system to a large extent, we can observe that the upper bound for ILC (7) reduces to a bound that does not depend ϵ_B . In other words, any modeling error ϵ_B in estimating the B_t matrices does not affect the upper bound (7) for ILC. In contrast, the upper bound for MM (5) reduces to an expression that has terms that depend on ϵ_B^2 , which can become significant when ϵ_B is large. Thus, for systems with $\|B_t\| \ll 1$, ILC is robust to any modeling errors ϵ_B in the B_t matrices, whereas MM degrades its worst case performance with increasing ϵ_B .

Modeling Error only at the first time step Consider a scenario where the model is inaccurate only at $t = 0$, i.e. $\|A_0 - \hat{A}_0\| \leq \epsilon_A$ and $\|B_0 - \hat{B}_0\| \leq \epsilon_B$, while $\hat{A}_t = A_t$ and $\hat{B}_t = B_t$ for all $t = 1, \dots, H - 1$. In this case, the upper bounds (5) and (7) simplify greatly as $\|P_t^* - P_t^{\text{ILC}}\| = \|P_t^* - P_t^{\text{MM}}\| = 0$ for all $t = 1, \dots, H - 1$, and we only have upper bounds on $\|P_0^* - P_0^{\text{MM}}\|$ and $\|P_0^* - P_0^{\text{ILC}}\|$ as given by Theorems 3 and 5 which when combined with Theorem 1 gives us the

suboptimality bounds:

$$\hat{V}_0^{\text{MM}}(x_0) - V_0^*(x_0) \leq \mathcal{O}(1)d\Gamma^9\|x_0\|^2(\epsilon_A + \epsilon_A^2 + \epsilon_B + \epsilon_B^2)^2 \quad (8)$$

$$\hat{V}_0^{\text{ILC}}(x_0) - V_0^*(x_0) \leq \mathcal{O}(1)d\Gamma^9\|x_0\|^2(\epsilon_A + \epsilon_B)^2 \quad (9)$$

The above two cost suboptimality bounds highlight the differences between MM and ILC in worst case performance. As described in Section 4, if ϵ_A and ϵ_B are small, then MM and ILC worst case performances match up to constants as we can ignore higher order terms. However, in cases where modeling errors ϵ_A and ϵ_B are large and higher order terms like $\epsilon_A^2\epsilon_B$, ϵ_A^4 etc. start becoming significant, the worst case performance of ILC tends to be better than MM as indicated by equations (8) and (9). Furthermore, the conditions for stability under synthesized control inputs, as stated in Theorem 1 (and in Lemma 6 in Appendix A,) is harder to satisfy for MM when compared to ILC, especially when modeling errors are large.

5. Empirical Results

In this section, we present three empirical experiments: a linear dynamical system with an approximate model, a nonlinear inverted pendulum system with misspecified mass, and a nonlinear planar quadrotor system in the presence of wind. The aim of these experiments is to show that under high modeling errors, ILC is more efficient than MM, thus backing our theoretical findings.²

5.1. Linear Dynamical System with Approximate Model

In this experiment, we use a linear dynamical system with states $x \in \mathbb{R}^2$ and control inputs $u \in \mathbb{R}$. The dynamics of the system are specified by matrices: $A_t = \begin{bmatrix} 1 & 1 \\ -3 & 1 \end{bmatrix}$, $B_t = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$. The approximate model we use is constructed by perturbing the dynamics as follows: $\hat{A}_t = A_t + \epsilon\mathbb{I}$, $\hat{B}_t = B_t + \epsilon \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ for any $\epsilon \geq 0$. Observe that this satisfies $\|\hat{A}_t - A_t\| \leq \epsilon$ and $\|\hat{B}_t - B_t\| \leq \epsilon$. We use a quadratic cost as specified in equation 1 with matrices: $Q = Q_f = \mathbb{I}$, $R = 1$ (more details in Appendix G.1). We can solve for the optimal controller K^* in closed form using true dynamics A_t, B_t as specified in Section 2. We compare MM controller K^{MM} and iterative learning controller K^{ILC} with approximate model \hat{A}_t, \hat{B}_t in Figure 1(a) where we vary ϵ along the X-axis (in log scale) and report the cost suboptimality gap $V_0(x_0) - V_0^*(x_0)$ on the Y-axis (in log scale) where $V_0(x_0)$ is the cost incurred by K^{MM} or K^{ILC} . To ensure that Assumption 3 is not violated, the X-axis is capped at $\epsilon = \frac{\sigma(B_t^T R B_t)}{\|B_t^T R\|}$. It is important to note that to generate the plot in Figure 1(a) we directly used the closed form solution for K^{ILC} (as described in Section 2) and did not run an iterative learning control algorithm. This was done to ensure that our results do not have any dependence on how well the step size sequence was tuned for ILC.

We can observe that for small modeling errors $\epsilon < 10^{-1}$, ILC outperforms MM by a constant factor (about 4×10^2) as evidenced by the linear trend in log scale. However in the regime of high modeling errors $\epsilon > 10^{-1}$ we observe that the gap between ILC and MM is not a constant factor anymore and grows very quickly as ϵ increases. This can be explained by the fact that for high ϵ , the higher order terms in the gap between ILC and MM starts becoming significant and results in poor

2. The code for all experiments can be found at <https://github.com/vvanirudh/ILC.jl>.




Figure 1: (a) Cost suboptimality gap with varying modeling error ϵ for a linear dynamical system. Note that both X-axis and Y-axis are in log scale. (b) Cost suboptimality gap with varying mass misspecification Δm for a nonlinear inverted pendulum system. (c) Cost suboptimality gap for planar quadrotor control with varying magnitude of wind η .

performance for MM when compared to ILC. For large epsilons, we also observe that the cost for MM blows up to really big values as the system is not stable anymore under K^{MM} due to violation of the condition in Theorem 1 (and in Lemma 6 in Appendix A.) This experiment validates our claim from the analysis that ILC tends to perform better in terms of cost and is more robust when modeling errors are high.

5.2. Nonlinear Inverted Pendulum with Misspecified Mass

For the second experiment, we use the nonlinear dynamical system of an inverted pendulum. The state space is specified by $x = [\theta \ \dot{\theta}] \in \mathbb{R}^2$ where θ is the angle between the pendulum and the vertical axis. The control input is $u = \tau \in \mathbb{R}$ specifying the torque τ to be applied at the base of the pendulum. The dynamics of the system are given by the ODE, $\ddot{\theta} = \frac{\bar{\tau}}{m\ell^2} - \frac{g \sin(\theta)}{\ell}$ where m is the mass of the pendulum, ℓ is the length of the pendulum, g is the acceleration due to gravity, and $\bar{\tau} = \max(\tau_{\min}, \min(\tau_{\max}, \tau))$ is the clipped torque based on torque limits (more details in Appendix G.2). We use an approximate model of the dynamics where the mass of the pendulum is perturbed as $\hat{m} = m + \Delta m$. This results in dynamics that are nonlinearly perturbed from the true dynamics. Since the dynamics are nonlinear, we cannot obtain optimal controls, and MM controls in closed form. Instead, we approximate these controllers by running iLQR Li and Todorov (2004) (both forward and backward pass) on the true dynamics and the approximate dynamics (obtained by linearizing the true dynamics) respectively for 200 iterations. To obtain ILC control inputs, we run iLQR with forward pass (or rollouts) using the true dynamics, and backward pass computed using the approximate linearized dynamics at each iteration. We chose step sizes for all iLQR runs using backtracking line search.

Figure 1(b) shows the cost suboptimality gap of MM and ILC as the perturbation Δm varies. Similar to our previous experiment, we observe that for small modeling errors $\Delta m < 0.07$ both ILC and MM perform similarly with ILC outperforming slightly. But as Δm grows, the cost of MM quickly grows saturating at a suboptimality gap around 57. In contrast, we observe that ILC is still able to compute near-optimal controls until $\Delta m = 0.15$ showcasing the robustness of ILC to higher modeling errors. Beyond $\Delta m = 0.15$, ILC performance also degrades significantly as the approximate model is not representative of the true dynamics anymore. Although our analysis in the previous sections was restricted to linear dynamical systems, we notice a similar trend between

ILC and MM in the presence of nonlinear dynamics namely, in the regime of large modeling errors, ILC tends to perform better than MM.

5.3. Nonlinear Planar Quadrotor Control in Wind

In our final experiment, we compare MM and ILC on a planar quadrotor control task in the presence of wind. A similar setting was used in [Agarwal et al. \(2021\)](#). The quadrotor is controlled using two propellers that provide upward thrusts (u_1, u_2) and allows movement in the 3D planar space described as (p_x, p_y, θ) where p_x, p_y are X, Y positions, and θ is the yaw of the quadrotor. The dynamics of the planar quadrotor is specified using a state vector $x \in \mathbb{R}^6$, and control input $u \in \mathbb{R}^2$ (more details in Appendix G.3). The quadrotor is flying in the presence of wind which is not captured in modeled dynamics, but affects the true dynamics of the quadrotor as a dispersive force field $(\eta p_x \mathbf{i} + \eta p_y \mathbf{j})$ resulting in overall dynamics given by: $\ddot{p}_x = \frac{1}{m}(u_1 + u_2) \sin(\theta) + \eta p_x$, $\ddot{p}_y = \frac{1}{m}(u_1 + u_2) \cos(\theta) - g + \eta p_y$ where $\eta \in \mathbb{R}^+$ is the magnitude of the wind force field.

The objective of the task is to move the quadrotor from an initial state x_0 to a final state x_f . Similar to previous experiment, the dynamics are nonlinear and we cannot obtain optimal controls and MM controls in closed form. We compute MM and ILC controllers similar to Section 5.2, and we initialize control inputs as the hover controls. Figure 1(c) compares MM and ILC for planar quadratic control with varying magnitude of wind η . For small wind magnitudes, we observe that both MM and ILC have good performance. As the wind magnitude increases, MM quickly diverges and the cost of synthesized control inputs blows up quickly as the modeled dynamics are incapable of capturing the dispersive force field exerted by the wind. ILC, on the other hand, manages to keep the cost from blowing up even at large wind magnitudes. This reinforces our conclusion that ILC is robust to large modeling errors while MM can quickly result in the cost blowing up when the model is highly inaccurate.

6. Discussion

Our analysis shows that the gap between ILC and MM is in higher order terms that can become significant when the modeling error ϵ_A, ϵ_B is large. This is backed by our empirical experiments where we observe that as the magnitude of modeling error increases, the performance gap between ILC and MM grows rapidly as MM is incapable of handling large modeling errors and the resulting cost diverges. Furthermore, the conditions needed for stability of the system under synthesized control inputs, are easier to satisfy for ILC when compared to MM, especially in the regime of large modeling errors. This explains the robustness of ILC over MM for complex control tasks when given access to highly inaccurate dynamical models.

While our current analysis is restricted to the linear quadratic control setting, exploring similar suboptimality bounds in more complex and possibly, nonlinear settings is an exciting direction for future work. Recent work by [Simchowit and Foster \(2020\)](#) uses a self-bounding ODE method to establish perturbation bounds that sharpens previous bounds in the infinite horizon setting by only depending on natural control-theoretic quantities and not relying on controllability assumptions. It remains to be seen if we can rely on similar techniques to sharpen the bounds presented in this work. Finally, comparing iterative learning control and robust control approaches such as [Dean et al. \(2020\)](#) would allow us to understand the regime of modeling errors in which ILC is more suitable than robust control approaches, and vice versa.

Acknowledgments

AV would like to thank Horia Mania for helping with understanding some technical lemmas in [Mania et al. \(2019\)](#). AV is supported by the CMU Presidential Fellowship endowed by TCS.

References

- Pieter Abbeel, Morgan Quigley, and Andrew Y. Ng. Using inaccurate models in reinforcement learning. In William W. Cohen and Andrew W. Moore, editors, *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 1–8. ACM, 2006. doi: 10.1145/1143844.1143845. URL <https://doi.org/10.1145/1143844.1143845>.
- Naman Agarwal, Elad Hazan, Anirudha Majumdar, and Karan Singh. A regret minimization approach to iterative learning control. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 100–109. PMLR, 2021. URL <http://proceedings.mlr.press/v139/agarwal21b.html>.
- Chae H An, Christopher G Atkeson, and John M Hollerbach. *Model-based control of a robot manipulator*. MIT press, 1988.
- Brian DO Anderson and John B Moore. *Optimal control: linear quadratic methods*. Courier Corporation, 2007.
- Suguru Arimoto, Sadao Kawamura, and Fumio Miyazaki. Bettering operation of robots by learning. *J. Field Robotics*, 1(2):123–140, 1984. doi: 10.1002/rob.4620010203. URL <https://doi.org/10.1002/rob.4620010203>.
- Karl J Åström and Björn Wittenmark. *Adaptive control*. Courier Corporation, 2013.
- Alon Cohen, Avinandan Hassidim, Tomer Koren, Nevena Lazic, Yishay Mansour, and Kunal Talwar. Online linear quadratic control. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 1028–1037. PMLR, 2018. URL <http://proceedings.mlr.press/v80/cohen18b.html>.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Found. Comput. Math.*, 20(4):633–679, 2020. doi: 10.1007/s10208-019-09426-y. URL <https://doi.org/10.1007/s10208-019-09426-y>.
- Michail M. Konstantinov, Petko Hr. Petkov, and Nicolai Christov. Perturbation analysis of the discrete riccati equation. *Kybernetika*, 29(1):18–29, 1993. URL <http://www.kybernetika.cz/content/1993/1/18>.
- Tae-Yong Kuc, Kwanghee Nam, and Jin S. Lee. An iterative learning control of robot manipulators. *IEEE Trans. Robotics Autom.*, 7(6):835–842, 1991. doi: 10.1109/70.105392. URL <https://doi.org/10.1109/70.105392>.

- Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In Helder Araújo, Alves Vieira, José Braz, Bruno Encarnação, and Marina Carvalho, editors, *ICINCO 2004, Proceedings of the First International Conference on Informatics in Control, Automation and Robotics, Setúbal, Portugal, August 25-28, 2004*, pages 222–229. INSTICC Press, 2004.
- Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalent control of LQR is efficient. *CoRR*, abs/1902.07826, 2019. URL <http://arxiv.org/abs/1902.07826>.
- Kevin L. Moore, Mohammed Dahleh, and S. P. Bhattacharyya. Iterative learning control: A survey and new results. *J. Field Robotics*, 9(5):563–594, 1992. doi: 10.1002/rob.4620090502. URL <https://doi.org/10.1002/rob.4620090502>.
- Fabian L. Mueller, Angela P. Schoellig, and Raffaello D’Andrea. Iterative learning of feed-forward corrections for high-performance tracking. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*, pages 3276–3281. IEEE, 2012. doi: 10.1109/IROS.2012.6385647. URL <https://doi.org/10.1109/IROS.2012.6385647>.
- Angela P. Schoellig, Fabian L. Mueller, and Raffaello D’Andrea. Optimization-based iterative learning for precise quadcopter trajectory tracking. *Auton. Robots*, 33(1-2):103–127, 2012. doi: 10.1007/s10514-012-9283-2. URL <https://doi.org/10.1007/s10514-012-9283-2>.
- Max Simchowitz and Dylan J. Foster. Naive exploration is optimal for online LQR. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 8937–8948. PMLR, 2020. URL <http://proceedings.mlr.press/v119/simchowitz20a.html>.