

The landscape of deterministic and stochastic optimal control problems: One-shot Optimization versus Dynamic Programming

Jihun Kim, Yuhao Ding, Yingjie Bi, and Javad Lavaei, *Fellow, IEEE*

Abstract—Optimal control problems can be solved via a one-shot (single) optimization or a sequence of optimization using dynamic programming (DP). However, the computation of their global optima often faces NP-hardness, and thus only locally optimal solutions may be obtained at best. In this work, we consider the discrete-time finite-horizon optimal control problem in both deterministic and stochastic cases and study the optimization landscapes associated with two different approaches: one-shot and DP. In the deterministic case, we prove that each local minimizer of the one-shot optimization corresponds to some control input induced by a locally minimum control policy of DP, and vice versa. However, with a parameterized policy approach, we prove that deterministic and stochastic cases both exhibit the desirable property that each local minimizer of DP corresponds to some local minimizer of the one-shot optimization, but the converse does not necessarily hold. Nonetheless, under different technical assumptions for deterministic and stochastic cases, if there exists only a single locally minimum control policy, one-shot and DP turn out to capture the same local solution. These results pave the way to understand the performance and stability of local search methods in optimal control.

Index Terms—Optimal Control, Landscape, One-shot optimization, Dynamic programming

I. INTRODUCTION

Dynamic Programming (DP) has been widely used in a variety of fields with a rich theoretical foundation and many mathematical and algorithmic aspects [2], [3]. One classic area of DP is to solve optimal control problems, with applications in communication systems [4], inventory control [5], power-train control [6], and many more. Furthermore, many recent

This work was supported by grants from ARO, ONR, AFOSR, NSF, and the UC Noyce Initiative.

J. Kim, Y. Ding, Y. Bi, and J. Lavaei are with the Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA 94720 USA (e-mail: jihun.kim@berkeley.edu; yuhao.ding@berkeley.edu; yb236@cornell.edu; lavaei@berkeley.edu).

A preliminary version of this paper has appeared in 2021 American Control Conference, New Orleans, USA, May 25-28, 2021 [1]. The previous version mainly discussed the deterministic problem with control inputs, while this journal version has significantly extended the results to include both the deterministic and the stochastic problems under a parameterized policy to study a closed-loop system. To address the parameterized problem, our new notion of a local minimizer of the one-shot optimization optimizes the objective function over the parameters modeling the inputs. Furthermore, the notion of a locally minimum control policy of DP is replaced with a local minimizer of DP, which considers a neighborhood in the parameter space instead of the action space. These new notions enable the investigation of the stochastic dynamics as a finite-dimensional problem.

successes in artificial intelligence, especially in reinforcement learning (RL) [7], [8], are also deeply rooted in DP. In the challenging domain of classic Atari 2600 games, the work [9] has demonstrated that the Q-learning method based on the generalized policy iteration along with a deep neural network as the function approximator for the Q-values surpasses the performance of all previous algorithms and achieves a level comparable to that of a professional human games tester.

Despite a strong theoretical framework of DP, the exact solutions of large-scale optimal control problems are often impossible to obtain using DP in practice [7]. Apart from suffering the “curse of dimensionality” when the state space is large, solving DP accurately could also be highly complex. The reason is that DP requires solving optimization sub-problems to global optimality, and the computation of their global optima is NP-hard in general, due to the non-linearity of the dynamics and the non-convexity of the cost function.

Therefore, although DP theory relies on global optimization solvers, practitioners routinely use local optimization solvers based on first- and second-order numerical algorithms. As a result, the theoretical guarantee of DP could break down as soon as a non-global local solution is found in any of the sub-problems. Understanding the performance of local search methods for non-convex problems has been a focal area in machine learning in recent years. This is performed under the notion of spurious solution, which refers to a local minimum that is not a global solution. The specific application areas are neural networks [10], [11], deep learning [12], [13], mixtures of regressions [14], [15], matrix sensing/recovery [16]–[19], phase retrieval [15], [20], and online optimization [21], [22].

Recently, there has been an increasing interest in understanding the global convergence of exact or approximate DP algorithms in policy gradient methods for RL, such as projected policy gradient, natural policy gradient, and mirror descent with or without regularizers [23]–[27]. Prior to them, the work [28] identified some general algorithm-independent properties of the policy gradient method by establishing a direct connection between policy gradient (one-shot) and policy iteration (DP) objectives. They showed that the global convergence of the policy gradient method is guaranteed if the policy iteration objectives have no sub-optimal stationary points. However, the literature lacks a rigorous analysis of the spurious solutions of the DP method.

In this paper, we analyze the spurious solutions of the DP method by focusing on the following fundamental question:

What if the globally optimal solution of each sub-problem of DP is replaced with a solution obtained by a local search method? A challenge in this analysis is that policy optimization even towards the spurious solutions can be problematic if the action space is continuous [29]. One can think of the policy iteration with function approximation [30] where the Q-function approximation error is zero. This is a reasonable assumption since a close-to-zero error can be obtained with a sufficiently rich and expressive policy class such as deep neural networks, which naturally yields the existence of the local minimizer of DP. That motivates our analysis on the comparison between the solutions of one-shot method and DP if they are only solved to the spurious local minimizers, and hence, our algorithm-agnostic study offers a clear understanding on the landscapes for the optimal control problem without considering the secondary issue of the approximation error.

We focus on both deterministic and stochastic discrete-time finite-horizon optimal control problems whose goal is to find an optimal input sequence minimizing the total cost subject to the dynamics. One approach to solving the problem is by formulating it as a one-shot optimization problem, a single entire-period problem, and another approach is using the DP to formulate it as a sequential decision-making problem with multiple single-period sub-problems and solve it in a backward way. Although it is known that the one-shot method and the DP method return the same globally optimal control sequence for the deterministic optimal control problem [3], it is not yet known what would occur if the global optimizer needed for solving each sub-optimization problem in DP is replaced by a local optimizer. In our work, we compare the two optimization landscapes: one induced by the DP method based on local search algorithms, and the other induced by its corresponding one-shot optimization based on local search methods.

Contribution and Outline. We address the relationship between the two landscapes holistically for three types of control systems:

1. In Section II, we first study *deterministic* systems under a non-parameterized policy. We introduce the notion of locally minimum control policy of DP and prove that under some mild conditions, each (spurious) local minimizer of the one-shot optimization corresponds to the control input induced by a (spurious) locally minimum control policy of DP, and vice versa. This indicates that DP with local search can successfully solve the optimal control problem to global optimality if and only if the one-shot problem is free of spurious solutions.

2. In Section III, we analyze *deterministic* systems under a *parameterized* policy. The necessity to study this problem arises in RL algorithms, where the control policy used by DP is parameterized by neural networks or other means. Thus, we generalize the results of Section II to optimization with respect to the parameters rather than the control inputs themselves. We prove that each local minimizer of DP corresponds to some local minimizer of the one-shot optimization, whereas its converse may not hold. Moreover, we show that if there exists only a single locally minimum control policy with a specific parameterized policy class, namely a linear combination of independent basis functions, each local minimizer of the one-shot optimization corresponds to a local minimizer of DP.

TABLE I
THEOREMS AND THE CORRESPONDING ASSUMPTIONS WITH RESULTS.

Theorems		Deterministic			Deterministic + Parameterized			Stochastic + Parameterized		
		1	2	3	4	5	6	7	8	9
Convex	action space									
	parameter space									
Policy class	C^1									
	C^2									
	Defined by Definition 13									
	contains a single locally minimum control policy									
Interior policy										
Strict local minimizer										
Continuous Random state										
Large parameter space										
Result	DP to one-shot									
	DP to one-shot (stationarity)									
	one-shot to DP									

3. In Section IV, we extend the result to *stochastic* systems under a *parameterized* policy. The stochasticity brings up the challenge to handle an uncountable number of realizations of random variables. We show that surprisingly a similar relationship in the deterministic parameterized problem holds. For both cases, we conclude that the optimization landscape of the one-shot problem is more complex than its DP counterpart in terms of the number of spurious solutions. This implies that if the one-shot problem has a *low* complexity, so does the DP problem. Another result says that if the DP problem has a *very low* complexity, the same holds for the one-shot problem. In this paper, our notion of “complexity” of an optimization problem is based on *the number of spurious local minima*. For example, convex optimization problems have very low complexity in light of having no spurious solutions. However, problems with an exponential number of spurious solutions are hard to solve [31]. Note that a reformulation of an optimization problem via a change of variables does not normally change the number of local minima, which justifies why the number of spurious solutions can serve as a complexity measure.

Finally, concluding remarks are provided in Section V. Table I summarizes the main results of the paper.

In various applications arising in machine learning and model-free approaches for which the model is unknown and simulations are expensive, DP is the only viable choice compared to the one-shot optimization approach. Hence, it is essential to understand when DP combined with a local search solver works. The results of this paper explain that the success of DP is closely related to the optimization landscape of a single optimization problem. For instance, the success of the DP method highly depends on the number of spurious solutions of the one-shot optimization problem.

Notation. Let \mathbb{R} denote the set of real numbers. We use $B(c; r)$ to denote the open ball centered at c with radius r and use $\bar{B}(c; r)$ to denote the closure of $B(c; r)$. The notation $x \in A \setminus B$ means that x is in the set A but not in the set B . Let $\| \cdot \|$ denote the Euclidean norm and $\| \cdot \|_F$ denote the Frobenius norm. Let $\nabla_x f(x; y)$ denote the gradient of $f(x; y)$ with respect to x and $\nabla_x^2 f(x)$ denote the Hessian of $f(x)$. For the matrix K , $K \succ 0$ means that K is positive definite. The notation C^n means that the function is n -times continuously differentiable. The notation \mathbb{E} denotes the expectation operator.

II. DETERMINISTIC PROBLEM

A. Problem Formulation

Consider a general discrete-time finite-horizon deterministic optimal control problem with time steps:

$$\begin{aligned} \min_{u_0, \dots, u_{n-1} \in A} & \sum_{i=0}^{n-1} c_i(x_i; u_i) + c_n(x_n) \\ \text{s.t.} & x_{i+1} = f_i(x_i; u_i); \quad i = 0; \dots; n-1; \\ & x_0 \text{ is given} \end{aligned} \quad (P1)$$

where $x_i \in \mathbb{R}^N$ is the state at time i and u_i is the control input at time i that is constrained to be in an action space $A \subseteq \mathbb{R}^M$. The state transition is governed by the dynamics $f_i : \mathbb{R}^N \times \mathbb{R}^M \rightarrow \mathbb{R}^N$. Each time instance is associated with a stage cost $c_i : \mathbb{R}^N \times \mathbb{R}^M \rightarrow \mathbb{R}$ or the terminal cost $c_n : \mathbb{R}^N \rightarrow \mathbb{R}$. Given an initial state x_0 , the goal of the optimal control problem is to find an optimal control input $(u_0; \dots; u_{n-1})$ minimizing the sum of the stage costs and the terminal cost.

In this paper, the dynamics f_i and the cost functions c_i are assumed to be at least twice continuously differentiable over $\mathbb{R}^N \times \mathbb{R}^M$, and the action space A is assumed to be compact.

The optimal control problem can be solved by two common approaches. The first approach directly solves (P1) as a one-shot optimization problem that simultaneously solves for all $(u_0; \dots; u_{n-1})$ variables. To simplify the analysis, we eliminate the equality constraints in (P1) via the notation $J_k(x_k; u_k; \dots; u_{n-1})$ defined as the cost-to-go started at the time step k with the initial state x_k and control inputs $(u_k; \dots; u_{n-1})$. In other words,

$$\begin{aligned} C(x) &= c_n(x); \\ C(x; u_k; \dots; u_{n-1}) &= c_k(x; u_k) \\ &\quad + C(f_k(x; u_k); u_{k+1}; \dots; u_{n-1}); \end{aligned}$$

for $k = 0; \dots; n-1$. The one-shot optimization problem (P1) can be equivalently written as

$$\min_{u_0, \dots, u_{n-1} \in A} C(x_0; u_0; \dots; u_{n-1}) \quad (P2)$$

The second approach to solving the optimal control problem is based on DP. Let $J_k(x_k)$ denote the optimal cost-to-go at the time step k with the initial state x_k , i.e.,

$$J_k(x_k) = \min_{u_k, \dots, u_{n-1} \in A} C(x_k; u_k; \dots; u_{n-1})$$

Then, J_k can be computed in a backward fashion from the time step $n-1$ to time 0 through the following recursion:

$$\begin{aligned} J_n(x) &= c_n(x); \\ J_k(x) &= \min_{u \in A} c_k(x; u) + J_{k+1}(f_{k+1}(x; u)); \end{aligned} \quad (P3)$$

for $k = 0; \dots; n-1$. It is worth noting that (P3) yields a set of functions that solve the problem for all initial states, whereas (P1) produces a vector specific to a given x_0 . The optimal cost $J_0(x_0)$ equals the optimal objective value of (P1).

However, due to the non-convexity of the function, it is generally NP-hard to obtain globally optimal solutions of (P3) for all states and at all times. Specifically, when using the DP to solve the optimal control problem (P1), the first step is to compute

for every $x_{n-1} \in \mathbb{R}^N$, which requires solving nonconvex optimization problems if the cost function or the dynamic is nonconvex. Since these intermediate problems are normally solved via local search methods, the best expectation is to obtain a local minimizer u_{n-1} as a function of $x_{n-1} \in \mathbb{R}^N$, denoted by the policy $\pi_{n-1}(x)$. As a result, instead of working with truly optimal cost-to-go functions, one may arrive at a sub-optimal cost-to-go at time $n-1$ as follows:

$$\begin{aligned} J_{n-1}(x_{n-1}) &= c_{n-1}(x_{n-1}; \pi_{n-1}(x_{n-1})) \\ &\quad + c_n(f_n(x_{n-1}; \pi_{n-1}(x_{n-1}))); \end{aligned}$$

which is obtained based on the local minimizer $\pi_{n-1}(x)$. Subsequently, it is required to solve the optimal decision-making problem $\min_{u \in A} c_{n-2}(x_{n-2}; u) + J_{n-1}(f_{n-1}(x_{n-2}; u))$ for every $x_{n-2} \in \mathbb{R}^N$. By repeating this procedure in a backward fashion toward the time step 0, we obtain a group of policy functions π_k and sub-optimal cost-to-go functions J_k for $k = 0; \dots; n-1$. Given the initial state x_0 , let

$$\begin{aligned} u_0 &= \pi_0(x_0); \quad x_1 = f_0(x_0; u_0); \quad u_1 = \pi_1(x_1); \quad x_2 = f_1(x_1; u_1) \\ &\quad \vdots \\ u_{n-1} &= \pi_{n-1}(x_{n-1}); \quad x_n = f_{n-1}(x_{n-1}; u_{n-1}); \end{aligned}$$

be the control inputs and the states induced by the policies $(\pi_0; \dots; \pi_{n-1})$. Then, $(u_0; \dots; u_{n-1})$ is a sub-optimal solution to the original optimal control problem (P1) with the sub-optimal objective value $J_0(x_0)$. This motivates us to define locally minimum control policies based on solving (P3) to local optimality.

Definition 1: Given a control policy $\pi = (\pi_0; \dots; \pi_{n-1})$, the associated Q-function $Q_k(\cdot; \cdot)$ and cost-to-go functions $J_k(\cdot)$ under the policy π are defined in a backward way from the time step $n-1$ to 0 through the following recursion:

$$\begin{aligned} J_n(x) &= c_n(x); \\ Q_k(x; u) &= c_k(x; u) + J_{k+1}(f_{k+1}(x; u)); \quad k = 0; \dots; n-1; \\ J_k(x) &= Q_k(x; \pi_k(x)); \quad k = 0; \dots; n-1; \end{aligned}$$

Definition 2 (local minimizer): A vector $(u_0; \dots; u_{n-1})$ is said to be a local minimizer of the one-shot optimization problem (P2) if there exists $\delta > 0$ such that

$$C(x_0; u_0; \dots; u_{n-1}) \leq C(x_0; \theta_0; \dots; \theta_{n-1})$$

for all $\theta_i \in B(u_i; \delta) \setminus A$ where $i = 0; \dots; n-1$. It is further called a spurious (non-global) local minimizer of the one-shot optimization problem if $C(x_0; u_0; \dots; u_{n-1}) > J_0(x_0)$.

Definition 3 (locally minimum control policy): A control policy $\pi = (\pi_0; \dots; \pi_{n-1})$ is said to be a locally minimum control policy of DP if for all $k \in \{0; \dots; n-1\}$ and for all $x \in \mathbb{R}^N$, the policy $\pi_k(x)$ is a local minimizer of the

$$Q_k(x; \pi_k(x)) \leq Q_k(x; \theta); \quad \forall \theta \in B(\pi_k(x); \delta) \setminus A;$$

such that it is further called a spurious locally minimum control policy if $J_0(x_0) > J_0(x_0)$.

In the following subsections, we will show that in the deterministic problem, both approaches capture the same local solutions under mild assumptions.

B. Local minimizers: From DP to one-shot optimization

It is well-known that the input sequence induced by a globally minimal control policy is a global minimizer of the input of a locally minimum control policy of DP corresponds one-shot problem [3]. In this subsection, we will show that the input sequence induced by a spurious locally minimum control policy of DP also corresponds to a spurious local minimizer of the one-shot problem if some mild conditions are satisfied.

Definition 4: Given a set S and a continuously differentiable function g , a point $s \in S$ is said to be a stationary point of the optimization problem $\min_{s \in S} g(s)$ if

$$r_s g(s) \in N_S(s);$$

where $N_S(s)$ denotes the normal cone of the set S at the point s [32]. We branch off into two specific notions of stationarity below.

Definition 5 (Stationary point): A vector of control inputs (u_0, \dots, u_{n-1}) is said to be a stationary point of the one-shot optimization if for all $k \in \{0, \dots, n-1\}$, it holds that

$$r_{u_k} C(x_0; u_0, \dots, u_{n-1}) \in N_A(u_k).$$

Proof: First, we will use induction to find positive numbers $\epsilon_0, \dots, \epsilon_n$ and $\delta_0, \dots, \delta_{n-1}$ such that

$$r_u^2 Q_k(x; u) \geq 0; \quad (1)$$

$$x_k(x) \in B(u_k; \delta_k); \quad (2)$$

$$f_k(x; u) \in B(x_{k+1}; \delta_{k+1}); \quad (3)$$

for every $x \in B(x_k; \delta_k)$, $u \in B(u_k; \delta_k) \setminus A$, and $k \in \{0, \dots, n-1\}$. At the base step $k = n$, we choose an arbitrary $\epsilon_n > 0$. At the induction step, since f_k is continuous and $r_u^2 Q_k$ is continuous at $(x_k; u_k)$, there exist $\delta_k > 0$ such that both (1) and (3) are satisfied for all $x \in B(x_k; \delta_k)$ and $u \in B(u_k; \delta_k) \setminus A$. Moreover, as f_k is continuous at (x_k, u_k) , (2) will be satisfied by further reducing δ_k .

For every (u_0, \dots, u_{n-1}) with $u_k \in B(u_k; \delta_k) \setminus A$, let (x_0, \dots, x_n) be its corresponding state sequence (note that $x_0 = x_0$). It follows from (3) that $x_k \in B(x_k; \delta_k)$ for all $k \in \{0, \dots, n-1\}$; which together with (2) implies that

$$x_k \in B(x_k; \delta_k); \quad \forall k \in \{0, \dots, n-1\}.$$

In light of (1), $Q_k(x_k; \cdot)$ is a convex function on the convex set $B(u_k; \delta_k) \setminus A$. Because $x_k \in B(u_k; \delta_k) \setminus A$ is a local minimizer of the function $Q_k(x_k; \cdot)$, it must be a global minimizer of this function over $B(u_k; \delta_k) \setminus A$. Thus, for $k \in \{0, \dots, n-1\}$, we have

$$Q_k(x_k; u_k) + J_{k+1}(x_{k+1}) = Q_k(x_k; u_k) - Q_k(x_k; x_k) = J_k(x_k);$$

By adding all of the above inequalities, one can obtain

$$C(x_0; u_0, \dots, u_{n-1}) - J_0(x_0) = C(x_0; u_0, \dots, u_{n-1});$$

which shows that (u_0, \dots, u_{n-1}) is a local minimizer of the one-shot problem. Also, if (u_0, \dots, u_{n-1}) is a spurious locally minimum control policy of DP, namely $J_0(x_0) > J_0(x_0)$, then

$$C(x_0; u_0, \dots, u_{n-1}) = J_0(x_0) > J_0(x_0);$$

As a result, (u_0, \dots, u_{n-1}) is also a spurious local minimizer of the one-shot problem. ■

Remark 1: By taking the contrapositive, one can immediately conclude that the DP method cannot produce any spurious locally minimum control policies that satisfy the regularity conditions in Theorem 1 as long as the one-shot problem has no spurious local minima.

C. Stationary points: From DP to one-shot optimization

In this subsection, we will show that the induced controlled stationary point of the one-shot problem, under some conditions milder than the assumptions of Theorem 1.

Definition 4: Given a set S and a continuously differentiable function g , a point $s \in S$ is said to be a stationary point of the optimization problem $\min_{s \in S} g(s)$ if

$$r_s g(s) \in N_S(s);$$

where $N_S(s)$ denotes the normal cone of the set S at the point s [32]. We branch off into two specific notions of stationarity below.

Definition 5 (Stationary point): A vector of control inputs (u_0, \dots, u_{n-1}) is said to be a stationary point of the one-shot optimization if for all $k \in \{0, \dots, n-1\}$, it holds that

$$r_{u_k} C(x_0; u_0, \dots, u_{n-1}) \in N_A(u_k).$$

Definition 6 (Stationary control policy): A control policy (u_0, \dots, u_{n-1}) is said to be a stationary control policy of DP if for all $k \in \{0, \dots, n-1\}$ and for all $x \in \mathbb{R}^N$, it holds that $r_x Q_k(x; u_k(x)) \in N_A(u_k(x))$.

Now, we will prove that a stationary control policy (which involves a locally minimum control policy) implies a stationary point of the one-shot optimization under mild assumptions. Let $D_k(x)$ be the Jacobian matrix of $f_k(\cdot)$ at point x , $D_k^{f;x}(x; u)$ be the Jacobian matrix of the function $f_k(\cdot; u)$ at point x while viewing u as a constant, and $D_k^{f;u}(x; u)$ be the Jacobian matrix of $f_k(x; \cdot)$ at point u while viewing x as a constant.

Theorem 2: Consider a stationary control policy (u_0, \dots, u_{n-1}) , and let the associated input and state sequences with the initial state x_0 be denoted as (u_0, \dots, u_{n-1}) and (x_0, \dots, x_n) . If for every $k \in \{0, \dots, n-1\}$:

- 1) f_k is continuously differentiable in a neighborhood of x_k ,
- 2) either $x_k(x_k)$ is in the interior of A or $D_k(x_k) = 0$,

then (u_0, \dots, u_{n-1}) is a stationary point of the one-shot optimization.

Proof: First, we will apply induction to prove that

$$r_x J_k(x_k) = r_x C(x_k; u_k; \dots, u_{n-1}) \quad (4)$$

holds for $k \in \{0, \dots, n\}$. The base step $k = n$ is obvious. For the induction step, observe that

$$\begin{aligned} r_x Q_k(x; u) &= r_x \alpha_k(x; u) + D_k^{f;x}(x; u)^T r_x J_{k+1}(f_k(x; u)); \\ r_x J_k(x) &= r_x [Q_k(x; u_k(x))] \\ &= r_x Q_k(x; u_k(x)) + D_k(x)^T r_u Q_k(x; u_k(x)); \end{aligned}$$

Therefore,

$$\begin{aligned} r_x J_k(x_k) &= r_x \alpha_k(x_k; u_k) + D_k^{f;x}(x_k; u_k)^T r_x J_{k+1}(x_{k+1}) \\ &\quad + D_k(x_k)^T r_u Q_k(x_k; u_k); \end{aligned} \quad (5)$$

If u_k is in the interior of A , we have $r_u Q_k(x_k; u_k) = 0$ by stationarity. Otherwise, by the assumption, we have $x_k(x_k) = 0$. In either case, the last term of (5) is zero. Meanwhile,

$$\begin{aligned} r_x C(x; u_k; \dots, u_{n-1}) &= r_x \alpha_k(x; u_k) + r_x [C(f_k(x; u_k); u_{k+1}; \dots, u_{n-1})] \\ &= r_x \alpha_k(x; u_k) + D_k^{f;x}(x; u_k)^T r_x C(f_k(x; u_k); u_{k+1}; \dots, u_{n-1}); \end{aligned}$$

Now, (4) can be obtained by taking $x_k = x_k$ in the above continuous. For $k < n$, assume that x_{k+1}, \dots, x_{n-1} and equality and then combining it with the induction hypothesis x_{k+1}, \dots, x_n with the above properties have been found. and (5). Finally, for $k = 0, \dots, n-1$, one can write

$$\begin{aligned} & r_u C(x_0; u_0; \dots; u_{n-1}) \\ &= r_u Q_k(x_k; u_k) + D_k^{f;u}(x_k; u_k)^T r_x C(x_{k+1}; u_{k+1}; \dots; u_{n-1}) \\ &= r_u Q_k(x_k; u_k) + D_k^{f;u}(x_k; u_k)^T r_x J_{k+1}(x_{k+1}) \\ &= r_u Q_k(x_k; u_k); \end{aligned}$$

in which the second equality is due to (4). Since u_k is a stationary point of $Q_k(x_k; \cdot)$, $r_u Q_k(x_k; u_k) \in N_A(u_k)$. Thus, $r_u C(x_0; u_0; \dots; u_{n-1}) \in N_A(u_k)$, which proves that $(u_0; \dots; u_{n-1})$ is a stationary point of the one-shot optimization. ■

D. Local minimizers: From one-shot optimization to DP

In this subsection, we will show that each strict local minimizer of the one-shot problem is induced by a locally minimum control policy of DP. Before proving the theorem, we first provide the following useful lemma.

Lemma 1: Given a function $g: \mathbb{R}^N \rightarrow \mathbb{R}$, a point $x \in \mathbb{R}^N$ and a number $\epsilon > 0$, if $u \in A$ is a strict local minimizer of the function $g(x; \cdot)$ and g is continuous in a neighborhood of $(x; u)$, then there exist $\delta > 0$ and a function $h: B(x; \delta) \rightarrow A$ such that $h(x) = u$ and that the following statements hold for all $x_k \in B(x; \delta)$:

- 1) $h(x)$ is a local minimizer of $g(x; \cdot)$.
- 2) $h(x) \in B(u; \delta)$.
- 3) The function $g(x; h(x))$ is continuous at x .

Proof: The proof is given in [1] (see Lemma 1). ■

Theorem 3: If the one-shot problem has a (spurious) strict local minimizer $(u_0; \dots; u_{n-1})$, then there exists a (spurious) locally minimum control policy of DP with the property that $u_k(x_k) = u_k$ for all $k = 0, \dots, n-1$, where $(x_0; \dots; x_n)$ is the state sequence associated with the (spurious) solution of the one-shot problem.

Proof: Let $(u_0; \dots; u_{n-1})$ be a strict local minimizer of the one-shot problem. There exists $\delta > 0$ such that

$$C(x_0; u_0; \dots; u_{n-1}) < C(x_0; u_0; \dots; u_{n-1}); \quad (6)$$

for every control sequence $(u_0; \dots; u_{n-1}) \in (u_0; \dots; u_{n-1})$ with the property that $u_i \in B(u_i; \delta) \cap A$ for $i = 0, \dots, n-1$. In what follows, we will prove by a backward induction that there exist policies $u_0; \dots, u_{n-1}$, positive numbers $\delta_0; \dots, \delta_{n-1}$, and corresponding cost-to-go functions $J_0; \dots, J_n$ such that they jointly satisfy the following properties:

- 1) $u_k(x_k)$ is a local minimizer of the function $Q_k(x_k; \cdot)$ for all $x_k \in \mathbb{R}^N$.
- 2) $u_k(x_k) = u_k$.
- 3) For all $x_k \in B(x_k; \delta_k)$, it holds that

$$u_k(x_k) \in B(u_k; \delta_k); \quad f_k(x_k; u_k(x_k)) \in B(x_{k+1}; \delta_{k+1});$$

- 4) J_k is lower semi-continuous on \mathbb{R}^N and continuous on $B(x_k; \delta_k)$.

For the base step $k = n$, we choose an arbitrary $\delta_n > 0$ and notice that $J_n(x) = c_n(x)$, implying that J_n is always

First, by the continuity of f_k , there exist $\delta_k > 0$ and $0 < \delta_k < \delta_{k+1}$ such that

$$f_k(x_k; u_k) \in B(x_{k+1}; \delta_{k+1}); \quad g(x_k; u_k) \in S_k; \quad (7)$$

where $S_k = B(x_k; \delta_k) \cap (B(u_k; \delta_k) \cap A)$. Since $Q_k(x_k; u_k) = c_k(x_k; u_k) + J_{k+1}(f_k(x_k; u_k))$ and J_{k+1} is continuous on $B(x_{k+1}; \delta_{k+1})$, Q_k is continuous on S_k . Next, for every $u_k \in B(u_k; \delta_k) \cap A$, if we define

$$\begin{aligned} x_{k+1} &= f_k(x_k; u_k); \quad u_{k+1} = u_{k+1}(x_{k+1}); \\ x_{k+2} &= f_{k+1}(x_{k+1}; u_{k+1}); \quad u_{k+2} = u_{k+2}(x_{k+2}); \\ &\vdots \\ x_{n-1} &= f_{n-2}(x_{n-2}; u_{n-2}); \quad u_{n-1} = u_{n-1}(x_{n-1}); \end{aligned}$$

by applying (7) and then the third property above repeatedly, we arrive at

$$u_i \in B(u_i; \delta_i) \cap A; \quad \delta_i \leq \delta_{k+1}; \quad i = 0, \dots, n-1;$$

When $u_k \in B(u_k; \delta_k)$, it follows from (6) and the second property above that

$$\begin{aligned} Q_k(x_k; u_k) &= C(x_k; u_k; \dots; u_{n-1}) \\ &= C(x_0; u_0; \dots; u_{k-1}; u_k; \dots; u_{n-1}) \\ &> C(x_0; u_0; \dots; u_{n-1}) \\ &= C(x_k; u_k; \dots; u_{n-1}) = Q_k(x_k; u_k); \end{aligned}$$

As a result, u_k is a strict local minimizer of $Q_k(x_k; \cdot)$. Applying Lemma 1 to the function Q_k with x_k and u_k , one can find $0 < \delta_k < \delta_{k+1}$ and a function $h_k: B(x_k; \delta_k) \rightarrow A$ such that $h_k(x_k) = u_k$ and that the following statements hold for every $x_k \in B(x_k; \delta_k)$:

- 1) $h_k(x_k)$ is a local minimizer of $Q_k(x_k; \cdot)$.
- 2) $h_k(x_k) \in B(u_k; \delta_k) \cap B(u_k; \delta_k)$, which together with (7) implies that $f_k(x_k; h_k(x_k)) \in B(x_{k+1}; \delta_{k+1})$.
- 3) The function $Q_k(x_k; h_k(x_k))$ is continuous at x_k .

Let h_k be the extension of the function h_k by setting $h_k(x_k)$ to be any global minimizer of the lower semi-continuous function $Q_k(x_k; \cdot)$ over the compact set A if $x_k \notin B(x_k; \delta_k)$. Obviously, h_k satisfies the first three properties. To verify the last property, observe that

$$J_k(x_k) = \begin{cases} Q_k(x_k; h_k(x_k)); & \text{if } x_k \in B(x_k; \delta_k), \\ H_k(x_k); & \text{otherwise,} \end{cases}$$

in which $H_k(x_k) = \min_{u \in A} Q_k(x_k; u)$, and therefore J_k is continuous on the set $B(x_k; \delta_k)$. In addition, note that A is compact and thus Q_k is lower semi-continuous, while A is compact. Hence, it follows from the Berge maximum theorem [33] that H_k is also lower semi-continuous on \mathbb{R}^N , which implies that J_k is lower semi-continuous on $B(x_k; \delta_k)$. For every point x_k on the boundary of $B(x_k; \delta_k)$, since H_k is lower

semi-continuous at x_k , for every $\epsilon > 0$ there exists $\delta > 0$ such that

$$J_k(x_k) - H_k(x_k) > H_k(x_k) - J_k(x_k)$$

holds for all $x_k \in B(x_k; \delta)$. Therefore, J_k is also lower semi-continuous at x_k .

By the first and second properties, $\mu = (u_0; \dots; u_{n-1})$ is a locally minimum control policy of DP. Also, if $(u_0; \dots; u_{n-1})$ is a spurious local minimizer of the one-shot problem, then $J_0(x_0) = C(x_0; u_0; \dots; u_{n-1}) > J_0(x_0)$; which implies that μ is also a spurious locally minimum control policy of DP.

Remark 2: Theorem 3 shows that, under mild conditions, DP is a reformulation from a single one-shot optimization problem to a sequence of optimization problems that preserves local minimizers. By taking the contrapositive of Theorem 3, one can immediately obtain the result that the one-shot problem has no spurious strict local minimizers as long as DP has no spurious locally minimum control policies.

Remark 3: Pontryagin's minimum principle implies that a global minimizer of the one-shot problem achieves a global optimality of each DP problem minimizing Hamiltonian. One can restrict the domain to apply the principle to a local minimizer of the one-shot problem; it achieves a local optimality for each DP problem if J_k are evaluated at the associated state x_k . Theorem 3 is a generalization of Pontryagin's principle in the sense that from each local minimizer of the one-shot problem, we obtain a locally minimum control policy, instead of u_k ; i.e., a set of functions that achieves a local optimality of every DP problem for all $x \in \mathbb{R}^N$. We further require a "strict" local minimizer of the one-shot problem to ensure that a local optimality is obtained at all points in the neighborhood of x_k . Meanwhile, one can now anticipate that

Theorem 1 would correspond to the converse of Pontryagin's principle. The principle provides sufficient conditions for the one-shot problem if we have a convex action space, convex cost functions, and linear dynamics [3], [34]. In contrast, Theorem 1 assumes a convex action space but still has general nonlinear transition dynamics. Theorem 1 instead requires locally "strictly" convex Q-functions (Hamiltonian) for each DP sub-problem. The connection between our results and Pontryagin's principle suggests the possibility of the extension of the above results to the continuous-time setting.

Remark 4: In fact, all results of our paper can be naturally generalized to the continuous-time setting, but the analysis is left as future work due to space restrictions. To outline the pathway for generalization, note that the Hamilton-Jacobi-Bellman equation for a given continuous-time system can be obtained from developing a discrete-time model, obtaining the Bellman equation for that model, and then closing the gap between the continuous-time and discrete-time system by taking a limit [3]. Moreover, the infinite-horizon case is also treated in [3] as the stationary limit of a finite-horizon problem, which again allows us to extend our results to the infinite-horizon case.

Considering Theorems 1 and 3 altogether, one can conclude that under mild conditions, each local minimizer of the one-shot optimization corresponds to some control input induced by a locally minimum control policy, and vice versa.

(a) Example 1

(b) Example 2

Fig. 1. Landscape of the one-shot optimization: (a) Each local minimizer is equivalent to a set of control inputs induced by each locally minimum control policy. (b) $(0; 0)$ is a control input induced by a locally minimum control policy but not a local minimizer of the one-shot optimization. However, it is indeed a stationary point of the one-shot optimization.

Numerical Examples

To effectively demonstrate the results of this section via visualization, we will provide two low-dimensional examples.

Example 1 Consider an optimal control problem with the control constraint $A = [-10; 10]$ and

$$\begin{aligned} c_0(x; u) &= 0; \\ c_1(x; u) &= \frac{1}{4}u^4 - \frac{3x+4}{3}u^3 + \frac{3x^2+8x+3}{2}u^2 \\ &\quad - x(x+1)(x+3)u + \exp(x^4); \\ c_2(x) &= 0; f_0(x; u) = x + u; f_1(x; u) = x + u; \end{aligned}$$

At the initial state $x_0 = 0$, the one-shot problem is written as

$$\min_{u_0 \in A; u_1 \in A} \frac{1}{4}u_1^4 - \frac{3u_0+4}{3}u_1^3 + \frac{3u_0^2+8u_0+3}{2}u_1^2 - u_0(u_0+1)(u_0+3)u_1 + \exp(u_0^4) :$$

This one-shot optimization problem has 3 spurious local minimizers $(0; 0.523)$; $(0; 5.23)$; $(0.938; 0.938)$ and one globally optimal minimizer $(0.938; 3.938)$. The landscape of this objective function is shown in Fig. 1a.

The optimal control problem can also be solved sequentially by DP. At the time step 1, the Q-function is $Q_1(x; u_1) = c_1(x; u_1)$, which has the maximum point $x + 1$, the spurious local minimizer x and the global minimizer $x + 3$. One can choose between the two different continuous policies

$$\mu_1(x) = \begin{cases} x; & |x| \leq 10; \\ 10 \operatorname{sgn}(x); & \text{otherwise} \end{cases}$$

$$\mu_1(x) = \begin{cases} x + 3; & 13 \leq x \leq 7; \\ 10 \operatorname{sgn}(x); & \text{otherwise} \end{cases}$$

where $\operatorname{sgn}(x)$ denotes the sign of x . The first policy has the cost-to-go function $J_1(x) = \frac{1}{12}(3x^4 + 16x^3 + 18x^2) + \exp(x^4)$ for $|x| \leq 10$ and the second policy has $J_1(x) = \frac{1}{12}(3x^4 + 16x^3 + 18x^2 + 27) + \exp(x^4)$ for $13 \leq x \leq 7$.

At the time step 0 and the initial state $x_0 = 0$, the Q-function is $Q_0(0; u_0) = J_1(u_0)$. For the first policy, the Q-function has a spurious local minimizer at $u_0 = 0.523$ and a global minimum at $u_0 = 0.938$. If we choose $u_0(0) = 0.523$, then the induced input under of DP is $(0.523; 0.523)$ and if we choose $u_0(0) = 0.938$, then the induced input under

of DP is (0:938; 0:938). Both of these input sequences are solved in a one-shot fashion. However, this method becomes spurious local minimizers of the one-shot problem.

The Q-function of the second policy has a spurious local minimizer at $u_0 = 0:523$ and a global minimum at $u_0 = 0:938$. If we choose $u_0(0) = 0:938$, then the locally minimum control policy is non-spurious and its induced action as well, which would still be a closed-loop control. Note that input (0:938; 3:938) is the global minimizer of the one-shot problem. However, if we choose $u_0(0) = 0:523$, then its spurious and its induced input (0:523; 2:477) is the spurious minimizer of the one-shot problem.

In this example, one can observe that each strictly local minimizer of the one-shot problem corresponds to a local minimum control policy of DP, which validates the result of Theorem 3. In addition, it can be noticed that since $r_u^2 Q_0(0; 0:523)$ and $r_u^2 Q_0(0; 0:938)$ are both strictly positive for each of the two policies, Theorem 1 clearly holds.

Example 2 Consider the problem in Example 1 but change $c_1(x; u)$ to $\frac{1}{4}u^4 - \frac{x}{3}u^3 - x^2u^2 + \exp(x^4)$. At the initial state $x_0 = 0$, the one-shot problem can be written as

$$\min_{u_0 \in A; u_1 \in A} \frac{1}{4}u_1^4 - \frac{u_0}{3}u_1^3 - u_0^2u_1^2 + \exp(u_0^4) \quad (8)$$

It has 3 stationary points $(0; 0)$ and $(\frac{1}{4}(\log(\frac{8}{3}))^{\frac{1}{4}}; 2(\log(\frac{8}{3}))^{\frac{1}{4}})$ and $(-\frac{1}{4}(\log(\frac{8}{3}))^{\frac{1}{4}}; 2(\log(\frac{8}{3}))^{\frac{1}{4}})$. The latter two are the global minimizers of this one-shot problem. For $(0; 0)$, we take $u_0 = u_1 = 0$ and use the Taylor expansion of the exponential function to arrive at $\frac{1}{4}u^4 - \frac{x}{3}u^3 - x^2u^2 + \exp(x^4) = \frac{1}{12}u^4 + 1 + \alpha(x^4)$; which is strictly less than 1 for sufficiently small values of x . This implies that $(0; 0)$ is not a local minimizer of the one-shot problem. The landscape of this objective function is shown in Fig. 1b. It can also be solved sequentially by DP. For the initial state x_0 , it has 3 different induced input sequences under the locally minimum control policy $(\frac{1}{4}(\log(\frac{8}{3}))^{\frac{1}{4}}; 2(\log(\frac{8}{3}))^{\frac{1}{4}})$, $(-\frac{1}{4}(\log(\frac{8}{3}))^{\frac{1}{4}}; 2(\log(\frac{8}{3}))^{\frac{1}{4}})$ and $(0; 0)$. The first two points are the global minimizers of the one-shot problem. $(0; 0)$ is not a local minimizer of the one-shot problem.

In this example, $r_u^2 Q_1(0; 0) = r_u^2 c_1(0; 0) = 0$ violates the assumptions in Theorem 1, and $(0; 0)$ is not a local minimizer of the one-shot problem. This clarifies the role of the regularity conditions needed in the theorem. On the other hand, $Q_1(x; \cdot)$ has three stationary control policies $x; 2x$. Consistent with Theorem 2, $(0; 0)$ is a saddle point (which is a stationary point) of the one-shot optimization.

III. DETERMINISTIC PROBLEM UNDER A PARAMETERIZED POLICY

A. Problem Formulation

In Section II, the one-shot optimization approach is referred to as an open-loop control, in the sense that it determines the control inputs at once, only given an initial state. On the other hand, the dynamic programming approach is referred to as a closed-loop control, in the sense that the control input at each time step is the function of the output of the previous step [3]. In this section, we formulate both approaches to a closed-loop control. To achieve this, we can replace the control inputs of the one-shot optimization with a parameterized policy. We still optimize over a vector at once, which means that it can

Now, we modify the deterministic problems (P1), (P2), and (P3) to a discrete-time finite-horizon deterministic optimal control problem under a parameterized policy as follows:

$$\begin{aligned} \min_{\theta_0, \dots, \theta_{n-1}} \sum_{i=0}^{n-1} c_i(x_i; \theta_i(x_i)) + c_n(x_n) \\ \text{s.t. } x_{i+1} = f_i(x_i; \theta_i(x_i)); \quad i = 0, \dots, n-1; \\ x_0 \text{ is given} \end{aligned} \quad (\text{PP1})$$

Definition 8: Given a control policy parameter vector $\theta = (\theta_0, \dots, \theta_{n-1})$, the associated Q-function $Q_k(\cdot; \theta)$ and cost-to-go functions $J_k(\cdot)$ under the policy θ are defined in a backward way from the time step n to the time step 0 through the following recursion:

$$\begin{aligned} J_n(x) &= c_n(x); \\ Q_k(x; \theta(x)) &= c_k(x; \theta(x)) + J_{k+1}(f_k(x; \theta(x))); \\ & \quad k = 0, \dots, n-1; \\ J_k(x) &= Q_k(x; \theta(x)); \quad k = 0, \dots, n-1; \end{aligned}$$

Then, the one-shot optimization problem (PP1) can be equivalently written as

$$\min_{\theta_0, \dots, \theta_{n-1}} J_0(x_0) \quad (\text{PP2})$$

and DP approach can be written as the following backward recursion:

$$\begin{aligned} J_n(x) &= c_n(x); \\ J_k(x) &= \min_{\theta} c_k(x; \theta(x)) + J_{k+1}(f_k(x; \theta(x))); \end{aligned} \quad (\text{PP3})$$

for $k = 0, \dots, n-1$. Note that θ was previously defined as a control policy $(\theta_0, \dots, \theta_{n-1})$, but we use the equivalent definition $(\theta_0, \dots, \theta_{n-1})$ in the parameterized case. We also call it control policy parameter vector alternatively.

Definition 9 (local minimizer of the one-shot optimization): A control policy parameter vector $\theta = (\theta_0, \dots, \theta_{n-1})$ is said to be a local minimizer of the one-shot optimization if there exists $\delta > 0$ such that

$$\begin{aligned} J_0(x_0) &\leq \tilde{J}_0(x_0) \\ \text{for all } \tilde{\theta} &= (\tilde{\theta}_0, \dots, \tilde{\theta}_{n-1}) \in (B(\theta; \delta) \setminus \theta) \end{aligned}$$

Definition 10 (local minimizer of DP): A control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$ is said to be a local minimizer of DP if for all $k \in \{0, \dots, n-1\}$ and for all $x \in \mathbb{R}^N$, the policy parameter θ_k is a local minimizer of $Q_k(x; \theta_{-k}(x))$, meaning that there exists $\delta_k > 0$ such that

$$Q_k(x; \theta_k(x)) \leq Q_k(x; \theta_{-k}(x)); \quad \forall x \in B(\theta_k; \delta_k) \setminus \{\theta_k\} \quad (9)$$

Definition 11 (Stationary point of the one-shot optimization): A control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$ is said to be a stationary point of the one-shot optimization if for all $k \in \{0, \dots, n-1\}$, it holds that $\theta_k \in \arg \min_{\theta_k} J_0(x_0; \theta_{-k})$.

Definition 12 (Stationary point of DP): A control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$ is said to be a stationary point of DP if for all $k \in \{0, \dots, n-1\}$ and for all $x \in \mathbb{R}^N$, it holds that $\theta_k \in \arg \min_{\theta_k} Q_k(x; \theta_{-k}(x))$.

Remark 5: By comparing (P2) with (PP2) as well as comparing Definition 2 with Definition 9, notice that one-shot optimization now considers $J_0(x_0)$ instead of $C(x_0; \theta_0; \dots; \theta_{n-1})$, since the two definitions are equivalent when the parameterized policy is incorporated.

We can compare Definition 10 with the following definition:

Let $k \in \{0, \dots, n-1\}$; $x \in \mathbb{R}^N$; $\delta_k > 0$ such that

$$Q_k(x; \theta_k(x)) \leq Q_k(x; \theta_{-k}(x)); \quad \forall x \in B(\theta_k; \delta_k) \setminus \{\theta_k\} \quad (10)$$

Definition 10 considers the open ball centered at the policy parameter in the parameter space, while (10) considers the corresponding open ball in the action space. Proposition 1 establishes the relationship between these definitions.

Proposition 1: If an arbitrary control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$ satisfies (10) with $\inf_{x \in \mathbb{R}^N} Q_k(x) > 0$ for all $k \in \{0, \dots, n-1\}$, then it is a local minimizer of DP.

Proof: Since $\inf_{x \in \mathbb{R}^N} Q_k(x) > 0$, by the continuity assumption, for every $k \in \{0, \dots, n-1\}$, there exists $\delta_k > 0$ such that

$$\sup_{x \in B(\theta_k; \delta_k)} Q_k(x) - Q_k(\theta_k) < \inf_{x \in \mathbb{R}^N} Q_k(x)$$

That is, for all $x \in B(\theta_k; \delta_k) \setminus \{\theta_k\}$, $Q_k(x) > Q_k(\theta_k)$ for all $x \in \mathbb{R}^N$. Notice that Definition 7 implies that $\theta_k \in \arg \min_{\theta_k} Q_k(x)$ for all $x \in \mathbb{R}^N$. Thus, it holds for all $x \in \mathbb{R}^N$ that

$$\theta_k \in \arg \min_{\theta_k} Q_k(x) \quad \forall x \in \mathbb{R}^N$$

Thus, given a control policy parameter vector satisfying (10), for all $k \in \{0, \dots, n-1\}$ and for all $x \in \mathbb{R}^N$, (9) holds if one substitutes θ_k with $\theta_k(x)$. This completes the proof. ■

Remark 6: The converse of Proposition 1 does not hold. For example, suppose there exists $\delta_k > 0$ such that $Q_k(x)$ takes the same value for all $x \in B(\theta_k; \delta_k) \setminus \{\theta_k\}$. While this control policy satisfies the continuity assumption, it is clearly a local minimizer of DP, which satisfies (9). However, it is even possible that $\theta_k(x)$ is a strict local maximizer of $Q_k(x)$.

Note that the condition $\inf_{x \in \mathbb{R}^N} Q_k(x) > 0$ is necessary for Proposition 1. Thus, the proposition implies that if we use our notion of a local minimizer of DP, we no longer need to assume $\inf_{x \in \mathbb{R}^N} Q_k(x) > 0$ while establishing the relationship from DP to one-shot optimization, which was the case in the (non-parameterized) deterministic case presented in our conference paper (see Theorem 2 in [1]).

B. From DP to one-shot optimization

In this subsection, we will show that in the deterministic case with a parameterized policy, each local minimizer (stationary point) of DP directly corresponds to some local minimizer (stationary point) of the one-shot optimization.

Theorem 4: Consider a local minimizer of DP $\theta = (\theta_0; \dots; \theta_{n-1})$. Then, θ is also a local minimizer of the one-shot optimization.

Proof: Since $\theta = (\theta_0; \dots; \theta_{n-1})$ is a local minimizer of DP, there exist $\delta_0; \dots; \delta_{n-1} > 0$ such that

$$\begin{aligned} J_0(x_0) &= Q_0(x_0; \theta_0(x_0)) \leq Q_0(x_0; \theta_{-0}(x_0)) \\ &= c_0(x_0; \theta_{-0}(x_0)) + Q_1(x_1; \theta_1(x_1)) \\ &\quad (x_1 = f_0(x_0; \theta_{-0}(x_0))) \\ &\leq c_0(x_0; \theta_{-0}(x_0)) + Q_1(x_1; \theta_1(x_1)) \\ &= c_0(x_0; \theta_{-0}(x_0)) + c_1(x_1; \theta_1(x_1)) + Q_2(x_2; \theta_2(x_2)) \\ &\quad (x_2 = f_1(x_1; \theta_1(x_1))) \\ &\leq J_0^-(x_0) \end{aligned}$$

where $\tilde{\theta} = (\tilde{\theta}_0; \dots; \tilde{\theta}_{n-1}) \in (B(\theta_0; \delta_0) \setminus \{\theta_0\}) \times \dots \times (B(\theta_{n-1}; \delta_{n-1}) \setminus \{\theta_{n-1}\})$.

Choose $\tilde{\theta} = \arg \min_{\tilde{\theta}} J_0(x_0)$. Then, $J_0(x_0) = J_0^-(x_0)$ for all $\tilde{\theta} = (\tilde{\theta}_0; \dots; \tilde{\theta}_{n-1}) \in (B(\theta_0; \delta_0) \setminus \{\theta_0\}) \times \dots \times (B(\theta_{n-1}; \delta_{n-1}) \setminus \{\theta_{n-1}\})$. This completes the proof. ■

Theorem 5: Consider a stationary point of DP $\theta = (\theta_0; \dots; \theta_{n-1})$. Let the corresponding state sequence be $(x_0; \dots; x_n)$. If for every $k \in \{0, \dots, n-1\}$, $\theta_k(x_k)$ is continuously differentiable with respect to x_k in a neighborhood of $(x_k; \theta_k)$, then θ is also a stationary point of the one-shot optimization.

Proof: Notice that $r_{\theta_k} J_0(x_0) = r_{\theta_k} J_k(x_k) = r_{\theta_k} Q_k(x_k; \theta_{-k}(x_k))$. Thus, $r_{\theta_k} J_0(x_0) \in N(\theta_k)$ for all $k \in \{0, \dots, n-1\}$, which means that θ is a stationary point of the one-shot optimization. ■

Remark 7: The converse of Theorem 5 clearly does not hold since one can generally find a point $x \in \mathbb{R}^N$ such that $r_x Q_k(x; \theta_{-k}(x)) \notin r_{\theta_k} Q_k(x; \theta_{-k}(x))$.

C. From one-shot optimization to DP

In this subsection, we first show that a local minimizer of the one-shot optimization does not necessarily correspond to a local minimizer of DP, i.e., the converse of Theorem 4 does not hold. Then, with Remark 7, it is clear that the optimization landscape of the one-shot optimization is more complex than that of DP. As a by-product, if the one-shot problem has a complexity, so does the DP problem.

To develop a clear counterexample, we restrict the parameterized policy to a certain class as given below, which automatically satisfies the continuity assumption defined in Definition 7.

Definition 13: Define our parameterized policy to be a linear combination of arbitrary linearly independent basis functions, while satisfying Definition 7; i.e., Given m functions $f_i: \mathbb{R}^N \rightarrow \mathbb{R}^m$, $i = 1; \dots; m$ and $s = [s_1; \dots; s_m]^T \in \mathbb{R}^m$,

$$x_k = \sum_{i=1}^m s_i f_i(x) \in A; \quad (11)$$

where there does not exist $(s_1, \dots, s_m) \in \mathbb{R}^m$ such that for all x in any set of non-zero measure, the following equation holds [36]:

$$\sum_{i=1}^m s_i f_i(x) = 0 \tag{12}$$

Remark 8: Since a set of isolated points is a set of measure zero, it is exempt from determining the independence of the one-shot optimization for all basis functions. When x has a continuous distribution, the independence of basis functions implies that if (12) holds for all x in the support of the distribution, then $s_1 = \dots = s_m = 0$. When x has a discrete distribution, since a set of all the possible values of x is a set of measure zero, the independence of basis functions does not guarantee $s_1 = \dots = s_m = 0$ even if (12) holds for all possible values of x .

Applications of a parameterized policy defined by Definition 13 arise in a piecewise polynomial function as well as a linear combination of kernels fully represents the solution of minimizing the empirical risk. It switches the optimization problem in finite-dimensional function space to finding the finite number of coefficients. The minimum number of parameters needed is the number of data points, which is generally much greater than the dimension of the output. Applying this to our problem, the number of parameters needs to be greater than the dimension of the action space to cover all data points. For the remainder of this section, we call a policy satisfying $m > M$ as an overparameterized policy.

We now provide some evidence to refute the converse of Theorem 4, specially if the parameterized policy class is a linear combination of basis functions. It turns out that a local minimizer of the one-shot optimization does not necessarily imply a local minimizer of DP in the overparameterized case.

Proposition 2: Consider an overparameterized policy class defined by Definition 13. Let (x_0, \dots, x_{n-1}) be a local minimizer of DP. If there exists at least one $k \in \{0, \dots, n-1\}$ such that x_k is in the interior of \mathcal{X} , then there exists an infinite number of local minimizers of the one-shot optimization corresponding to each local minimizer of DP.

Proof: Consider the state sequence (x_0, \dots, x_n) induced by a local minimizer of DP. Let k be an index for which x_k is in the interior of \mathcal{X} . Then, one can express the action taken at step k as $a_k(x_k) = \sum_{i=1}^m s_i f_i(x_k)$ with $s_k = (s_0, s_1, \dots, s_m)$ by Definition 13. Since the policy is overparameterized, m is greater than the dimension of the action space M . Now, consider the matrix equation

$$f_0(x_k) \ f_1(x_k) \ \dots \ f_m(x_k) \cdot k = k(x_k); \tag{13}$$

where k is an $m-1$ vector variable, and let F_k denote the first matrix in the left-hand side, which is an $m \times m$ constant matrix given by x_k . (13) has at least one solution.

The dimension of the null space of F_k is greater than 0 due to $m > M$. We take any nonzero element from the null space. Then, for all $v \in \mathbb{R}^{m-1}$, $k + v$ satisfies (13). Since x_k is in the interior of \mathcal{X} , one can pick $\epsilon > 0$ such that $B(x_k; \epsilon) \subset \mathcal{X}$. Thus, for all $v \in \mathbb{R}^{m-1}$, $k + v \in B(x_k; \epsilon)$ preserves the state and action sequence associated with (13). The induced cost is also indeed preserved.

By Theorem 4, a_k is a local minimizer of the one-shot optimization. Now, we select $\delta > 0$ such that $J_0(x_0) = J_0^-(x_0)$ for all $\tilde{x} = (x_0, \dots, x_{k-1}, \tilde{x}, \dots, x_{n-1})$ where $\tilde{x} \in B(x_k; \delta) \setminus \mathcal{X}$. Let $\tilde{s} := \min_{\tilde{x} \in B(x_k; \delta) \setminus \mathcal{X}} \sum_{i=1}^m s_i f_i(\tilde{x}) > 0$. Then, for all $\tilde{x} \in B(x_k + v; \delta/2)$, we have $\sum_{i=1}^m (s_i + \tilde{s}) f_i(\tilde{x}) > 0$. Since $k + v$ preserves the induced cost $J_0(x_0, \dots, x_{k-1}, \tilde{x}, \dots, x_{n-1})$ is a local minimizer of the one-shot optimization for all $\tilde{x} \in B(x_k + v; \delta/2)$. This completes the proof. ■

Proposition 2 implies that for every $f \in C(\mathcal{X}; \mathbb{R})$, a_k of a "strict" local minimizer of the one-shot optimization does not lie in the interior of \mathcal{X} . Thus, one can think of constructing a strict local minimizer by restricting the area of interest that given a strict local minimizer of the one-shot optimization and the induced input sequence, no other points can retrieve the same input sequence if it is convex.

Lemma 2: Consider a strict local minimizer of the one-shot optimization $a_k = (a_0, \dots, a_{n-1})$. Let (x_0, \dots, x_n) be the induced state sequence. Suppose that a_k is convex and the parameterized policy is defined by Definition 13. Then, a_k is the unique control policy parameter vector that achieves the minimum cost $J_0(x_0, \dots, x_{n-1}, a_k)$.

Proof: For every $k \in \{0, \dots, n-1\}$, $(x_k) = \sum_{i=1}^m s_i f_i(x_k)$ where $s_k = (s_1, \dots, s_m)$. Let $k = (s_1, \dots, s_m)$. Since a_k is a strict local minimizer of the one-shot optimization, $(x_k) \in \mathcal{X}$ in the neighborhood of x_k if $\tilde{s} \in \mathcal{X}$; i.e., there exists $\delta > 0$ such that

$$\forall \tilde{s} \in B(s_k; \delta) \setminus \mathcal{X} : \sum_{i=1}^m \tilde{s}_i f_i(x_k) = \sum_{i=1}^m s_i f_i(x_k) \tag{14}$$

Assume that there exists $\tilde{s} \in \mathcal{X}$ such that $\sum_{i=1}^m \tilde{s}_i f_i(x_k) = \sum_{i=1}^m s_i f_i(x_k)$ where $\tilde{s} = (\tilde{s}_1, \dots, \tilde{s}_m)$. Then, for $\alpha \in [0, 1]$, one can obtain $\sum_{i=1}^m (\alpha \tilde{s}_i + (1-\alpha) s_i) f_i(x_k) = \sum_{i=1}^m s_i f_i(x_k)$ by linearity and $\alpha \tilde{s} + (1-\alpha) s_k \in \mathcal{X}$ by convexity. Letting $\alpha = 1$, one can construct an element of the left-hand side of (14) distinct from s_k . By contradiction, s_k is the unique point that achieves $\sum_{i=1}^m s_i f_i(x_k)$. ■

Note that Lemma 2 does not necessarily imply that a strict local minimizer of the one-shot optimization is a local minimizer of DP even if \mathcal{X} is convex. A simple counterexample can be constructed by considering the 1-step problem

$$\begin{aligned} f_0(x; (d_1, d_2)) &= \frac{1}{4} (x)^4 - \frac{1}{3} (x^2 + 2x) (x)^3 + \\ &\quad \frac{1}{2} (2x^3 + x - 1) (x)^2 - (x^4 - x^3 + x^2 - x) (x); \\ f_1(x; (d_1, d_2)) &= 0; \quad f_2(x; (d_1, d_2)) = x + (x); \end{aligned}$$

with the parameterized policy $a(x) = d_1 x + d_2$ where $(d_1, d_2) \in \mathbb{R}^2$ and $f(d_1, d_2) := \frac{1}{4} (2d_1 - d_2)^2 - \frac{1}{3} (d_1 + d_2)^3 + (d_1 + d_2)^2$ which is convex. At the initial state $x_0 = 1$, the one-shot problem can be written as

$$\min_{(d_1, d_2) \in \mathbb{R}^2} \frac{1}{4} (d_1 + d_2)^4 - (d_1 + d_2)^3 + (d_1 + d_2)^2 :$$

Each vector $(d_1, d_2) \in \mathbb{R}^2$ which satisfies $d_1 + d_2 = 0$ or $d_1 + d_2 = 2$ is a local minimizer of the one-shot optimization. Since $\{(d_1, d_2) : d_1 + d_2 = 0\} = f^{-1}(1; -1)g$ and $\{(d_1, d_2) : d_1 + d_2 = 2\} = f^{-1}(1; 1)g$, we have

(a) Domain (b) Landscape

Fig. 2. The domain and the landscape of the one-shot optimization for a deterministic parameterized problem: (a) The gray-colored area is the domain of the parameter space. The intersection between the dotted lines and the domain is $f(1; 1); (1; 1)g$. (b) Both $(1; 1)$ and $(1; 1)$ are a strict local minimizer of the one-shot optimization but only $(1; 1)$ is a local minimizer of DP.

$(1; 1)$ and $(1; 1)$ as strict local minimizers of the one-shot optimization. On the other hand, since $Q_0(x; (x)) = r c_0(x; (x)) = [g(x; (x)); g(x; (x))]^T$ where $g(x; (x)) = (x(x) - (x^2 + 1))(x(x) - x)(x(x) - (x - 1))$, a local minimizer of DP should be the parameter that yields $x = x - 1$ or $(x) = x^2 + 1$ for all $x \in \mathbb{R}^N$. Since a linear policy cannot contain $x^2 + 1$, $(1; 1) \in \mathbb{A}$ is the only local minimizer of DP. Thus, $(1; 1)$ is a strict local minimizer of the one-shot optimization but not a local minimizer of DP. Fig. 2 shows the domain and the landscape of the one-shot optimization.

In light of the above counterexample, one can think of the situation where the parameterized policy contains every locally minimum control policy of DP (see Definition 3). It turns out that if such a situation is possible, given a convex parameter space, each strict local minimizer of the one-shot optimization is a local minimizer of DP under the following assumptions.

Assumption 1: Given a local minimizer of the one-shot optimization π , let $(x_0; \dots; x_n)$ be the associated state sequence. Then, for all $k \in \{0; \dots; n - 1\}$, the $M \times m$ matrix $[f_0(x_k) \ f_1(x_k) \ \dots \ f_m(x_k)]$ has a full row rank.

Assumption 2: Assume that $\mathbb{A} \setminus \bigcup_{k=1}^n (x_k)$, where (x_k) is the image of π through (x_k) : $\pi(x_k) \in \mathbb{A}$.

Lemma 3: Assume that \mathbb{A} is convex. Consider a strict local minimizer of the one-shot optimization $\pi = (\pi_0; \dots; \pi_{n-1})$. Suppose that the parameterized policy defined by Definition 13 satisfies Assumptions 1 and 2. If the parameterized policy class contains every locally minimum control policy of DP and $\inf_{x \in \mathbb{R}^N} J_k(x) > 0$ for all $k \in \{0; \dots; n - 1\}$, then π is a local minimizer of DP.

Proof: Let $(x_0; \dots; x_n)$ be the state sequence associated with π . Recall that $J_0(x_0) = \sum_{i=0}^{n-1} c_i(x_i; \pi_i(x_i)) + Q_k(x_k; \pi_k(x_k))$. One can vary all parameters except π_k to derive that $J_0(x_0) - J_0^*(x_0) = Q_k(x; \pi_k(x)) - Q_k(x; \pi_k^*(x))$, where $\pi^* = (\pi_0; \dots; \pi_{k-1}; \pi_{k+1}; \dots; \pi_{n-1})$. Thus, a local minimizer of the one-shot optimization implies that for all $k \in \{0; \dots; n - 1\}$, there exists $\delta_k > 0$ such that

$$Q_k(x_k; \pi_k(x_k)) - Q_k(x_k; \pi_k^*(x_k)) \geq \delta_k \quad \forall \pi_k^* \in \mathbb{B}(\pi_k; \delta_k) \setminus \{\pi_k\} \quad (15)$$

Now, let F_k be the $M \times m$ matrix $[f_0(x_k) \ f_1(x_k) \ \dots \ f_m(x_k)]$, where its smallest singular value is denoted by σ_k . Given an arbitrary direction $v \in \mathbb{R}^m$, one can take a point $x_k + v$ that is farthest from x_k in the direction of v since the action space is compact. Let v be the value that achieves $x_k + v = x_k + \frac{v}{\|v\|} \|v\|$. By Assumption 2, there exists $v \in \mathbb{B}(0; \delta_k)$ satisfying $u_v = \pi_v(x_k)$, and by Definition 13, $\pi_v(x_k)$ is defined by $F_k^{-1} v$.

Case 1 $\sigma_k = 0$: There does not exist $\delta > 0$ such that $\pi(x_k) + v \in \mathbb{A}$.

Case 2 $\sigma_k > 0$ and $v \in \mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2})$: Due to the linearity of policy and the convexity of \mathbb{A} , there exists $\delta \in \mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{0\}$ such that $\pi(x_k) = \pi(x_k) + v$ for all $0 < \|v\| < \delta$.

Case 3 $\sigma_k > 0$ and $v \notin \mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2})$: Consider $\pi(x_k) + \frac{\delta_k}{2\|v\|} (\pi_v(x_k) - \pi(x_k))$. The corresponding parameter is definitely in $\mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{\pi(x_k)\}$ by the linearity of policy and the convexity of \mathbb{A} . Then, as in Case 2, there exists $\delta \in \mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{0\}$ such that $\pi(x_k) = \pi(x_k) + v$ for all $0 < \|v\| < \frac{\delta_k}{2\|v\|} \delta$.

In Case 3, notice that $\frac{\delta_k}{2\|v\|} (\pi_v(x_k) - \pi(x_k)) \in \mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{\pi(x_k)\}$ where the last inequality is from Assumption 1 and the second last inequality is from the basic property of singular value [38].

Considering all three cases, $\mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{\pi(x_k)\} \cap \mathbb{A}$ implies that at least one corresponding parameter for each π_k is in $\mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{\pi_k\}$. Thus, one can notice that (15) implies

$$Q_k(x_k; \pi_k(x_k)) - Q_k(x_k; \pi_k^*(x_k)) \geq \delta_k \quad \forall \pi_k^* \in \mathbb{B}(\frac{\delta_k}{2}; \frac{\delta_k}{2}) \setminus \{\pi_k\} \quad (16)$$

We select an arbitrary locally minimum control policy $\tilde{\pi} = (\tilde{\pi}_0; \dots; \tilde{\pi}_{n-1})$ with the property that $\inf_{x \in \mathbb{R}^N} J_k(x) > 0$. Let $\tilde{\pi} = (\tilde{\pi}_0; \dots; \tilde{\pi}_{n-1})$ be the policy such that for all $k \in \{0; \dots; n - 1\}$,

$$\tilde{\pi}_k(x_k) = \begin{cases} \pi_k(x_k); & \text{if } x_k = x_k, \\ \pi_k^*(x_k); & \text{otherwise.} \end{cases}$$

Such $\tilde{\pi}$ is also a locally minimum control policy by (16).

This implies that the parameterized policy contains $\tilde{\pi}$. Also, $\tilde{\pi}$ achieves the same input sequence $(x_0; \dots; x_{n-1})$ as the strict local minimizer. Therefore, by Lemma 2, $\tilde{\pi}$ is a local minimizer of DP. Since $\inf_{x \in \mathbb{R}^N} J_k(x)$ induced by π_k is greater than 0, $\inf_{x \in \mathbb{R}^N} J_k(x)$ induced by $\tilde{\pi}_k$ is also greater than 0. Then, by Proposition 1, $\pi = (\pi_0; \dots; \pi_{n-1})$ is a local minimizer of DP. \blacksquare

Remark 9: With a given set of parameters $(\pi_0; \dots; \pi_{n-1})$, there exists only one associated state sequence for the deterministic parameterized problem. Assumptions 1 and 2 are thus only required for that specific state sequence, where one can readily check the assumptions in advance with known dynamics, parameter space, action space, and policy class. Assumption 1 is a type of regularity condition, which can be regarded as the extension of an overparameterized policy. Assumption 2 implies that \mathbb{A} should be large enough to contain relevant parameters to cover the action space. Since $\pi(x)$ is designed to be in \mathbb{A} by Definition 7, Assumption 2 is equivalent to saying that $\pi(x_0) = \pi(x_n)$.

Meanwhile, suppose that there exist two different locally minimum control policies in a set of non-zero measure, μ_1 and μ_2 , where $\mu_1(x) \neq \mu_2(x)$ for all $x \in I$ where I is a set of non-zero measure. Then, there exists an infinite number of locally minimum control policies made up of μ_1 and μ_2 by alternating between $\mu_1(x)$ and $\mu_2(x)$ along $x \in I$, and the parameterized policy class cannot contain all these policies. We now present the situation that the parameterized policy contains every locally minimum control policy of DP.

Theorem 6: Assume that $J_k(x) = Q_k(x; w_k(x))$; $k = 0; \dots; n-1$ is convex. Consider a strict local minimizer of the one-shot optimization $J_0(x_0)$ $(x_0 \in \mathcal{B}(x_0; \delta))$. Suppose that the parameterized policy defined by Definition 13 satisfies Assumptions 2 and 1. If there exists only a single locally minimum control policy of DP μ^* $(\mu^* \in \mathcal{P}(I))$ and the parameterized policy class contains μ^* , then μ^* is a local minimizer of DP.

Proof: Let $\theta = (\theta_0; \dots; \theta_{n-1})$ be the parameters associated with μ . For all $k \in \{0; \dots; n-1\}$ and for all $x \in \mathbb{R}^N$, $w_k(x)$ is the unique local minimizer of $Q_k(x; u)$. Having no spurious local minima implies that $\nabla_x Q_k(x; w_k(x)) = 0$. Moreover, the parameterized policy class contains every locally minimum control policy of DP. Since these facts satisfy the preconditions of Lemma 3, this completes the proof.

Considering both Theorem 4 and 6, one can conclude that under the assumptions of Theorem 6, a local minimizer of DP is equivalent to a local minimizer of the one-shot optimization.

IV. STOCHASTIC PROBLEM UNDER A PARAMETERIZED POLICY

A. Problem Formulation

In this section, we will show that the results obtained for the deterministic problem under a parameterized policy also hold for the stochastic problem under a parameterized policy. Since we now take the expectation of the sum of the costs over the trajectories, the issue of strictness, as in Proposition 2, does not take place. Before presenting the theorems, we first define the problem setting in the stochastic case.

Definition 14: Given a complete probability space $(\Omega; \mathcal{F}; P)$, let x_0 be a \mathcal{F} -measurable, \mathbb{R}^N -valued random variable, which has an initial distribution. Also, let w_k be an \mathcal{F} -measurable, \mathbb{R}^W -valued random variable for all $k \in \{0; \dots; n-1\}$ such that $x_0; w_0; \dots; w_{n-1}$ are mutually independent. The state transition is now governed by the dynamics $f_i: \mathbb{R}^N \times \mathbb{R}^W \rightarrow \mathbb{R}^N$, $i = 0; \dots; n-1$. The dynamics are again defined to be at least twice continuously differentiable.

Now, we modify the deterministic problems under a parameterized policy, i.e., (PP1), (PP2), and (PP3), to a discrete-time finite-horizon stochastic optimal control problem under a parameterized policy:

$$\min_{\theta \in \Theta} E_{x_0; w_0; \dots; w_{n-1}} \sum_{i=0}^{n-1} C_i(x_i; w_i) + C_n(x_n);$$

where $x_{i+1} = f_i(x_i; w_i)$; $i = 0; \dots; n-1$;

(SP1)

Notice that for stochastic problems, x_0 may not be given as a point, but has an initial distribution. Afterwards, x_{i+1} is a random variable induced by $(x_i; w_i)$.

Definition 15: Given a control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$, the associated Q-function $Q_k(\cdot; \theta)$ and cost-to-go functions $J_k(\cdot)$ under the policy are defined in a backward way from the time step $n-1$ to the time step 0 through the following recursion:

$$J_n(x) = C_n(x);$$

$$Q_k(x; \theta) = E_{w_k} [C_k(x; w_k) + J_{k+1}(f_k(x; w_k))];$$

$k = 0; \dots; n-1$;

Then, the one-shot optimization problem (SP1) can be equivalently written as

$$\min_{\theta \in \Theta} E_{x_0} [J_0(x_0)]; \tag{SP2}$$

as long as the cost functions C_i , $i = 0; \dots; n-1$, are uniformly bounded, due to the product measure Theorem and Fubini's Theorem [39]. In the remainder of the paper, we assume that the two problems are equivalent.

The DP approach can be written as the following backward recursion:

$$J_n(x) = C_n(x);$$

$$J_k(x) = \min_{w_k} E_{w_k} [C_k(x; w_k) + J_{k+1}(f_k(x; w_k))];$$

$k = 0; \dots; n-1$;

(SP3)

Definition 16 (local minimizer of the one-shot optimization): A control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$ is said to be a local minimizer of the one-shot optimization if there exists $\delta > 0$ such that

$$E_{x_0} [J_0(x_0)] \leq E_{x_0} [J_0^{\sim}(x_0)]$$

for all $\tilde{\theta} = (\tilde{\theta}_0; \dots; \tilde{\theta}_{n-1}) \in (\mathcal{B}(x_0; \delta) \times \Theta) \setminus \{\theta\}$.

Definition 17 (Stationary point of the one-shot optimization): A control policy parameter vector $\theta = (\theta_0; \dots; \theta_{n-1})$ is said to be a stationary point of the one-shot optimization if for all $k \in \{0; \dots; n-1\}$, it holds that $\nabla_{w_k} E_{x_0} [J_0(x_0)] = 0$.

While the one-shot method aims for optimizing the expectation over all steps in the stochastic dynamics, DP studies optimizing Q-function at every step both in the deterministic and stochastic cases. Since we have modified the definition of Q-function to incorporate the expectation, it is natural that the definition of a local minimizer (stationary point) of DP is exactly the same as Definition 10 (12).

B. From DP to one-shot optimization

In this subsection, we will show that, in the stochastic case with a parameterized policy, each local minimizer (stationary point) of DP directly corresponds to some local minimizer (stationary point) of the one-shot optimization, just as in the deterministic case. However, it turns out that for the stationary points, the policy needs to be continuously differentiable with respect to both states and parameters since the expectation is over all trajectories rather than a single trajectory.

Theorem 7: Consider a local minimizer of DP $\theta = (\theta_0; \dots; \theta_{n-1})$. Then, θ is also a local minimizer of the one-shot optimization.

Proof:

Since (x_0, \dots, x_{n-1}) is a local minimizer of DP, there exist $\epsilon_0, \dots, \epsilon_{n-1} > 0$ such that

$$\begin{aligned} E_{x_0}[J_0(x_0)] &= E_{x_0}[Q_0(x_0; x_0(x_0))] - E_{x_0}[Q_0(x_0; \tilde{x}_0(x_0))] \\ &= E_{x_0}[c_0(x_0; \tilde{x}_0(x_0)) + E_{w_0}[Q_1(x_1; x_1(x_1))] \\ &\quad (x_1 = f_0(x_0; \tilde{x}_0(x_0); w_0))] \\ &= E_{x_0}[c_0(x_0; \tilde{x}_0(x_0)) + E_{w_0}[Q_1(x_1; \tilde{x}_1(x_1))] \\ &= E_{x_0}[c_0(x_0; \tilde{x}_0(x_0)) + E_{w_0}[c_1(x_1; \tilde{x}_1(x_1)) \\ &\quad + E_{w_1}[Q_2(x_2; x_2(x_2))] \\ &\quad (x_2 = f_1(x_1; \tilde{x}_1(x_1); w_1))] \\ &= E_{x_0; w_0, \dots, w_{n-1}}[J_0^{\sim}(x_0)] \end{aligned}$$

where $\tilde{x} = (\tilde{x}_0, \dots, \tilde{x}_{n-1}) \in (B(x_0; \epsilon_0) \setminus \{x_0\}) \cap (B(x_{n-1}; \epsilon_{n-1}) \setminus \{x_{n-1}\})$. The last inequality is due to the assumption that the two problems (SP1) and (SP2) are equivalent.

Choose $\tilde{x} = \min_{f_0, \dots, f_{n-1}} g$. Then, $J_0(x_0) = J_0^{\sim}(x_0)$ for all $\tilde{x} = (\tilde{x}_0, \dots, \tilde{x}_{n-1}) \in (B(x_0; \epsilon_0) \setminus \{x_0\}) \cap (B(x_{n-1}; \epsilon_{n-1}) \setminus \{x_{n-1}\})$. This completes the proof. ■

Now, let $D_x(\cdot)$ be the Jacobian matrix of $(f_k)(x)$ at point x , $D_k^{f;x}(x; (x); w)$ be the Jacobian matrix of the function $f_k(\cdot; (\cdot); w)$ at point x while viewing (x) as a constant, and similarly $D_k^{f;x}(x; (x); w)$ be the Jacobian matrix of $f_k(x; (\cdot)(x); w)$ at point (x) while viewing x as a constant.

Theorem 8: Consider a stationary point of DP $= (x_0, \dots, x_{n-1})$. If for all $k \in \{0, \dots, n-1\}$,

- 1) $x_k(x_k)$ is continuously differentiable with respect to x_k in a neighborhood of x_k for all $x_k \in \mathbb{R}^N$;
- 2) $x_k(x_k)$ is continuously differentiable with respect to x_k everywhere,

then (x_0, \dots, x_{n-1}) is a stationary point of the one-shot optimization.

Proof:

First, we will apply induction to prove that for every $k \in \{0, \dots, n-1\}$, $J_k(x)$ is continuously differentiable. For the base step, $J_n(x) = c_n(x)$ is continuously differentiable. For the induction step, observe that

$$\begin{aligned} r_x J_k(x) &= r_x [Q_k(x; x_k(x))] \\ &= r_x [c_k(x; x_k(x)) + \int_{\mathcal{Z}} J_{k+1}(f_k(x; x_k(x); w_k)) dp(w_k)] \\ &= r_x [c_k(x; x_k(x))] + r_x [J_{k+1}(f_k(x; x_k(x); w_k))] dp(w_k) \\ &= r_x [c_k(x; x_k(x))] \\ &\quad + D_k^{f;x}(x; x_k(x); w_k)^T r_x J_{k+1}(f_k(x; x_k(x); w_k)) dp(w_k) \end{aligned}$$

This observation is based on the existence and continuity of the Jacobian matrix $D_k^{f;x}(x; x_k(x); w_k)$ due to assumption 2, continuity of $r_x J_{k+1}(f_k(x; x_k(x); w_k))$ due to the induction step, and therefore the continuity of $r_x [J_{k+1}(f_k(x; x_k(x); w_k))]$. This allows us to interchange integration and differentiation in the second equality by Leibniz's integration rule.

Now, for $k \in \{0, \dots, n-1\}$, observe that

$$r_k Q_k(x_k; x_k(x_k))$$

$$\begin{aligned} &= r_k [c_k(x_k; x_k(x_k)) + \int_{\mathcal{Z}} J_{k+1}(f_k(x_k; x_k(x_k); w_k)) dp(w_k)] \\ &= \int_{\mathcal{Z}} r_k (c_k(x_k; x_k(x_k)) + D_k^{f;x}(x_k; x_k(x_k); w_k)^T r_x J_{k+1}(f_k(x_k; x_k(x_k); w_k)) dp(w_k)) \\ &\quad \text{which is valid because for } k \in \{0, \dots, n-1\}, J_k(x) \text{ is continuously differentiable and assumption 1 implies the existence and continuity of } D_k^{f;x}(x_k; x_k(x_k); w_k). \text{ Thus, } r_k Q_k(x_k; x_k(x_k)) \text{ is continuous in a neighborhood of } x_k \text{ for all } x_k \in \mathbb{R}^N. \text{ Then, for } k \in \{0, \dots, n-1\}, \\ &= \int_{\mathbb{R}^N} r_k E_{x_0}[J_0(x_0)] = \int_{\mathbb{R}^N} r_k Q_k(x_k; x_k(x_k)) dp(w_{k-1}) \dots dp(w_0) dp(x_0); \end{aligned}$$

Now, note that $N(x_k)$ is nonempty, closed, and convex [32]. By the definition of a stationary point of DP, we have $r_k Q_k(x_k; x_k(x_k)) \in N(x_k)$ for all $x_k \in \mathbb{R}^N$. To prove by contradiction, assume that $r_k E_{x_0}[J_0(x_0)] \notin N(x_k)$. Let a_k denote the dimension of $N(x_k)$. By the separating hyperplane theorem, there exist $\rho \in \mathbb{R}^{a_k}$ and $\gamma \in \mathbb{R}$ such that

$$\rho^T r_k Q_k(x_k; x_k(x_k)) < \gamma < \rho^T r_k E_{x_0}[J_0(x_0)];$$

for all $x_k \in \mathbb{R}^N$. Then, observe that

$$\begin{aligned} &\rho^T r_k E_{x_0}[J_0(x_0)] \\ &= \rho^T \int_{\mathbb{Z}} \int_{\mathbb{R}^N} r_k Q_k(x_k; x_k(x_k)) dp(w_{k-1}) \dots dp(x_0) \\ &= \int_{\mathbb{R}^N} \int_{\mathbb{Z}} \rho^T r_k Q_k(x_k; x_k(x_k)) dp(w_{k-1}) \dots dp(x_0) \\ &< \int_{\mathbb{R}^N} \rho^T r_k E_{x_0}[J_0(x_0)] dp(w_{k-1}) \dots dp(x_0) \\ &= \rho^T r_k E_{x_0}[J_0(x_0)]; \end{aligned}$$

which is a contradiction. Thus, $r_k E_{x_0}[J_0(x_0)] \in N(x_k)$, which shows that (x_0, \dots, x_{n-1}) is a stationary point of the one-shot optimization. ■

C. From one-shot optimization to DP

In this subsection, we first show that a local minimizer (stationary point) of the one-shot optimization does not necessarily correspond to a local minimizer (stationary point) of DP; the converse of Theorem 7 and that of Theorem 8 do not hold. Then, the optimization landscape of the one-shot optimization is more complex than that of DP. In other words, if the one-shot problem has low complexity, so does the DP problem.

To provide a counterexample, we use the basic parameterized policy that follows Definition 13: $x_k(x) = a_k x + b_k$, where $k = (a_k, b_k)$. Consider the 2-step problem

$$\begin{aligned} x_0 &= 0; c_0(x; x_0(x)) = 0; \\ f_0(x; x_0(x); w_0) &= x + a_0 x + b_0 + w_0; \\ c_1(x; x_1(x)) &= \frac{1}{4}(a_1 x + b_1)^4 - \frac{1}{2}(a_1 x + b_1)^2 + x^2; \\ f_1(x; x_1(x); w_1) &= x + a_1 x + b_1 + w_1; \\ c_2(x; x_2(x)) &= 0 \text{ where } w_0, w_1 \stackrel{iid}{\sim} \text{Uniform} \left(\frac{5}{3}; \frac{5}{3} \right); \end{aligned}$$

Fig. 3. Landscape of the one-shot optimization for a stochastic parameterized problem: b_0 is fixed to 0 in the figure. $(a_1; b_1) = (1; 0)$, $(0; 1)$ are strict local minimizers of the one-shot optimization but only $(0; 1)$ is a local minimizer of DP.

where $\alpha = [2; 2] [2; 2]$. The associated one-shot problem can be written as

$$\min_{b_0; a_1; b_1} E_{w_0} \frac{1}{4} a_1 (b_0 + w_0) + b_1 g^4$$

$$\frac{1}{2} a_1 (b_0 + w_0) + b_1 g^2 + (b_0 + w_0)^2$$

It turns out that there are 9 stationary points of the one-shot optimization in the interior of $(b_0; a_1; b_1) = (0; 0; 1)$, $(0; 1; 0)$, $(0; 0; 0)$. Among them, there are 4 strict local minimizers of the one-shot optimization: $(0; 1; 0)$, $(0; 0; 1)$. On the other hand, considering $r_{c_1}(x; \alpha(x)) = [g(x; a_1; b_1); g(x; a_1; b_1)]$ where $g(x; a_1; b_1) = (a_1 x + b_1)(a_1 x + b_1 - 1)(a_1 x + b_1 + 1)$, there are 3 stationary points of DP: $(0; 0; 1)$, $(0; 0; 0)$ and 2 strict local minimizers of DP: $(0; 0; 1)$. This verifies that a local minimizer (stationary point) of DP is indeed a local minimizer (stationary point) of the one-shot optimization but not the other way around. Fig. 3 shows the landscape of the one-shot optimization when b_0 is fixed to 0.

Now, we present the specific case that a local minimizer of the one-shot optimization implies a local minimizer of DP, similar to Theorem 6. The preconditions of theorems are similar in the sense that they both consider the case when DP has a very low complexity in the sense that there is no spurious local minima at each step of DP. The main difference between the theorems comes from whether we consider a single trajectory or the expectation over infinitely many trajectories. We consider this in the view of stationarity. (see Definitions 6, 11, and 12)

Assumption 3: There exists only a single stationary control policy $\alpha = (\alpha_0; \dots; \alpha_{n-1})$ which is also a locally minimum control policy in the interior of A for all $x \in \mathbb{R}^N$. The parameterized policy defined by Definition 13 contains with the associated parameters denoted by $\alpha = (\alpha_0; \dots; \alpha_{n-1})$.

Theorem 9: Assume that Assumption 3 holds. Consider a local minimizer of the one-shot optimization $\alpha = (\alpha_0; \dots; \alpha_{n-1})$ in the interior of A . If x_k is a continuous random variable for all $k \in \{0; \dots; n-1\}$, where $(x_0; \dots; x_n)$ is the random state process associated with α , then α is a local minimizer of DP.

Proof: Since α is a single locally minimum control policy, $Q_k(x; u)$ has no spurious local minima for all $k \in \{0; \dots; n-1\}$.

$f_0; \dots; f_{n-1}$. Thus, by Proposition 1, the corresponding $\inf_{x \in \mathbb{R}^N} Q_k(x) = 1 > 0$ makes α_0 be a local minimizer of DP.

Consider a stationary point of the one-shot optimization $\alpha = (\alpha_0; \dots; \alpha_{n-1})$ in the interior of A . We will now prove by a backward induction that should always be $\alpha_k = \alpha_k^0$ for all $k \in \{0; \dots; n-1\}$.

For the base step, at step $n-1$, since the parameterized policy contains α_{n-1} , it can be expressed as $\alpha_{n-1}(x) = \prod_{i=1}^m s_i f_i(x)$, where $f_i: \mathbb{R}^N \rightarrow \mathbb{R}^M$, $i = 1; \dots; m$, $f_i(x) = [f_{i1}(x); \dots; f_{iM}(x)]^T$ and $\alpha_{n-1} = (s_1; \dots; s_m)$. Notice that $Q_{n-1}(x; \alpha_{n-1}(x)) = Q_{n-1}(x; \alpha_{n-1}^0(x))$ since α_{n-1}^0 is the final parameter of the whole system to determine the control inputs and the state transition. Now, observe that

$$r_{c_1} Q_{n-1}(x; \alpha_{n-1}(x)) = r_{c_1} Q_{n-1}(x; \alpha_{n-1}^0(x)) = 0$$

and $\alpha_{n-1}^0(x)$ is the unique solution for $Q_{n-1}(x; \alpha_{n-1}^0(x)) = 0$ since $\alpha_{n-1}^0(x)$ is the unique stationary point located within the interior of A due to Assumption 3. This yields the following expression with $\alpha(x) = (u_1; \dots; u_M)^T$:

$$r_{c_1} Q_{n-1}(x; \alpha(x)) = r_{c_1} \prod_{i=1}^m (u_1 \dots u_M s_i f_{iM}(x)) g_M(x; \alpha(x)) = 0$$

where $g_j(x; \alpha(x)); j = 1; \dots; M$, are nonnegative at $\alpha = \alpha_{n-1}^0$ and positive at all the other points since α_{n-1}^0 is a local minimizer of DP that yields the unique stationary control policy.

Now, let $\alpha_{n-1}(x)$ be $\prod_{i=1}^m d_i f_i(x)$ where $\alpha_{n-1} = (d_1; \dots; d_m)$. Observe that according to the chain rule, the following expression holds:

$$r_{c_1} Q_{n-1}(x; \alpha_{n-1}(x)) = r_{c_1} \prod_{i=1}^m (d_i s_i f_{i1}(x)) g_1(x; \alpha_{n-1}(x)) \dots \prod_{i=1}^m (d_i s_i f_{iM}(x)) g_M(x; \alpha_{n-1}(x)) = 0$$

Now, notice that $r_{c_1} E_{x_0} [J_0(x_0)] = 0$ since α is a stationary point of the one-shot optimization in the interior of A . Then, observe that

$$r_{c_1} E_{x_0} [J_0(x_0)] = r_{c_1} E_{x_0; w_0; \dots; w_{n-2}} [Q_{n-1}(x_{n-1}; \alpha_{n-1}(x_{n-1}))] = E_{x_0; w_0; \dots; w_{n-2}} [r_{c_1} Q_{n-1}(x_{n-1}; \alpha_{n-1}(x_{n-1}))] = 0$$

The second equality comes from $\alpha_{n-1}(x; \alpha(x))$ being differentiable with respect to the parameters due to the linearity of the policy defined by Definition 13. Now, we substitute (17) into (18) to derive an m -dimensional vector equation, and multiply $(d_k - s_k)$ with k^{th} component as follows:

$$E_{x_0; w_0; \dots; w_{n-2}} \prod_{j=1}^m (d_k - s_k) f_{kj}(x_{n-1}) \prod_{i=1}^m (d_i - s_i) f_{ij}(x_{n-1}) g_j(x_{n-1}; \alpha_{n-1}(x_{n-1})) = 0$$

for all $k = 1; \dots; m$: We sum up them equations and rearrange the terms to derive the following equation:

$$\sum_{j=1}^M E_{x_0; w_0; \dots; w_{n-2}} \sum_{i=1}^n (d_i - s_i) f_{ij}(x_{n-1}) g_j(x_{n-1}; w_{n-1}) = 0: \quad (19)$$

The term inside the expectation is always nonnegative regardless of the distribution of $x_0; w_0; \dots; w_{n-2}$. Now, suppose that $w_{n-1} \in \mathbb{R}^m$; i.e., $d_i \in s_i$ for some $i \in \{1; \dots; m\}$. Then, we have $g_j(x_{n-1}; w_{n-1})$ to be strictly positive. As a result, for (19) to be satisfied, $\sum_{i=1}^m (d_i - s_i) f_{ij}(x_{n-1})$ should be 0 for every $j \in \{1; \dots; M\}$ for all possible values of x_{n-1} . Recall from Remark 8 to note that it is impossible to satisfy (19) since x_{n-1} is a continuous random variable and the policy is defined by Definition 13, i.e., a linear combination of some independent basis functions. Thus, $d_i = s_i$ for all $i \in \{1; \dots; m\}$, which means $w_{n-1} = \mathbf{0}$.

For the induction step, assume that $w_k = \mathbf{0}$. Again, $Q_k^0(x; w_k(x)) = Q_k(x; w_k(x))$ holds, and thus one can apply the same logic as the base step to obtain $w_{k+1} = \mathbf{0}$. Thus, $w_k = \mathbf{0}$ holds, which implies that w_k is a local minimizer of DP since $\mathbf{0}$ is a local minimizer of DP. ■

Remark 10: The results of both Theorems 6 and 9 state that a local minimizer of the one-shot optimization is indeed a local minimizer of DP under the common assumption that no spurious local minima exist at each step of DP. By taking the contrapositive, one can observe that under such a condition, there is at most one local minimizer of the one-shot optimization, indicating that no spurious local minima exist; i.e., if DP has a very low complexity, the same holds for the one-shot problem.

Remark 11: To determine the form of $Q_{n-1}^0(x; u)$, it was necessary to argue that $w_{n-1}(x)$ should be the unique solution for $r = Q_{n-1}^0(x; u) = 0$. For this to be true, there should certainly be only a single stationary control policy, which necessitates Assumption 3. In fact, it may be difficult to satisfy the precondition that a single stationary control policy should be in the interior of \mathcal{A} for all $x \in \mathbb{R}^N$. Instead, we can relax this condition to apply only within the domain of the set of values that at least one of the states $x_1; \dots; x_n$ can take. For example, if the state space is finite, satisfying the condition becomes relatively straightforward.

Remark 12: The challenging part of a backward induction in the proof arises from the fact that the state at step k is determined by the previous steps but one cannot look at the previous steps in the backward induction. Thus, the main idea of the proof leverages equation (19), which incurs the fact that $w_k = \mathbf{0}$ regardless of the distribution of $x_0; w_0; \dots; w_{k-1}$. Thus, we only need the assumption that there is a single stationary control policy with respect to the given distribution of $x_0; w_0; \dots; w_{n-1}$. This is a big improvement from the work [28] (see Condition 4 of Section 5.4) in the sense that Condition 4 needs no sub-optimal stationary point with respect to any possible distribution.

Now, we present the pictorial example of Theorem 9. Consider the example of 2-step problem presented above

Fig. 4. Landscape of the one-shot optimization under the assumptions of Theorem 9: b_0 is fixed to 0 in the figure. $(a_1; b_1) = (1; 0.5)$ is the only stationary point (local minimizer) of DP and also the only stationary point (local minimizer) of the one-shot optimization.

Fig. 3, but modify $c_1(x; w_1(x))$ to $\frac{1}{4}(a_1 x + b_1 - x - 0.5)^4 + x^4$. The associated one-shot problem can be written as

$$\min_{b_0; a_1; b_1} E_{w_0} \frac{1}{4} f(a_1 - 1)(b_0 + w_0) + b_1 - 0.5 g^4 + (b_0 + w_0)^4$$

$(a_1; b_1) = (0; x + 0.5)$ is the only locally minimum control policy, and the parameterized policy class contains this policy as $(b_0; a_1; b_1) = (0; 1; 0.5)$. Clearly, it is a local minimizer of DP. It turns out that the corresponding one-shot problem also has a single stationary point $(0; 1; 0.5)$, which is also a local minimizer of the one-shot optimization. Fig. 4 shows the landscape of the one-shot optimization where b_0 is fixed to 0. Considering both Theorems 7 and 9, one can conclude that under the assumptions of Theorem 9, a local minimizer of DP is equivalent to a local minimizer of the one-shot optimization.

D. Numerical Experiments

In this subsection, we will present a high-dimensional experiment on the classical linear quadratic regulator (LQR):

$$\begin{aligned} f_k(x_k; u_k) &= A_k x_k + B_k u_k; \quad k = 0; \dots; n-1; \quad x_0 \in \mathcal{D}; \\ Q_k(x_k; u_k) &= x_k^T Q_k x_k + u_k^T R_k u_k; \quad k = 0; \dots; n-1; \\ U_k &= K_k x_k; \quad k = 0; \dots; n-1; \quad c_n(x_n) = x_n^T Q_n x_n; \end{aligned}$$

whose goal is to find the optimal parameters $K_0; \dots; K_{n-1}$. To solve the problem using DP, we use Xpress Optimizer v9.3.0 [40]. To solve the problem in a one-shot fashion, we use Gurobi Optimizer v11.0.0 [41] with the tolerance of 10^{-4} .

Let $K_{k; DP}$ and $K_{k; OS}$ denote an observed local solution of the k^{th} step parameter K_k obtained by DP and the one-shot problem, respectively. We aim to determine whether each local solution of DP corresponds to some local solutions of the one-shot problem, and vice versa. One can verify this by first solving DP or the one-shot problem and then providing its solution as the initial parameter values when solving its counterpart. This method is often referred to as ‘‘warm start.’’ We expect to observe unchanged values from the initial guess if there is indeed a correspondence. Let $(K_{k; DP}, OS)$ and $(K_{k; OS}, DP)$ denote the solution of K_k obtained by the one-shot problem (DP) using warm start with DP (one-shot) solution as an initial guess. The initial distribution \mathcal{D} introduces the stochasticity to the system and induces the states to be continuous random variables, which obeys the assumption of Theorem 9.

TABLE II
RELATIONSHIP BETWEEN DP AND ONE-SHOT SOLUTIONS OF LQR

Numerical difference		Scenario	(a) Unconstrained	(b) Constrained
$kK_{0,DP}$	$K_{0,DP}$	OS $k_F = kK_{0,DP}k_F$	0	0
$kK_{1,DP}$	$K_{1,DP}$	OS $k_F = kK_{1,DP}k_F$	0	0
$kK_{0,OS}$	$K_{0,OS}$	DP $k_F = kK_{0,OS}k_F$	$2.901 \cdot 10^{-6}$	$5.399 \cdot 10^{-4}$
$kK_{1,OS}$	$K_{1,OS}$	DP $k_F = kK_{1,OS}k_F$	$2.291 \cdot 10^{-5}$	3.156

We perform 20 experiments for $x_k \in \mathbb{R}^3$ and $u_k \in \mathbb{R}^4$ with $n = 30$. We randomly generate A_k and B_k , whose entries are all in $[-100; 100]$. We also generate $Q_k = QQ^T$ and $R_k = RR^T + 100I$, where all entries of Q and R are in $[-20; 20]$ and I denotes the identity matrix. D is the normal distribution with the expectation $20\mathbb{1}$ and the variance VV^T , where $\mathbb{1}$ denotes the vector of ones and all entries in V are in $[-200; 200]$. We consider two scenarios: (a) unconstrained and (b) constrained by the last-step (nonconvex) condition $K_{n-1}^T \preceq K_{n-1} - 10000I$, where \preceq denotes the Loewner partial ordering (roughly speaking, this condition ensures that the controller has a high gain). Table II shows whether the correspondence holds between the solutions of DP and the one-shot problem using warm start under the two scenarios, presenting the results of the average of 20 experiments.

It turns out that for both scenarios, one can observe that a solution of DP corresponds to each solution of one-shot problem since the one-shot solver directly identifies DP solution as a local solution of the one-shot problem without any numerical update. This implication supports the findings of Theorem 7.

However, whether a one-shot solution implies some DP solutions depends on the problem setting. Our experiment for the unconstrained case shows that a solution of the one-shot problem indeed corresponds to that of DP, implying with the above result that a one-shot solution is equivalent to a DP solution. Previous studies have shown the global convergence of this one-shot problem by proving that LQR satisfies the gradient dominance property, even though the problem is generally nonconvex [42], [43]. Our general approach alternatively observes the DP counterparts: Since every DP sub-problem of LQR has no spurious local minima, our experiment implies that the one-shot LQR problem also has none of them and achieves the global convergence, which supports Theorem 9.

On the other hand, the nonconvex constraint on K_{n-1} creates spurious solutions for the $(n-1)$ th DP, independent of whether the $(n-2)$ th; \dots ; 0 th DP steps have any spurious local minima. Our experiment for the constrained case shows that having spurious local minima at the $(n-1)$ th DP propagates backward to K_1 , where the solver fails to guarantee that an observed local minimum of the one-shot optimization corresponds to that of DP. This illustrates that the landscape of the one-shot problem has a higher complexity than its DP counterpart, which was also shown in Fig. 2b and Fig. 3. This result serves as a counterexample of the converse of Theorems 7 and 8, and it implies that a single high-complexity DP step affects the landscape of the one-shot problem.

V. CONCLUSION

In this paper, we studied the optimization landscape of the optimal control problems via two different formulations: one-shot optimization aimed at solving for all input values at the same time, and DP method aimed at finding the input values sequentially. For the deterministic problem, we proved that under some mild conditions, each local minimizer of the one-shot optimization corresponds to an input sequence induced by some locally minimum control policy of DP, and vice versa.

To help better understand the quality of the local solutions obtained by reinforcement learning algorithms, we incorporated exact parameterized policies into the optimal control problem for both deterministic and stochastic dynamics. We showed that if the one-shot problem has a low complexity, so do the corresponding DP sub-problems, indicating the success of DP methods. Moreover, under the condition that there exists only a single locally minimum control policy, with different technical assumptions, both deterministic and stochastic cases yield that a local minimizer of the one-shot optimization is equivalent to a local minimizer of DP.

We focused on the discrete-time finite-horizon optimal control problem in this work. A natural future direction would be to extend this work to the continuous-time and infinite-horizon cases, which was discussed in Remark 4. For safety-critical systems, state constraints may also be enforced, with recursive feasibility being crucial to guarantee the success of DP. One may also want to extend the parameterized policy class beyond a linear combination of basis functions, such as composite functions widely used in deep neural networks.

REFERENCES

- [1] Y. Ding, Y. Bi, and J. Lavaei, "Analysis of spurious local solutions of optimal control problems: One-shot optimization versus dynamic programming," in *American Control Conference (ACC)*. IEEE, 2021.
- [2] R. E. Bellman, *Dynamic programming*. Princeton university press, 1957.
- [3] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Two Volume Set*, 4th ed. Athena scientific, MA, 2012.
- [4] R. Bellman and R. Kalaba, "On the role of dynamic programming in statistical communication theory," *IRE Transactions on Information Theory*, vol. 3, no. 3, pp. 197–203, 1957.
- [5] E. A. Feinberg, "Optimality conditions for inventory control," in *Optimization Challenges in Complex, Networked and Risky Systems*. INFORMS, 2016, pp. 14–45.
- [6] I. Kolmanovskiy, I. Siverguina, and B. Lygoe, "Optimization of power-train operating policy for feasibility assessment and calibration: Stochastic dynamic programming approach," in *American Control Conference (ACC)*, vol. 2. IEEE, 2002, pp. 1425–1430.
- [7] D. P. Bertsekas, *Reinforcement learning and optimal control*. Athena Scientific Belmont, MA, 2019.
- [8] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [9] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [10] M. Soltanolkotabi, A. Javanmard, and J. D. Lee, "Theoretical insights into the optimization landscape of over-parameterized shallow neural networks," *IEEE Transactions on Information Theory*, vol. 65, no. 2, pp. 742–769, 2019.
- [11] C. Liu, L. Zhu, and M. Belkin, "Loss landscapes and optimization in over-parameterized non-linear systems and neural networks," *Applied and Computational Harmonic Analysis*, vol. 59, pp. 85–116, 2022.

