

CACTO: Continuous Actor-Critic with Trajectory Optimization—Towards global optimality

Gianluigi Grandesso¹, Elisa Alboni², Gastone P. Rosati Papini³, *Member, IEEE*, Patrick M. Wensing⁴, *Member, IEEE* and Andrea Del Prete⁵, *Member, IEEE*

Abstract—This paper presents a novel algorithm for the continuous control of dynamical systems that combines Trajectory Optimization (TO) and Reinforcement Learning (RL) in a single framework. The motivations behind this algorithm are the two main limitations of TO and RL when applied to continuous nonlinear systems to minimize a non-convex cost function. Specifically, TO can get stuck in poor local minima when the search is not initialized close to a “good” minimum. On the other hand, when dealing with continuous state and control spaces, the RL training process may be excessively long and strongly dependent on the exploration strategy. Thus, our algorithm learns a “good” control policy via TO-guided RL policy search that, when used as initial guess provider for TO, makes the trajectory optimization process less prone to converge to poor local optima. Our method is validated on several reaching problems featuring non-convex obstacle avoidance with different dynamical systems, including a car model with 6D state, and a 3-joint planar manipulator. Our results show the great capabilities of CACTO in escaping local minima, while being more computationally efficient than the Deep Deterministic Policy Gradient (DDPG) and Proximal Policy Optimization (PPO) RL algorithms.

Index Terms—Trajectory optimization, reinforcement learning, continuous control.

I. INTRODUCTION

When a model of the system to be controlled is available, one of the most common and flexible techniques to compute optimal trajectories is gradient-based Trajectory Optimization (TO). Starting from a specific initial state, TO can find the control sequence that minimizes a cost function representing the task to be accomplished, where the system’s dynamics and possible state and control limits are considered as constraints. Such a powerful framework has led to excellent results when the problem is convex or slightly non-convex, especially when used in a Model Predictive Control (MPC) fashion. For example, it has been successfully employed to control

high-dimensional nonlinear systems such as quadrupeds and humanoids [1]–[4]. However, when the task to be accomplished requires a highly non-convex cost function and/or the dynamics is highly nonlinear, the presence of multiple local minima, some of which are of poor quality (i.e., associated to a cost that is significantly worse than the global minimum), often prevents TO from finding a satisfying control trajectory. A possible solution to this problem is providing TO with a good initial guess, which turns out to be very complex—not to say impossible—without a deep prior knowledge about the system, which is often unavailable in practice. The IREPA algorithm [5] tackles this problem by building a kinodynamic Probabilistic Road Map (PRM) and approximating the *Value function* and control policy, which is then used to warm-start an MPC. The method produced satisfying results on a 3-DoF system, but it is limited to problems with a fixed terminal state, and scaling to high dimensions seems not trivial due to its need to explicitly store the locally optimal trajectories in the PRM edges.

Approaches that can find the global optimum of non-convex problems exist, and are based on the Hamilton-Jacobi-Bellman equation [6] (for continuous-time problems) or Dynamic Programming (for discrete-time problems). However, these methods suffer from the curse of dimensionality, which restricts their applicability to systems with extremely few degrees of freedom. Some efficient solutions exist, but for specific problems, such as simple integrator dynamics with control-independent cost [7], [8]. Alternatively, the tensor-train decomposition [9] has been used to reduce the computational complexity of *value iteration*, by representing the *Value function* in a compressed form. However, the approach has been tested on stochastic optimal control problems with at most 7-dimensional state and scalar control. Subsequent improvements of that approach [10] could solve problems with a 12-dimensional state, but required some restrictions on the form of the dynamics and cost.

On the other hand, in the recent years deep RL has shown impressive results on continuous state and control spaces, which represented its greatest challenge until 2015, when the ground-breaking algorithm DDPG [11] was presented. Successively, many variants have been developed to further improve its performance, such as TD3, SAC, and RTD3 [12]–[14]. In 2021, Zhang *et al.* [15] managed to speed up DDPG’s training time and enhance the effectiveness of its learning with their expansion called AE-DDPG, by making the agent latch on “good” trajectories very soon through the use of asynchronous episodic control and improving exploration with

Manuscript received: November 10th, 2022; Revised: February 15th, 2023; Accepted: March 23rd, 2023.

This paper was recommended for publication by Editor Lucia Pallottino upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by the Italian MUR through the “Departments of Excellence” programme and by PRIN project DOCEAT (CUP n. E63C22000410001).

¹, ², ³ and ⁵ are with the Dept. of Industrial Engineering, University of Trento, Italy [gianluigi.grandesso; elisa.alboni; gastone.rosatipapini; andrea.delprete]@unitn.it

⁴ is with the Dept. of Aerospace and Mechanical Engineering, University of Notre Dame, Indiana, USA pwensing@nd.edu

Digital Object Identifier (DOI): see top of this page.

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

a new type of control noise. However, deep RL is intrinsically limited by its low sample efficiency, which implies the need for a considerable number of interactions with the environment to reach a good performance level.

To mitigate this problem, Levine *et al.* [16] proposed to guide the exploration process by using Differential Dynamic Programming (DDP) as a generator of guiding samples that push the policy search towards low-cost regions. However, the *imitation* component of this approach makes its capability to find an optimal control policy strongly dependent on the quality of the guiding samples. This downside applies also to the method proposed by Mordatch and Todorov [17] combining policy learning and TO through the Alternating Direction Method of Multipliers (ADMM), which involves imitation in that the policy learning problem is reduced to a sequence of trajectory optimization and regression problems. In general, all the methods that are *imitation-oriented* (e.g. [18]–[20]) suffer from the same limitation: the quality of what the RL algorithm can learn is limited by the quality of the demonstrations or guiding trajectories found by TO.

In a similar spirit to [16], but with an imitation-free approach, our algorithm aims to mitigate both problems: the local minima issue affecting TO, and the low sample efficiency of RL. Our main contribution is an algorithm named CACTO (Continuous Actor-Critic with TO). The algorithm combines TO and RL in such a way that their interplay guides the search towards the globally optimal control policy. Our main contributions are:

- 1) We present a novel RL algorithm that exploits TO to speed up the search, and that (contrary to previous work) does not rely on imitation.
- 2) Our tests show that initializing TO with the CACTO policy outperforms standard warm-starting techniques.
- 3) We have proved that, considering a discrete-space version of CACTO with look-up tables instead of deep neural networks (DNN), the policy approaches global optimality as the algorithm proceeds.

II. METHOD

This section presents an optimization algorithm to solve a finite-horizon discrete-time optimal control problem (OCP) that takes the following general form:

$$\underset{X,U}{\text{minimize}} \quad J(X,U) = \sum_{k=0}^{T-1} l_k(x_k, u_k) + l_T(x_T) \quad (1a)$$

$$\text{subject to} \quad x_{k+1} = f_k(x_k, u_k) \quad \forall k = 0 \dots T-1 \quad (1b)$$

$$|u_k| \in \mathcal{U} \quad \forall k = 0 \dots T-1 \quad (1c)$$

$$x_0 = x_{init} \quad (1d)$$

where the state and control sequences $X = x_{0..T}, U = u_{0..T-1}$, with $x_k \in \mathbb{R}^n$ and $u_k \in \mathbb{R}^m$, are the decision variables. The cost function $J(\cdot)$ is defined as the sum of the running costs $l_k(x_k, u_k)$ and the terminal cost $l_T(x_T)$. The dynamics, control limits and initial conditions are represented by (1b), (1c) and (1d), with $\mathcal{U} = \{u \in \mathbb{R}^m : |u| \leq u_{max}\}$.

Our algorithm combines ideas from RL and TO. The core idea is to exploit TO to guide the exploration towards low-cost

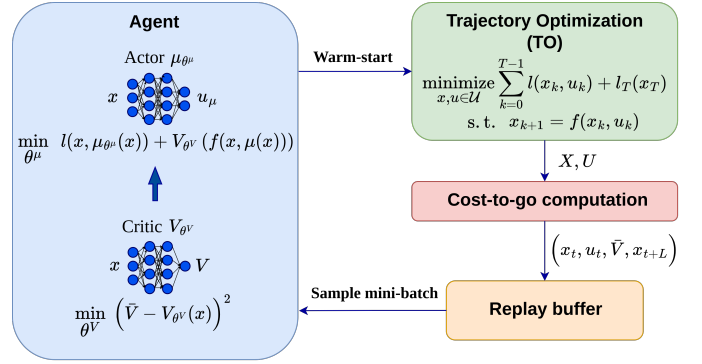


Fig. 1. Scheme of CACTO: at each episode, a TO problem is warm-started with a policy rollout and solved, then we compute the cost-to-go associated with each state of the trajectory found by TO and store the related transition in the replay buffer. Finally, we update critic and actor by sampling a mini-batch of transitions from the replay buffer.

regions, so as to make the learning process more efficient. In turn, a rollout of the currently learned policy is used to initialize the next TO problem. In this way, the RL and TO components of the algorithm help each other, making convergence faster: as training proceeds, the TO solver is provided with better initial guesses and this increases the chance of obtaining better solutions; in turn, these trajectories drive the agent along lower-cost paths, pushing the critic towards a better approximation of the optimal *Value function* and, consequently, the actor towards the optimal policy. Fig. 1 shows a scheme of CACTO.

A. TO and RL Notation

Besides the control sequence $[u_0, \dots, u_{T-1}]$, we also make use of the term “*policy*” $\pi(x)$, that belongs to the RL community and refers to the mapping from states x to controls u (“actions” in the RL language). In the optimal control (OC) community, the *Value function* of a policy π describes the cost-to-go following π from x at time t , namely:

$$V_t^{\pi}(x_t) = \sum_{k=t}^{T-1} l_k(x_k, \pi(x_k)) + l_T(x_T)$$

In RL, the *Value function* is the sum of the rewards starting from a state x at time t and following a policy π ; therefore, we aim to maximize it, opposite to what we do in OC. For reasons of clarity and to keep the notation simple, from now on we always refer to the *Value function* $V(x_t)$ in an OC sense. The same reasoning applies to the *Q-value function* $Q(x_t, u_t)$, which, in RL, is given by the sum of the rewards after applying u from x at time t and following π thereafter, whereas in the OC community, it represents the cost-to-go after performing u_t from x_t and following π thereafter:

$$Q_t^{\pi}(x_t, u_t) = l_t(x_t, u_t) + \sum_{k=t+1}^{T-1} l_k(x_k, \pi(x_k)) + l_T(x_T)$$

In the following, to simplify the notation, we consider time as the last component of the state vector ($x[n] = t$), so that we can drop the explicit time dependency for the *Value function*, policy, dynamics, and running cost.

B. Algorithm Description

As shown in Fig. 1, our algorithm can be broken down into four phases. In the TO phase (green block), the optimal state and control trajectories are computed considering a random initial state. Then, the cost-to-go is computed (red block) for each state of the optimal trajectory and stored in a buffer (orange block). This buffer is then sampled to update both the critic and actor (blue block). Finally, to close the loop, a rollout of the actor is used to warm-start the next TO problem.

A more detailed description is reported in Algorithm 1. In the initial phase (lines 1-4), the critic $V(x|\theta^V)$ and actor $\mu(x|\theta^\mu)$ networks (“Agent” blue block in Fig. 1) are initialized, as well as the target critic network $V'(x|\theta^{V'})$ (copying the weights of the critic), and a buffer R is created. This buffer will store the transitions $(x_t, u_{TO,t}, \hat{V}_t, x_{\min(t+L+1,T)})$, where \hat{V}_t is the partial cost-to-go until state $x_{\min(t+L+1,T)}$ after a rollout of either L steps, or the number of steps to reach T (whichever is lower). In this way, to update the critic one can choose between n-step Temporal Difference (TD) ($L = n - 1$) and Monte-Carlo ($L = T - 1 - t$) by setting the hyperparameter L [23]. After the initialization phase, M episodes are performed (lines 5-29) starting from a random initial state x_0 (line 6). The state vector has a dimension of n and includes the time as its last component. The starting time index for each episode (and partial cost-to-go computation) is based on $x_0[n]$ (lines 14 and 17), so the length of each episode can vary. At the beginning of each episode, the state and control variables $(x_{TO}^\circ, u_{TO}^\circ)$ of a TO problem are initialized (line 11) with a rollout $(x_\mu^\circ, u_\mu^\circ)$ of the policy network $\mu(x|\theta^\mu)$ (lines 7-10, “Warm-start” arrow in Fig. 1). The TO problem is then solved, the control inputs u_{TO} are applied starting from x_0 , and the resulting costs are computed and saved (lines 14-16). Then, for each step of the episode the partial cost-to-go \hat{V}_t is computed and the related transition is saved in R (lines 17-20, from “Cost-to-go computation” block to “Replay Buffer” one in Fig. 1). Finally, every e_{update} episodes the critic and the actor are updated: for K times a minibatch of S transitions is sampled from R (line 23, “Sample minibatch” arrow in Fig. 1), for each of them the complete cost-to-go \bar{V}_i is computed (line 24) by either adding the tail of the Value to \hat{V}_i or copying the partial cost-to-go, which is the cost-to-go itself in case the lookahead window exceeds the episode length T . Then the critic and actor loss functions are minimized (lines 24-25). Considering that CACTO deals with finite-horizon problems, it is worth noting that V' is used only when $i + L + 1 \neq T$, namely when the costs-to-go are computed with n-step TD. The critic and actor loss functions are respectively the mean squared error between the costs-to-go and the values predicted by the critic $(\bar{V}_i - V(x_i))^2$ and the Q-value, namely $Q(x_t, u_t) = l(x_t, u_t) + V(x_{t+1})$, which represents the policy’s performance.

C. Differences with respect to DDPG

Our approach takes inspiration from DDPG, but presents a few key differences, which we highlight in this subsection. A first difference is that we replaced the Q -value with the Value function, which is approximated by the critic network.

Algorithm 1: CACTO

```

1 Inputs: dynamics  $f(\cdot, \cdot)$ , running cost  $l(\cdot, \cdot)$ , terminal
   cost  $l_T(\cdot)$ ,  $T$ ,  $M$ ,  $S$ ,  $K$ ,  $L$ 
2 Output: trained control policy  $\mu(x)$ 
3  $\theta^V \leftarrow$  random,  $\theta^\mu \leftarrow$  random,  $\theta^{V'} \leftarrow \theta^V$ 
4 Initialize replay buffer  $R$ 
5 for  $episode \leftarrow 1$  to  $M$  do
6    $x_0 \leftarrow$  random,  $x_{TO,0}^\circ \leftarrow x_0$ 
7   for  $t \leftarrow x_0[n]$  to  $T$  do (policy rollout)
8      $u_{\mu,t}^\circ \leftarrow \mu(x_{\mu,t}^\circ | \theta^\mu)$ 
9      $x_{\mu,t+1}^\circ \leftarrow$  Environment( $x_{\mu,t}^\circ, u_{\mu,t}^\circ$ )
10  end
11   $(x_{TO}^\circ, u_{TO}^\circ) \leftarrow (x_\mu^\circ, u_\mu^\circ)$  (TO warm-start)
12  Solve TO problem and get control trajectory  $u_{TO}$ 
13  Agent’s initial state  $\leftarrow x_0$ 
14  for  $t \leftarrow x_0[n]$  to  $T$  do (episode rollout)
15     $x_{t+1}, l_t \leftarrow$  Environment( $x_t, u_{TO,t}$ )
16  end
17  for  $t \leftarrow x_0[n]$  to  $T$  do
18    Compute partial cost-to-go:  $\hat{V}_t = \sum_{j=t}^{\min(t+L, T-1)} l_j$ 
19     $R \leftarrow (x_t, u_{TO,t}, \hat{V}_t, x_{\min(t+L+1, T)})$ 
20  end
21  if  $episode \% e_{update} = 0$  then
22    for  $k \leftarrow 1$  to  $K$  do (critic & actor update)
23      Sample minibatch of  $S$  transitions
        ( $x_i, u_{TO,i}, \hat{V}_i, x_{\min(i+L+1, T)}$ )
24      Compute cost-to-go:
        
$$\bar{V}_i = \begin{cases} \hat{V}_i & \text{if } i + L + 1 > T \\ \hat{V}_i + V'(x_{i+L+1}) & \text{otherwise} \end{cases}$$

        Update critic by minimizing the loss over  $\theta^V$ :
        
$$L_c = \frac{1}{S} \sum_{i=1}^S (\bar{V}_i - V(x_i | \theta^V))^2$$

25      Update actor by minimizing the loss over  $\theta^\mu$ :
        
$$L_a = \frac{1}{S} \sum_{i=1}^S Q(x_i, \mu(x_i | \theta^\mu))$$

26      Update target critic:  $\theta^{V'} \leftarrow \tau \theta^V + (1 - \tau) \theta^{V'}$ 
27    end
28  end
29 end

```

By doing so, the complexity is reduced since the critic’s input is only the state, thus there is no need to explore the control space. Keeping a Q -value approach instead, one could not explore the control space using only TO because only the locally-optimal control inputs would be chosen. Therefore, one should alternate the use of TO trajectories with other exploration techniques (*e.g.*, acting greedily *w.r.t.* the Q -value function and adding some noise), which would dilute the benefits of our algorithm, making it more similar to standard RL. Contrary to DDPG, this algorithm is on-policy, meaning that the critic estimates the Value of the exploratory policy being followed. This is the policy obtained by initializing

TO with rollouts of the current policy network. This implies that it is important to size R not too large, so that only the most recent TO trajectories obtained by warm-starting TO with similar policies are stored; in this way, the critic, after its update, will approximate the *Value function* associated to that policy. Otherwise, the risk is having a critic that represents a meaningless *Value function*, because it used trajectories generated with very different policies.

Another difference with respect to DDPG is that we consider finite-horizon problems because TO cannot be used to solve arbitrary infinite-horizon problems. Therefore, the *Value function* is time-dependent. We address this by considering time as the last component of the state vector ($x[n] = t$).

As in DDPG, we also make use of a target network V' to improve the stability of our algorithm. It is a copy of the critic network, whose weights $\theta^{V'}$ are updated slower than the critic's ones θ^V by performing only partially the update of the critic, that is $\theta^{V'} \leftarrow \tau\theta^V + (1 - \tau)\theta^{V'}$ with $\tau \ll 1$ being the target learning rate.

D. Implementation Details

Each TO problem was solved using collocation, which was available in the *Pyomo* library [25] and solved with the nonlinear programming solver *IpOpt* [26].

The following details are not fixed as part of the proposed method. They are reported here for the sake of clarity and to help understand how the results in Section III have been obtained. Concerning the critic's neural network, we used a structure with two small preprocessing fully-connected layers with respectively 16 and 32 neurons, followed by two other fully-connected hidden layers with 256 neurons each, all with *ReLU* activation functions. For the actor instead we used a residual neural network with two fully-connected hidden layers with 256 neurons each and *ReLU* activation functions, whose outputs are added and passed as inputs to the last layer with a *tanh* activation function to bound the final controls in $[-1, +1]$. We chose a residual neural network to better propagate the gradient to the first layer and limit the effects of the potential combination of the vanishing gradient problem due to the *tanh* in the last layer and the dying ReLU problem [27] which would prevent the network from continuing to learn. We finally multiply the output of this last layer by the control upper bound. To stabilize the training of these networks, the inputs are normalized and L2 weight and bias regularizers are used in each layer, with weight equal to 10^{-2} .

The critic and actor loss functions were minimized with a stochastic gradient descent optimizer *Adam* [28]. We set the maximum number of episodes to 80000 and stopped early the training of the neural networks if the results were satisfactory.

E. Global convergence proof for discrete spaces

As for most continuous-space RL algorithms, it is hard to give any formal guarantee of convergence for CACTO. However, we can show that, considering a discrete-space version of CACTO using look-up tables instead of DNN, the algorithm converges to a globally optimal policy. This version of CACTO performs sweeps of the entire state space as in classic Dynamic Programming (DP) algorithms (*e.g.*, Policy

Iteration [24]), rather than the asynchronous approach characterizing RL algorithms. Moreover, we consider the Policy Iteration version of our algorithm meaning that each phase of policy evaluation and policy improvement converges before the other begins. This proof does not extend easily to the original CACTO algorithm, but it gives us an insight into the soundness of its key principle.

Theorem II.1. *Consider the following assumptions.*

- *State and control spaces are finite.*
- *The optimal Value function V^* is bounded.*
- *We have access to a discrete-space TO algorithm that can perform a local search and return a trajectory with a cost not greater than the cost of the initial guess.*

Let us define ${}^k\pi_{TO}$ as the control policy obtained solving a TO problem using a rollout of the policy ${}^k\pi$ as initial guess. Then, starting from an arbitrary initial policy ${}^0\pi$, the following algorithm converges to the optimal Value function V^ and an optimal policy π^* :*

$$\begin{aligned} {}^kV^{\pi_{TO}} &= \text{PolicyEvaluation}({}^k\pi_{TO}) \\ {}^{k+1}\pi &= \arg \min_{u \in \mathcal{U}} [l + {}^kV^{\pi_{TO}}] \end{aligned}$$

Proof. Let us define $\pi'(x)$ as the policy obtained by minimizing the *Action-value function* of π_{TO} , namely $Q^{\pi_{TO}}(x, u)$:

$$\pi'_t(x_t) \triangleq \arg \min_{u \in \mathcal{U}} [l_t(x_t, u) + V_{t+1}^{\pi_{TO}}(x_{t+1})] \quad \forall x_t, t \quad (2)$$

where $x_{t+1} = f_t(x_t, u)$, and $V_{t+1}^{\pi_{TO}}(x_{t+1})$ is the cost-to-go following π_{TO} from x_{t+1} at time $t + 1$.

Since TO always finds a solution that is at least as good as the provided initial guess, we have that:

$$V_t^\pi(x) \geq V_t^{\pi_{TO}}(x) \quad \forall x, t \quad (3)$$

Since π' is the minimizer of $Q^{\pi_{TO}}$ (see (2)), we know that:

$$V_t^{\pi_{TO}}(x) = Q_t^{\pi_{TO}}(x, \pi_{TO,t}(x)) \geq Q_t^{\pi_{TO}}(x, \pi'_t(x)) \quad (4)$$

Starting from (4), and following the same idea of the convergence proof of Policy Improvement [24], we can write:

$$V_t^{\pi_{TO}}(x_t) \geq \min_{u \in \mathcal{U}} [l_t(x_t, u) + V_{t+1}^{\pi_{TO}}(x_{t+1})] \quad (5)$$

$$\text{(by (2))} = l_t(x_t, \pi'_t(x_t)) + V_{t+1}^{\pi_{TO}}(x_{t+1}) \quad (6)$$

$$\text{(by (4))} \geq l_t(x_t, \pi'_t(x_t)) + Q_{t+1}^{\pi_{TO}}(x_{t+1}, \pi'_{t+1}(x_{t+1})) \quad (7)$$

$$\begin{aligned} &= l_t(x_t, \pi'_t(x_t)) + [l_{t+1}(x_{t+1}, \pi'_{t+1}(x_{t+1})) \\ &\quad + V_{t+2}^{\pi_{TO}}(x_{t+2})] \end{aligned} \quad (8)$$

...

$$\geq \sum_{k=t}^{T-1} l_k(x_k, \pi'_k(x_k)) + l_T(x_T) \triangleq V_t^{\pi'}(x_t) \quad (9)$$

$$\text{(by (3))} \geq V_t^{\pi_{TO}}(x_t) \quad (10)$$

To infer (9) from (8) we can iteratively apply the same reasoning we used to go from (6) to (8). This proves that the updated policy π'_{TO} cannot be worse than the previous one π_{TO} . Since $V^{\pi_{TO}}$ is nonincreasing and always bounded from below by the optimal value V^* , it follows from the *monotone*




Fig. 2. Cost function without the control effort term (15), considering a target point located at $[-7, 0]$ with weights $w_d = 100$, $w_p = 5 \cdot 10^5$ and $w_{ob} = 1 \cdot 10^6$. The green rectangle delimits the *Hard Region*.

convergence theorem that $V^{\pi^{TO}}$ converges to a constant value V^∞ . At convergence we must have:

$$V_t^\infty(x_t) = \min_{u \in \mathcal{U}} [l_t(x_t, u) + V_{t+1}^\infty(x_{t+1})] \quad \forall x_t, t$$

This is Bellman’s optimality equation, which is a sufficient condition for global optimality, so it follows that the algorithm converges to the optimal Value ($V^\infty = V^*$) and to an optimal policy ($\pi^\infty = \pi^*$). \square

III. RESULTS

This section presents the results of four different systems of increasing complexity: single integrator, double integrator, Dubins car, and 3-joint planar manipulator. For each system, the task consists of finding the shortest path to a target point (related to the end-effector’s position for the manipulator) while ensuring low control effort and avoiding an obstacle. The aim is verifying the capability of CACTO to learn a control policy to warm-start TO so that it can find “good” trajectories, where other warm-starting techniques, such as using the initial conditions (ICS) or random values, would make it converge to poor local minima. More precisely, the ICS warm-start uses the initial state (varying at each episode) as initial guess for the state variables in the OCP, for all time steps, and 0 as initial guess for the control variables.

For each system, we divided the XY-plane in a grid of 961 points and, starting from each point with 0 initial velocity, we compared the results of TO when warm-started with either the policy learned by CACTO, or a random initial guess, or the ICS for x and y (and 0 for the remaining variables). Table I reports the time and number of updates needed to train the critic and the actor for each test.

In the four tests, we changed the dynamics of the system under analysis keeping the same highly non-convex cost

TABLE I
Time, number of DNN updates, number of environment steps, and learning rates (LR_C for critic and LR_A for actor) used for each test.

	Time [h]	# updates	# env. steps	LR_C, LR_A
Single Int.	5.46	110k	3.4M	$5e^{-3}, 1e^{-4}$
Double Int.	7.29	130k	4.1M	$5e^{-3}, 5e^{-4}$
Dubins Car	10.45	260k	4.1M	$1e^{-3}, 5e^{-4}$
Manipulator	30.68	385k	6M	$1e^{-3}, 5e^{-5}$

function to ensure the presence of many local minima, that takes the following form:

$$l(\cdot) = \frac{1}{c_2} \left(\sum_{i=1}^4 l_i(\cdot) - c_1 \right) \quad (11)$$

$$l_1(\cdot) = w_d((x - x_g)^2 + (y - y_g)^2) \quad (12)$$

$$l_2(\cdot) = \frac{w_p}{\alpha_1} \ln(e^{-\alpha_1(\sqrt{(x-x_g)^2+c_2} + \sqrt{(y-y_g)^2+c_3+c_4})} + 1) \quad (13)$$

$$l_3(\cdot) = \frac{w_{ob}}{\alpha_2} \sum_{i=1}^3 -\ln(e^{-\alpha_2 \left(\frac{(x-x_{ob,i})^2}{(a_i/2)^2} + \frac{(y-y_{ob,i})^2}{(b_i/2)^2} - 1 \right)} + 1) \quad (14)$$

$$l_4(\cdot) = w_u \|u\|_2^2 \quad (15)$$

where (x_g, y_g) are the target point coordinates, $c_1 = 10000$ and $c_2 = 100$ are two constants, while $c_3 = 0.1$, $c_4 = -2c_3 - 2\sqrt{c_3}$, $\alpha_1 = 50$ and $\alpha_2 = 50$ are the parameters that define the smoothness of the *softmax* functions used to model respectively a cost valley in the neighborhood of the target and the three ellipses (centered at $(x_{ob,i}, y_{ob,i})$ and with principal axes a_i and b_i) forming the obstacle. The four terms composing the cost describe the task to be performed: (12) and (13) push the agent to reach the target point (with weights w_d and w_p , respectively), (14) makes it avoid the C-shaped obstacle represented by three overlapping ellipses (with weight w_{ob}) and (15) discourages it from using too much control effort (with weight w_u). Fig. 2 illustrates the cost function without the control effort term (15), where the cost peaks correspond to the obstacle penalties (14) and the cost valley to the term (13).

The critic was updated with Monte-Carlo (MC) for the 2D point, and with 50-step TD for the car and the manipulator. The reason for using TD rather than MC lies in the stability of the training of critic and actor. Indeed, in the early phase of training, using Monte-Carlo could lead to large variations of the critic, and indirectly also of the actor, because the target Values would be extremely different from the current Values. Moreover, in the early training phase TO could compute extremely poor trajectories because of the poor initial guess provided by the actor. In turn, these poor trajectories result in hard-to-learn *Value functions*. Therefore, we have empirically observed that these two effects can destabilize the training, leading to either longer training times, or even divergence of the algorithm.

To summarize the results, Table II reports the number of times that CACTO made TO find lower-cost solutions than the other two warm-starts, considering both the whole grid

TABLE II

Percentage of the time that warm-starting TO with CACTO leads to lower costs than using random initial guesses (CACTO vs. Random) and the initial conditions for x and y and 0 for the remaining variables (CACTO vs. ICS) as initial guess. Also the percentages when CACTO has lower or equal cost (*i.e.*, including ties) as its competitor are reported. The best result out of 5 runs is reported for random warm-start. For each system, we sampled from a 31×31 grid for the initial x and y coordinates (those of the end effector for the manipulator) and setting the remaining initial state components to 0 (the joint positions of the manipulator were obtained by fixing the orientation of the end-effector and inverting the kinematics). The *Hard Region* is the region delimited by $x \in [1, 15]$ m ($[1, 23]$ m in the manipulator test) and $y \in [-5, 5]$ m.

System	Whole Space		Hard Region	
	CACTO < (\leq) Random	CACTO < (\leq) ICS	CACTO < (\leq) Random	CACTO < (\leq) ICS
Single Int.	99.88% (99.88%)	14.49% (99.88%)	99.11% (99.11%)	91.96% (99.11%)
Double Int.	99.88% (99.88%)	12.38% (99.88%)	99.11% (99.11%)	91.96% (99.11%)
Dubins Car	89.72% (98.83%)	15.65% (95.56%)	100% (100%)	92.86% (100%)
Manipulator	91.78% (91.91%)	77.94% (78.33%)	87.50% (87.50%)	100% (100%)

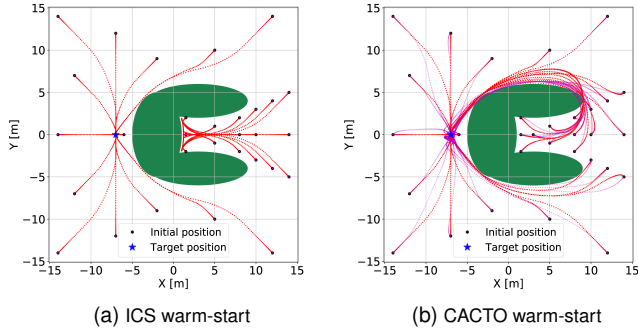


Fig. 3. Optimal trajectories (red) of the 2D double integrator obtained with ICS and CACTO warm-starts. In (b), the magenta lines represent the CACTO policy rollouts.




Fig. 5. Dubins car model: cost difference between CACTO warm-start and other two warm-starts normalized by the largest cost difference.




Fig. 4. Double integrator: cost difference between CACTO warm-start and other two warm-starts normalized by the largest cost difference.

and the region from which it is harder for TO to find “good” solutions (*Hard Region*).

A. Single Integrator

The first problem we considered is a simple 2D single integrator that has to reach a target point avoiding a C-shaped obstacle. The state is $[x, y, t] \in \mathbb{R}^3$, while the controls are the 2D velocities $[v_x, v_y] \in \mathbb{R}^2$, bounded in $[-4, 4] \frac{m}{s}$.

The task is simple except when the system starts from the *Hard Region* ($x \in [1, 15]$ m and $y \in [-5, 5]$ m), where TO can easily get stuck in local minima. Most of the times this means that the resulting trajectories point immediately towards the target, making the 2D point stay at the right boundary of the vertical ellipse. But since the C-shaped obstacle is modelled as three overlapping ellipses represented by three soft penalties (14) in the cost function, it can also occur that the 2D point passes through them to reach the target, albeit

at high cost. Table. II reports the percentage of the time that warm-starting TO with CACTO rollouts leads to lower costs than using random values and the initial conditions for x and y and 0 for the remaining variables as initial guess. CACTO warm-start wins over the other two techniques, particularly if we consider the agent starting from the *Hard Region* where warm-starting TO with CACTO leads to lower-cost solutions 99.11% and 91.96% of the time compared to using random values and the ICS as the initial guess, respectively.

B. Double Integrator

We run the same experiment with the 2D point considering a double integrator dynamics. Therefore, now the state includes also the velocities $[x, y, v_x, v_y, t] \in \mathbb{R}^5$ and the control inputs are the accelerations $[a_x, a_y] \in \mathbb{R}^2$. As illustrated in Fig. 3(b), the rollouts of the CACTO policy are already close to the globally optimal trajectories, therefore TO only needs to refine them when they are used as initial guess. We refer to trajectories as globally optimal if their cost is the lowest among those obtained solving several TO problems warm-started with random initial guesses. Fig. 4 shows instead the cost difference (normalized by the highest difference in absolute value) when TO is warm-started with rollouts of the CACTO policy in place of random values or the ICS, respectively. Fig. 4(a) clearly shows that using rollouts of the policy learned by CACTO as an initial guess makes TO find lower-cost solutions from almost any initial state compared to those found with a random initial guess. In Fig. 4(b) instead, we can notice that ICS is a good initial guess for the majority of initial states and TO finds the same solutions as when warm-started with CACTO rollouts. However, when starting the agent in the *Hard Region*,




Fig. 6. TO solutions considering 12 different initial configurations of the manipulator and ICS warm-start. The red dotted lines represent the trajectories performed by the end-effector (EE).




Fig. 7. 3-DoF Manipulator: normalized cost difference when warm-starting TO with CACTO compared to using the initial conditions as initial guess for the joint positions and 0 for the remaining variables.

the ICS warm-start makes TO find poor local minima, where the agent remains stuck in that region or passes through the obstacle, whereas CACTO enables TO to successfully bring the agent to the target without touching the obstacle. Also in this case, in the *Hard Region*, warm-starting TO with CACTO rollouts rather than with random values or ICS makes TO find lower costs with the same percentage as in the previous test.

C. Dubins Car

To test CACTO with a higher-dimensional system, we selected a jerk-controlled version of the so-called *Dubins car model* [30], while keeping the same environment and cost function of the previous tests. Now the state has size 6 because it includes the steering angle θ , the tangential velocity v , and acceleration a , in addition to the coordinates x and y of the car’s center of mass and time t : $s = (x, y, \theta, v, a, t) \in \mathbb{R}^6$. The control is still bi-dimensional and it consists in the steering velocity and the jerk $[\omega, j] \in \mathbb{R}^2$. When the car starts from the *Hard Region*, TO warm-started with CACTO always finds a lower cost than that obtained by warm-starting TO with random values, and it does so the 92.86% of the time when compared to using ICS warm-start, as reported in Table II. In addition, Fig. 5(b) shows that TO warm-started with ICS is not able to find the globally optimal solution also when the car starts from a point along the vertical line passing through




Fig. 8. Cost of the trajectory found by TO (starting the 2D point from $[5, 0]$ m with 0 velocity) when warm-started with rollouts of the policy learned with CACTO (green), our DDPG (DDPG-c in orange), Stable Baselines’ DDPG (DDPG-sb in blue) and PPO (red) as their trainings proceed. The shaded area denotes the standard deviation over 5 different runs.

the target point. This is due to the fact that in that region the gradient of the cost is zero along the initialization itself.

D. 3-DoF Planar Manipulator

Finally, we tested our algorithm on a problem with a 7D state and 3D control space. It consists of a 3-DoF planar manipulator with base fixed at $[-7, 0]$ m, working in the same environment of the previous tests, whose end-effector has to reach a target point located at $x_g = -20$ m and $y_g = 0$ m. The cost function is always (11), where x and y represent the coordinates of the end-effector. Fig. 6 shows some solutions found by TO when warm-started with the ICS. It may seem that TO succeeds in some cases in finding the globally optimal trajectories for those initial configurations, but actually all of them are only locally optimal, as shown by the negative cost difference in Fig. 7 when those solutions are compared to the ones obtained by warm-starting TO with CACTO. This test is harder, meaning that it is much easier for TO to find poor local minima, not only due to the larger state-action space, but also because the actor has to intrinsically learn the manipulator kinematics. Indeed, considering the whole manipulator workspace, CACTO warm-start wins over using the ICS or random values 77.94% and 91.78% of the time, respectively, as reported in Table II.

E. Comparison with DDPG and PPO performance

To compare CACTO and DDPG, we have used our implementation of DDPG with hand-tuned hyper-parameters (DDPG-c), as well as that from Stable Baselines (DDPG-sb), to find a warm-starting policy for the 2D double integrator test of Section III-B. In addition, we have made the comparison also with PPO from Stable Baselines. Fig. 8 shows the costs obtained by TO as functions of the computation time allocated to each algorithm to train its warm-starting policy. Clearly, CACTO was faster than the other RL algorithms in learning a policy enabling TO to find lower-cost solutions, and its training was also more stable (rollouts with lower variance). We have reported only the results with the double integrator because, despite our efforts, we did not manage to make DDPG converge in a reasonable number of episodes in the car and manipulator tests. Indeed, as stressed in [31], DDPG

is known to be very sensitive to its hyper-parameters making it prone to converge to poor solutions or even diverge.

Besides this quantitative comparison, we could also try to qualitatively compare our results with the ones reported for DDPG in [11]. Among their tests, the most similar to ours is the *fixedReacher*, where a 3-DoF arm must reach a fixed target; this has the same dimensions as our manipulator test. However, our cost/reward function is highly non-convex, particularly due to the obstacle avoidance term, while theirs was quadratic and consisted of only two terms. This makes it harder to reach convergence in our DNN training. Using CACTO in such a simple setting, with a convex cost-function, would not make sense: TO would converge to the global optimum even with a trivial warm-start (indeed DDPG converged to roughly the same policy found by iLQG). Consequently, it makes no sense to use these results for a comparison with CACTO, which should instead be based on highly non-convex problems, where TO cannot find the global optimum with a naive initial guess.

IV. CONCLUSIONS

This paper presented a new algorithm for finding quasi-optimal control policies. In particular, we addressed the open problems affecting Trajectory Optimization and deep Reinforcement Learning: the possibility of getting stuck in poor local minima when TO is not properly warm-started on one side, and the low sample efficiency of deep RL. The proposed algorithm relies on the combination of TO and RL in such a way that their interplay guides efficiently the RL state-control space exploration process towards the globally optimal control policy, to be used then as TO initial guess provider.

We provided a global convergence proof for a discrete-space version of our algorithm, which gives insight into its underlying theoretical principles.

We have shown the effectiveness of the algorithm testing it on four systems of increasing complexity, with highly non-convex cost functions, where TO struggles to find “good” solutions. Even though preliminary, our results validate our methodology and unlock a wide range of possible applications.

Despite our encouraging results, we believe that CACTO can still be improved, in particular concerning its computation time. We are investigating how to speed up the learning of the critic using Sobolev Training [32], which could exploit the derivatives of the *Value function* computed by a DDP-like TO algorithm [33]. We are also considering to implement CACTO using sampling-based multi-query planners, as done with PRMs in [34]. Moreover, besides using CACTO as initial guess provider for TO, where the RL and TO environments must match, we are interested in using it as a deep RL technique to find directly a control policy, where the two environments do not need to match (*e.g.*, the TO problem could be a simplified version of the RL environment without noise sources). Finally, we plan to extend CACTO to optimize also hardware parameters, to create a concurrent design (co-design) framework robust against the local minima problem, which is a crucial issue in co-design applications [35].

REFERENCES

- [1] B. Katz, J. Di Carlo and S. Kim, “Mini Cheetah: A Platform for Pushing the Limits of Dynamic Quadruped Control,” in *Proc. of '19 IEEE ICRA*, Montreal, Canada, pp. 6295–6301, 2019.
- [2] M. Neunert et al., “Whole-Body Nonlinear Model Predictive Control Through Contacts for Quadruped” in *IEEE RA-L*, 3(3), pp. 1458–1465, 2018.
- [3] T. Erez, K. Lowrey, Y. Tassa, V. Kumar, S. Koley and E. Todorov, “An integrated system for real-time model predictive control of humanoid robots,” in *Proc. of '13 IEEE-RAS Humanoids*, Atlanta, Georgia, USA, pp. 292–299, 2013.
- [4] N. Scianca, D. De Simone, L. Lanari and G. Oriolo, “MPC for Humanoid Gait Generation: Stability and Feasibility” in *IEEE T-RO*, 36(4), pp. 1171–1188, 2020.
- [5] N. Mansard, A. Del Prete, M. Geisert, S. Tonneau and O. Stasse, “Using a Memory of Motion to Efficiently Warm-Start a Nonlinear Predictive Controller,” *Proc. of '18 IEEE ICRA*, Brisbane, QLD, Australia, pp. 2986–2993, 2018.
- [6] M. Bardi and I. Capuzzo Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, 1st ed., 12, Birkhäuser, 1997.
- [7] J.N. Tsitsiklis, “Globally Optimal Trajectories” in *IEEE Trans. Automat. Contr.*, 40(9), pp. 1528–1538, 1995.
- [8] L.C. Polymenakos, D.P. Bertsekas and J.N. Tsitsiklis, “Implementation of efficient algorithms for globally optimal trajectories” in *IEEE Trans. Automat. Contr.*, 43(2), pp. 278–283, 1998.
- [9] A. Gorodetsky, S. Karaman and Y. Marzouk, “Efficient high-dimensional stochastic optimal motion control using tensor-train decomposition” in *Proc. of '15 RSS*, Rome, Italy, 2015.
- [10] E. Stefansson and Y.P. Leong, “Sequential alternating least squares for solving high dimensional linear Hamilton-Jacobi-Bellman equation” in *Proc. of '16 IEEE/RSJ IROS*, Daejeon, Korea, pp. 3757–3764, 2016.
- [11] T.P. Lillicrap, J.J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver and D. Wierstra, “Continuous control with deep reinforcement learning,” 2015. Online available: arXiv:1509.02971
- [12] S. Fujimoto, H. van Hoof, D. Meger, “Addressing Function Approximation Error in Actor-Critic Methods,” 2018. Online available: arXiv:1802.09477
- [13] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, S. Levine, “Soft Actor-Critic Algorithms and Applications,” 2019. Online available: arXiv:1812.05905v2
- [14] Y. Hou, H. Hong, Z. Sun, D. Xu and Z. Zeng, “The Control Method of Twin Delayed Deep Deterministic Policy Gradient with Rebirth Mechanism to Multi-DOF Manipulator” in *Electronics*, 10(7:870), 2021.
- [15] Z. Zhang, J. Chen, Z. Chen and W. Li, “Asynchronous Episodic Deep Deterministic Policy Gradient: Toward Continuous Control in Computationally Complex Environments” in *IEEE Trans. Cybern.*, 51(2), pp. 604–613, 2021.
- [16] S. Levine and V. Koltun, “Guided Policy Search,” in *Proc. of '13 ICML*, Atlanta, Georgia, USA, pp. III-1–III-9, 2013.
- [17] I. Mordatch and E. Todorov, “Combining the benefits of function approximation and trajectory optimization,” in *Proc. of '14 RSS*, 4, 2014.
- [18] M. Vecerik et al., “Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards,” 2017. Online available: arXiv:1707.08817
- [19] A. Filos et al., “PsiPhi-Learning: Reinforcement Learning with Demonstrations using Successor Features and Inverse Temporal Difference Learning,” in *Proc. of '21 ICML*, pp. 3305–3317, 2021.
- [20] T. Osa, A. M. G. Esfahani, R. Stolkin, R. Lioutikov, J. Peters and G. Neumann, “Guiding Trajectory Optimization by Demonstrated Distributions,” in *IEEE RA-L*, 2(2), pp. 819–826, 2017.
- [21] S. Gu, T. Lillicrap, I. Sutskever and S. Levine “Continuous Deep Q-Learning with Model-based Acceleration,” in *Proc. of '16 ICML*, New York, New York, USA, vol 48, 2016.
- [22] S. Fujimoto, D. Meger and D. Precu, “Off-Policy Deep Reinforcement Learning without Exploration,” in *Proc. of '19 ICML*, Long Beach, California, USA, pp. 2052–2062, 2019.
- [23] D.P. Bertsekas, *Reinforcement Learning and Optimal Control*, Athena Scientific, Belmont, Massachusetts, USA, pp. 256–266, 2019.
- [24] R. S. Sutton and A. G. Barto, *Reinforcement Learning - An Introduction*, The MIT Press, Cambridge, Massachusetts, USA, 2018.
- [25] M.L. Bynum, G.A. Hackebeil, W.E. Hart, C.D. Laird, B.L. Nicholson, J.D. Siirola, J.P. Watson and D.L. Woodruff, *Pyomo — Optimization Modeling in Python*, 3rd ed., 67, Springer, 2021.

- [26] A. Wächter and L. T. Biegler, "On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming," in *Math. Program.*, 106, pp. 25-57, 2006.
- [27] L. Lu, Y. Shin, Y. Su and G. E. Karniadakis, "Dying ReLU and Initialization: Theory and Numerical Examples," in *Commun. Comput. Phys.*, 28(5), pp. 1671-1706, 2020.
- [28] D.P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," 2017. Online available: arXiv:1412.6980v9
- [29] J. Nocedal, A. Wächter and R.A. Waltz, "Adaptive Barrier Update Strategies for Nonlinear Interior Methods," in *SIOPT*, 19, pp. 1674-1693, 2009.
- [30] L. E. Dubins, "On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents," in *Am. J. Math.*, 79(3), pp. 497–516, 1957.
- [31] R. Islam, P. Henderson, M. Gomrokchi and D. Precup, "Reproducibility of Benchmarked Deep Reinforcement Learning Tasks for Continuous Control," 2017. Online available: arXiv:1911.11679
- [32] W. M. Czarnecki, S. Osindero, M. Jaderberg, G. Swirszcz and R. Pascanu, "Sobolev Training for Neural Networks," in *Proc. of '17 NeurIPS*, Atlanta, Georgia, USA, 30, 2017.
- [33] C. Mastalli et al., "Crocodyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control," 2020. Online available: arXiv:1509.02971
- [34] E. Jelavic, F. Farshidian and M. Hutter, "Combined Sampling and Optimization Based Planning for Legged-Wheeled Robots," in *Proc. of '21 IEEE ICRA*, pp. 8366–8372, 2021.
- [35] G. Grandesso, G. Bravo-Palacios, P. M. Wensing, M. Fontana and A. Del Prete, "Exploring the Limits of a Redundant Actuation System Through Co-Design," in *IEEE Access*, 9, pp. 56802–56811, 2021.