



HAL
open science

Value learning from trajectory optimization and Sobolev descent: A step toward reinforcement learning with superlinear convergence properties

Amit Parag, Sébastien Kleff, Léo Saci, Nicolas Mansard, Olivier Stasse

► To cite this version:

Amit Parag, Sébastien Kleff, Léo Saci, Nicolas Mansard, Olivier Stasse. Value learning from trajectory optimization and Sobolev descent: A step toward reinforcement learning with superlinear convergence properties. International Conference on Robotics and Automation (ICRA 2022), May 2022, Philadelphia, United States. hal-03356261v1

HAL Id: hal-03356261

<https://hal.science/hal-03356261v1>

Submitted on 27 Sep 2021 (v1), last revised 3 Mar 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Value learning from trajectory optimization and Sobolev descent: A step toward reinforcement learning with superlinear convergence properties

Amit Parag^{1,2}, Sébastien Kleff^{2,3}, Léo Saci², Nicolas Mansard^{1,2} and Olivier Stasse^{1,2}

Abstract—The recent successes in deep reinforcement learning largely rely on the capabilities of generating masses of data, which in turn implies the use of a simulator. In particular, current progress in multi body dynamic simulators are underpinning the implementation of reinforcement learning for end-to-end control of robotic systems. Yet simulators are mostly considered as black boxes while we have the knowledge to make them produce a richer information. In this paper, we are proposing to use the derivatives of the simulator to help with the convergence of the learning. For that, we combine model-based trajectory optimization to produce informative trials using 1st- and 2nd-order simulation derivatives. These locally-optimal runs give fair estimates of the value function and its derivatives, that we use to accelerate the convergence of the critics using Sobolev learning. We empirically demonstrate that the algorithm leads to a faster and more accurate estimation of the value function. The resulting value estimate is used in model-predictive controller as a proxy for shortening the preview horizon. We believe that it is also a first step toward superlinear reinforcement learning algorithm using simulation derivatives, that we need for end-to-end legged locomotion.

I. INTRODUCTION

We consider the problem of computing the optimal control policy in the continuous domain, with application to the end-to-end control of polyarticulated robots, like torque-controlled manipulator arms or legged robots. Since the advent of deep reinforcement learning [1] and its generalization to continuous domain [2], [3], robot control has been recognized has one of the major challenges of the domain [4]. On the one hand, researches have targeted advanced robot scenarios and sensor outputs [5], [6] often relying on an underpinning low-level controller (impedance or Cartesian controller). On the other hands, efforts have been made to target more complex robots such as quadruped [7] or biped [8] robots by focusing on a more accurate modeling of the robot dynamics [9] and a more advanced (stronger, more robust, faster) simulation of the polyarticulated behavior [10], [11]. In parallel, RL solvers have made great progress and are now more stable [3] and faster to converge [12].

Yet solving end-to-end polyarticulated problem remains a difficult challenge for generic RL solvers. Indeed, sending reference torques directly to the simulator or to the model makes the system much more unstable [13]. The introduction of a lower level of control, hand-tuned to the particular robotic system, mitigates this effect [7]. Yet imposing a priori the low-level controllers structure and its corresponding impedance might be suboptimal or even lead to improper

behavior during demanding contacts, in particular for biped locomotion [14]. Following the recent advances in legged locomotion [15]–[17], we believe that computing a policy that would drive the robot directly in torque, if achievable, would have great practical interest.

While current RL solvers have difficulties to directly tackle end-to-end polyarticulated problems, several approaches have been proposed to make their work easier. Acknowledging for the improved capabilities of motion planner to discover complex movement, guided policy search bootstrap the reinforcement loop with locally optimal trajectories obtained from a trajectory optimizer [18]. The trajectories can later be refined using the trained policy to produce better demonstrations by warm-starting the trajectory optimizer [19], or the policy and the demonstrations be simultaneously optimized [20], [21]. These works show the correlation between policy learning and trajectory optimization, as optimal trajectories and optimal policies are two ways of considering the solution of the optimal control (or Markov decision) problem. Reciprocally, several works have tried to take advantage of this duality to learn a quantity that would help a trajectory optimizer. In [22], a trajectory optimizer is helped with a dedicated dynamic model learned on a specific task. In [23], a dedicated cost model is learned to cope with the prohibitive complexity of optimizing long-horizon trajectories. Shortening the trajectory horizon typically leads to observing the importance of obtaining a good approximation of the value function, both in trajectory optimization and policy learning. It was originally proposed to learn the value as a proxy for the cost corresponding to the end of the horizon [24]. While a fair approximation of the value function can be obtained offline by sampling long locally optimal trajectories, the cost might become prohibitive when the state space is large or implies clever sampling solutions [25]. In [26], the value approximation gets reinforced when short horizon trajectories are optimized with a terminal cost set to the approximated value. Producing new optimal trajectories while exploiting a previously optimized value approximation in a reinforcement loop can be considered as an actor-critic formulation, where the actor is optimally acting over an horizon: it can then be shown that the prediction reduces the noise in learning the value and speed up the training [27]. As shown in [28], deep policy gradients methods often fail to adhere to the key primitives of reinforcement learning: the role of the value function estimated by the critic is only marginal since the learned value function does not accurately model the underlying value function.

In this paper, we focus on learning the value function by

¹Artificial and Natural Intelligence Toulouse Institute, France

²LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

³New York University, USA

combining a trajectory optimizer and a supervised learning phase using value approximation and its derivatives. Our objective is to learn the value with a high accuracy, while avoiding the need of extensively sampling the state space. A trajectory optimizer is first used to sample optimal trajectories over a short horizon a get a lower bound of the value function. As we are using differential dynamic programming [29]–[31], a particular instance of (direct shooting) trajectory optimizer, we also get for free samples of the derivative of the value (see Sec. II). We then train a neural network to approximate the value and its derivative using a Sobolev loss [32] (see Sec. III). This particular loss accelerates the convergence to a more accurate approximation, while enabling better generalization despite sparse sampling. As a side effect, it also provides an accurate way of evaluating the approximate gradient of the value. As our trajectory optimizer also demands the Hessian of the value, we propose a particular network architecture, that we call Gauss-residual network, which provides a fair approximate of the second derivatives. We evaluate the interest of our approach on three classical control problems : inverted pendulum, cartpole and unicycle.

II. RECURSIVE OPTIMALITY

A. Problem formulation and notations

While optimal control problem are best written with continuous time and integrals, we directly consider here its transcription into a time-discrete system with autonomous dynamics:

$$x_{t+1} = f(x_t, u_t, \Omega) \quad (1)$$

where the states x are living in an n -dimensional vector space $x \in \mathcal{X}$, (possibly representing an element of a Lie group), and controls $u \in \mathcal{U}$, with Ω a parametrization of the environment where the system evolves and f , the time independent dynamic function that governs the state transition from the current state and control. The goal in optimal control is to find pairs of trajectories $X : t \rightarrow x(t) \in \mathcal{X}$ and $U : t \rightarrow u(t) \in \mathcal{U}$ that minimizes some cost functional $L(X, U)$. Typically, the cost functional is composed of an integral cost term and a terminal cost term. For a discrete system, the cost functional can be written as :

$$L(X, U) = \sum_{k=0}^{+\infty} l(x_k, u_k) + l_t(x_N) \quad (2)$$

where $l(x, u) : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is the running cost and $l_t(x, u) : \mathbb{R}^n \rightarrow \mathbb{R}$ is the terminal cost. Therefore given an initial state x_0 and a time horizon T , aim of optimal control is to find the pair of optimal state and control trajectories X^*, U^* that minimizes the cost functional.

We denote by $V : x \rightarrow V(x) \in \mathbb{R}$ the value function, i.e. the optimal value of $L(X, U)$ when the trajectory X starts from a given state x , and by $\pi : x \rightarrow \pi(x) \in \mathcal{U}$ the optimal policy.

B. Truncating the horizon

While we typically want to control the system during an indefinite (infinite) horizon, finite computational resources implies to shorten the horizon to a finite length T . We denote by V_k cost-to-go from a starting state x :

$$V_k(x) = \min_U \sum_{j=k}^{T-1} l(x_j, u_j) + l_t(x_t) \quad (3)$$

where k is the time step.

Recursive optimality can be obtained if the terminal cost is chosen to be the value function. In that case, the problem writes:

$$\begin{aligned} \min_{X, U} \sum_{k=0}^{T-1} l(x_k, u_k) + V(x_T) \\ \text{s.t. } \forall k = 0..T-1, x_{k+1} = f(x_k, u_k) \end{aligned} \quad (4)$$

This problem is strictly equivalent to the infinite horizon problem but can be solved with finite resources, given that $V(x)$ can be properly evaluated.

In that particular case, the cost-to-go is independent of the timestep and is equal to the value function:

$$V_k(x) = V(x), \quad \forall k = 0 \dots T$$

C. Differential Dynamic Programming

Differential dynamic programming (DDP) exploits the recursive nature of Bellman Optimality Principle by rewriting the optimal value function, Eq 3, as

$$V_k(x) = \min_u l(x, u) + V_{k+1}(f(x, u)) \quad (5)$$

with the terminal condition $V_T(x) = l_x$. It proceeds by iteratively solving the the optimal control problem described in Sec II-A to find the local solution by performing a backward pass to approximate a quadratic fit of the value function around the current candidate trajectory and a forward pass to compute a new nominal trajectory based on the value function computed in backward pass. For our experiments, we use Crocoddyl [31] as our model-based DDP solver.

A consequence of the backward pass is that it reconstructs a quadratic model of the value function along the horizon. The solver then returns an evaluation of the cost-to-go and its first and second order derivatives over the preview horizon.

DDP is a second-order algorithm which exhibits super-linear convergence rate [29], [33]. It then requires at each iteration an evaluation of the dynamic and cost function and their derivatives. In particular, we need to provide the first and second order derivatives of the cost $l(x, u)$.

D. Value at the horizon start and end

In the following, we will use an approximation of the value function as a proxy to represent the truncated horizon end. This implies that we should be able to evaluate an approximation of the value function and its first and second order derivatives at the end of the horizon. The solver will then return a refined approximation of the value function at the beginning of the horizon, and its derivatives. We will now use these two properties to build our algorithm.

III. DIFFERENTIAL VALUE PROGRAMMING

A. Algorithm principles

1) *Batch of trajectories and cost-to-go learning*: The first iteration of our algorithm, named Differential Value Programming (DVP), simply generates a batch of optimal trajectories of horizon length T without any terminal cost model, i.e. approximating the value function at the end of the horizon by 0. We then learn the value function by supervised learning, as explained with more details below. The result of this first iteration is a neural network approximating the cost-to-go for an horizon of T .

2) *Iterative value learning*: DVP then proceeds by iteratively building upon its estimates of value functions. In the subsequent iterations, we replace the terminal cost functional with the approximated value predicted by the neural network. So, after the initial iteration is complete, we replace (3) with

$$V_i(x) = \min_u \sum_{k=0}^{T-1} l(x, u) + V_\theta^{i-1}(x_T) \quad (6)$$

where i is the DVP iteration number and V_θ is a deep neural network parameterized by θ representing the value approximated at the previous iteration. Should each iteration result in a perfect training, iteration i would lead to the approximation of the cost-to-go for an horizon of $(i+1) \times T$, which would tend toward V as i increases.

B. Sobolev Regression

Classical regression invariably involves optimizing the parameters of a deep neural network such that the error between the learned function and the *ground truth* function decreases. While the function approximation theoretically guarantees that neural networks can learn arbitrarily well, the accuracy and generalization capabilities yet depend on the quality (exactness and density) of the training dataset. The premise of Sobolev learning, [34], is to use the derivatives of the function to be approximated, such that the derivatives of the network match the derivatives of the learned function. In our work, we further benefit from the advantage of Sobolev learning since we need to provide the derivatives of the value function to our optimal control solver.

The loss function for our regression is then composed of two terms:

$$loss_f = \sum_{i=1}^N \lambda(m(x_i | \theta), f(x_i)) \quad (7)$$

$$loss_d = \sum_{i=1}^N \sum_{j=1}^K \lambda_j(D_x^j m(x_i | \theta), D_x^j f(x_i)) \quad (8)$$

where $loss_f$ penalizes the differences between the model and the dataset with a norm λ (in our case the L_2 norm), $loss_d$ penalizes the difference between the derivatives of the model and the derivatives of the observations, and D_x^j are the higher order derivatives evaluated at x (we stay at order 1 in our implementation, although order 2 would be feasible) for a norm λ_j (also L_2 in our implementation).

Sobolev learning has been shown to lead to better generalization in robotics and reinforcement learning [35], albeit at a higher computation cost [36], as it constrains training by forcing neural networks to fit a target slope. However, encoding the target derivative information in neural networks has been shown, empirically, to increase robustness against noise as proven in [37] and mitigate the problem of increased computation cost by being more data efficient [38].

C. Evaluating second-order derivatives

As mentioned above, Sobolev learning allows us to obtain a more accurate convergence despite a sparser dataset, and also regularizes the derivatives of the model. This second point is important as these derivatives are needed by the DDP, as explained in Sec. II. Yet DDP also requires the Hessian of the value function. While Sobolev would theoretically benefit from second-order information, which is also available in practice from the DDP, the evaluation of the second-order loss (8) and the backpropagation of its gradient are not reasonable to evaluate.

We rather propose to set up a particular network architecture, called Gauss-residual network architecture, that we demonstrate to properly regularize the training of the Hessian. Our network model reads:

$$m(x|\theta) = r(x|\theta)^2 \quad (9)$$

where $r(x|\theta)$ is a vector-tailed network, called the residual network, whose output is squared and summed to produce the final value. The gradient of this model is then simply:

$$D_x^1 m(x|\theta) = 2D_x^1 r(x|\theta)^T r(x|\theta)$$

where $D_x^1 m$ and $D_x^1 r$ are the gradient of m and Jacobian of r respectively. The the Hessian of m is approximated as:

$$D_x^2 m(x|\theta) \approx 2D_x^1 r(x|\theta)^T D_x^1 r(x|\theta)$$

This approximation is known as the Gauss approximation, and leads to the famous superlinear algorithm Gauss-Newton (which is typically implemented in most of DDP frameworks). We can interpret this network as computing the value as the square of a latent vector r , and approximating the second order derivatives of the value as the derivatives of the network tail.

We will show empirically that this particular model, combined with first-order Sobolev learning, leads to an accurate, robust and efficient representation of the value function, which perfectly suits to the requirements of the DDP algorithm.

IV. EMPIRICAL EVALUATION

A. Experimental setup

We propose to benchmark our approach on 3 systems by using as a baseline (ground truth) the arbitrarily good approximations of the value function that can be obtained by sampling large datasets of long optimal trajectories in various configurations. We used the following problems:

a) *Unicycle*: The unicycle [39] features a kinematic model of evolving on the 2D horizontal plane either driving forward or turning on the spot. Denoting the configuration vector $q = (x, y, \theta)^T$ of dimension $n = 3$, the unicycle model reads:

$$\dot{x} = v \cos(\theta) \quad (10)$$

$$\dot{y} = v \sin(\theta) \quad (11)$$

$$\dot{\theta} = \omega \quad (12)$$

where the control $u = (v, \omega)$ includes the unconstrained longitudinal and angular velocities. The task is to reach the goal position $q = (0, 0, 0)^T$ while minimizing the residual sum of errors:

$$L = ||w_1 * q||^2 + ||w_2 * u||^2 \quad (13)$$

b) *Cartpole*: A cartpole¹ is a dynamical system where an underactuated pole is attached on top of a 1D actuated cart. The task is to balance the pole around its unstable equilibrium (upper position) by controlling the horizontal forces acting on the cart [40]. The cost functional to be minimized is:

$$L = w_1 ||x||^2 + w_2 ||u||^2 \quad (14)$$

where $x = (q, v)$ is the state and the control u is force exerted on the cart.

c) *Simple Pendulum*: The inverted pendulum¹ swingup problem consists in bringing the pendulum from a random position to its upper equilibrium and maintaining it upright. The cost functional we use is identical to the cartpole cost functional, with u representing the torque applied about the pendulum's rotation axis.

In our experiments, we empirically establish baselines of DVP. Concretely, we aim to quantify the convergence properties of DVP and the influence of Sobolev training. In order to characterize the convergence properties of our algorithm, we use a validation dataset as a substitute for the value function of the infinite horizon problem. This can be done, albeit at an extremely high computational cost, by solving the optimal control problem over extremely long horizons.

B. Convergence of DVP

1) *Overall convergence*: Fig. 1 shows the value function learned by our algorithm in its 10th iteration and the corresponding mean squared error with respect to the ground truth value function. Convergence to 10^{-5} is obtained for the pendulum and the cartpole, and convergence to 10^{-3} is obtained for the unicycle (which would eventually reach the same accuracy with more iterations).

The difference between the value function at each iteration and *ground truth* with respect to DVP iterations is depicted in Fig. 2. For the pendulum case, just 1 iteration is sufficient for DVP to achieve convergence. For cartpole, DVP takes a few more iterations to converge to a good enough approximation

¹We use the Open Ai gym implementation of the dynamical model.

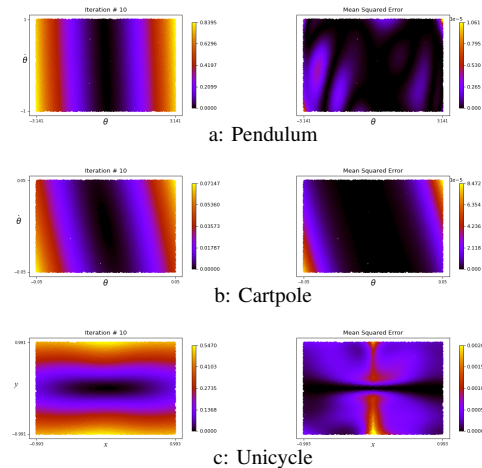


Fig. 1: The predicted value functions after 10 iterations (right) and mean squared error between the predicted function and ground truth value function (left).




Fig. 2: Evolution of the mean squared difference between value predictions of iterations 1, 5 and 10 and ground truth for the 10th iteration.

of the global value function. For systems with regions of local minima like unicycle, achieving convergence requires relatively more iterations.

Fig. 3 quantifies the convergence through validation loss at the end of every epoch. We used the idea of *Bellman residual* [41], [42] to establish the convergence criteria of DVP. It is easy to see that as the value function estimates come closer to the optimal value function, the difference in successive estimates decreases. Residuals between two successive value functions can be a good indication of prediction. Upon or near convergence, the higher iterations of DVP should not show much difference between their behaviors.

2) *Influence of the horizon length*: We discuss two results regarding the impact of the initial horizon length T . Short horizons lead to important differences between the cost-to-go and the value, hence to a poor approximation of the value in early iterations of DVP. This clearly appears in Fig. 4, where the truncation to small horizon leads to trajectories far from the optimum. Once the value is properly estimated,

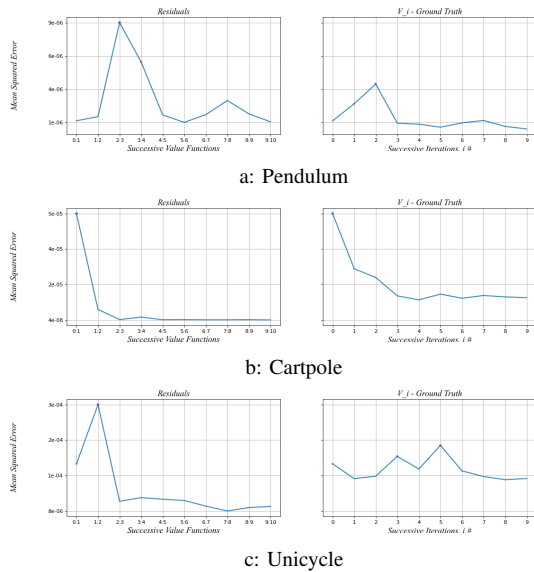


Fig. 3: Bellman residuals i.e difference between the value at two successive iterations (left) and difference between value function at each iteration and ground truth (right)




Fig. 4: (left) State trajectories computed by the solver with the value as terminal proxy (6), for $T = 20$ to 100. The trajectory in black is the reference ground truth state trajectory computed by the solver at horizon 1000. (right) State trajectories computed by the solver without the terminal value proxy, for different horizons: more variability (suboptimality) appears.

the bundle of trajectories converges closer to the optimum (on the unicycle, the convergence is not perfect despite an accurate convergence to the value, due to nonholonomy).

Consequently, DVP converges faster when T increases. For the considered system, the typical duration of an episode (until system steady state) is 150, and DVP shows proper convergences for $T \geq 40$, see Fig. 5. For smaller T , the convergence is slower or even fails to reach a global optimum. This is not surprising and can be mitigated by increasing the dataset size.

3) *Robust convergence*: We empirically establish the stability and robustness of our algorithm by forcing DVP to learn a noisy dataset at the first iteration. We find that DVP requires only a few iterations to learn the ground truth. Fig. 6 shows the convergence of DVP under various levels of initialization noise.

C. Influence of Sobolev learning

1) *Importance of the Sobolev loss*: Our experiments with Sobolev learning corroborate the generalization capabilities

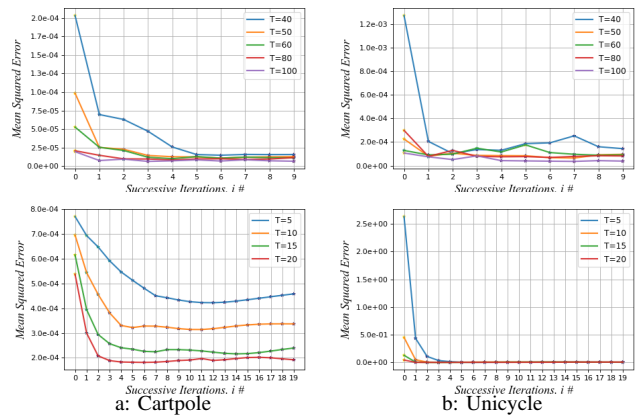


Fig. 5: Evolution of the difference between ground truth value function and different iterations, for high horizons (top) and small horizons (low).

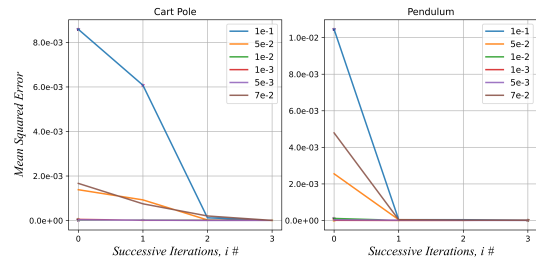


Fig. 6: Evolution of the algorithm under noise initialization. The y-axis measures the mean squared difference between each iteration and ground truth value function.

and confirm that Sobolev regression requires fewer training epochs than classical regression, see Fig. 7. Sobolev training requires only 50 samples to achieve a 10^{-3} accuracy on a validation dataset.

2) *Gauss-residual network*: For our experiments, we use a 3 layered Gauss-residual network with hyperbolic tangent as an activation function and 64 units in each hidden layers. The final residual layer contains 3 units. Empirically, we find that the advantage of modeling the value function as a squared residual leads to faster convergence during Sobolev training as shown in Fig. 8. The gradients of residual network are also more accurate than those of feed forward network as




Fig. 7: Comparison of sample size needed to achieve sufficient accuracy with classical and Sobolev regression (left) with Sobolev loss (right) without Sobolev loss. The Sobolev loss is less prone to be stuck in a local minimum when smaller dataset are provided when approximating the value (top). The derivative of the value is always better approximated with Sobolev loss (bottom).




Fig. 8: Sobolev training Loss comparison for Gauss Residual Network and Feed Forward Network. The red curve shows the training loss for Gauss Residual Network, while the green curve shows the training loss of feedforward network. Fig. a shows the training loss in predicting the value function while Fig. b shows the training loss in gradients for feedforward and residual network

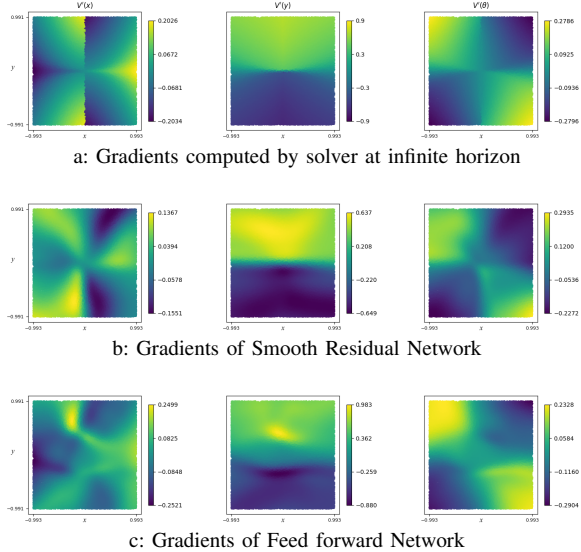


Fig. 9: Evaluation of the gradient approximation for the unicycle. (top) true gradients of the underlying function computed by the solver. (mid) gradients of Gauss residual network. (bottom) gradients when using a plain feedforward network (no Gauss approximation). The underlying function is $V(X, T)$ where $T = 1000$ and X is the 3-dimensional state vector, $X = (x, y, \theta)$. The two neural networks were trained for 50000 epochs on 100 samples generated by solver at horizon = 50.

shown on Fig. 9. Additionally, we benefit from being able to make Gauss approximation of the Hessian, since computing exact Hessians of neural network can become prohibitively expensive when the number of hidden layers increase.

D. Comparison with PPO

We compare the accuracy and rate of convergence of DVP against a classical actor-critic RL algorithm off the shelf, which approximates the value as a side quantity when computing the policy (critic network). To cope with the requirement of the RL solver, we have consider for both DVP and PPO a discounted version of the optimal control problem considered until now. Fig. 10 shows the qualitative comparison of the value functions predicted for pendulum, cartpole and unicycle by DDP, DVP and PPO. PPO properly captures the overall shape and the spread of the topology, but overestimates it. This is to be expected, since policy gradient methods often fail to guarantee anything about value function as empirically established in [28].




Fig. 10: Comparison between the values estimated by DVP, PPO and the ground truth for the unicycle. Differences between (left) ground truth and DVP, (mid) ground truth and PPO and (right) DVP and PPO. PPO produces a value approximation which is further to the ground truth than DVP.




Fig. 11: The state trajectories predicted by DVP and PPO for unicycle. The x-axis denotes the horizon in knots, while the y-axis shows how the state vector of unicycle evolves. The goal is to reach the origin $[0, 0, 0]$

Fig. 11 shows the state trajectories computed by DVP and PPO. We observe that the trajectories computed by DVP are smoother than the trajectories output by PPO. From our experience, PPO was also more sensitive to small changes, either to the environment parameters (e.g discount factor) and algorithm hyperparameters which limited the experiments we have been able to carry out.

V. CONCLUSION

In this paper, we have proposed an algorithm to accurately learn the value function of an optimal control problem. Our contribution relies on a trajectory optimizer which produces good estimates of the cost-to-go over a finite preview horizon. We then leverage on supervised learning using (i) a Sobolev loss and (ii) a particular network architecture that we named Gauss residual network, to learn an accurate approximation of the value function and its derivatives. We can then use our value model to evaluate an approximation of the value at the end of the horizon, and its derivatives. By alternating between production of new optimal trajectories and refinement of the value and its derivatives, we set up a reinforcement loop which leads to a quick convergence to an accurate value approximation. We have proposed a complete evaluation of our method, on three typical control systems, with more qualitative results in the video. We are currently working toward using this approach on higher dimensional systems such as manipulators. We see our algorithm as a first step toward building a complete reinforcement learning algorithm for the continuous domain, able to fully exploits the derivatives of the simulator to reach superlinear convergence rate.

REFERENCES

[1] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fiedjeland, G. Ostrovski

- et al.*, “Human-level control through deep reinforcement learning,” *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [2] 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,” in *Int. Conf. on Learning Representations*, 2015.
 - [3] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, “Trust region policy optimization,” in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.
 - [4] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, “Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning,” in *Conference on Robot Learning*. PMLR, 2020, pp. 1094–1100.
 - [5] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
 - [6] R. Strudel, A. Pashevich, I. Kalevtykh, I. Laptev, J. Sivic, and C. Schmid, “Learning to combine primitive skills: A step towards versatile robotic manipulation §,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4637–4643.
 - [7] T. Haarnoja, S. Ha, A. Zhou, J. Tan, G. Tucker, and S. Levine, “Learning to walk via deep reinforcement learning,” *arXiv preprint arXiv:1812.11103*, 2018.
 - [8] Z. Xie, P. Clary, J. Dao, P. Morais, J. Hurst, and M. Panne, “Learning locomotion skills for cassie: Iterative design and sim-to-real,” in *Conference on Robot Learning*. PMLR, 2020, pp. 317–329.
 - [9] M. Hutter, C. Gehring, A. Lauber, F. Gunther, C. D. Bellicoso, V. Tsounis, P. Fankhauser, R. Diethelm, S. Bachmann, M. Blösch *et al.*, “Anymal-toward legged robots for harsh environments,” *Advanced Robotics*, vol. 31, no. 17, pp. 918–931, 2017.
 - [10] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
 - [11] J. Hwangbo, J. Lee, and M. Hutter, “Per-Contact Iteration Method for Solving Contact Dynamics,” *IEEE Robotics and Automation Letters*, vol. PP, p. 1, 2018.
 - [12] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, “Asynchronous methods for deep reinforcement learning,” in *International conference on machine learning*. PMLR, 2016, pp. 1928–1937.
 - [13] A. Albu-Schäffer, C. Ott, and G. Hirzinger, “A unified passivity-based control framework for position, torque and impedance control of flexible joint robots,” *The international journal of robotics research*, vol. 26, no. 1, pp. 23–39, 2007.
 - [14] N. Hogan, “Impedance control: An approach to manipulation: Part I—theory,” *Journal of Dynamic Systems Measurement and Control-Transactions of The Asme*, vol. 107, pp. 1–7, 1985.
 - [15] F. Grimmering, A. Meduri, M. Khadiv, J. Viereck, M. Wüthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols, J. Fiene, A. Badri-Spröwitz, and L. Righetti, “An open torque-controlled modular robot architecture for legged locomotion research,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3650–3657, 2020.
 - [16] B. Katz, J. Di Carlo, and S. Kim, “Mini cheetah: A platform for pushing the limits of dynamic quadruped control,” in *International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 6295–6301.
 - [17] E. Dantec, R. Budhiraja, A. Roig, T. Lembono, G. Saurel, O. Stasse, P. Fernbach, S. Tonneau, S. Vijayakumar, S. Calinon *et al.*, “Whole body model predictive control with a memory of motion: Experiments on a torque-controlled talos,” in *International conference on robotics and automation (ICRA)*, 2021.
 - [18] S. Levine and V. Koltun, “Guided policy search,” in *International conference on machine learning*. PMLR, 2013, pp. 1–9.
 - [19] N. Mansard, A. DelPrete, M. Geisert, S. Tonneau, and O. Stasse, “Using a memory of motion to efficiently warm-start a nonlinear predictive controller,” in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2986–2993.
 - [20] I. Mordatch and E. Todorov, “Combining the benefits of function approximation and trajectory optimization,” in *Robotics: Science and Systems*, vol. 4, 2014.
 - [21] I. Mordatch, N. Mishra, C. Eppner, and P. Abbeel, “Combining model-based policy search with online model learning for control of physical humanoids,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 242–248.
 - [22] I. Lenz, R. A. Knepper, and A. Saxena, “Deepmpc: Learning deep latent features for model predictive control,” in *Robotics: Science and Systems*. Rome, Italy, 2015.
 - [23] A. Tamar, G. Thomas, T. Zhang, S. Levine, and P. Abbeel, “Learning from the hindsight plan—episodic mpc improvement,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 336–343.
 - [24] M. Zhong, M. Johnson, Y. Tassa, T. Erez, and E. Todorov, “Value function approximation and model predictive control,” in *2013 IEEE symposium on adaptive dynamic programming and reinforcement learning (ADPRL)*. IEEE, 2013, pp. 100–107.
 - [25] B. Landry, H. Dai, and M. Pavone, “Seagull: Sample efficient adversarially guided learning of value functions,” in *Learning for Dynamics and Control*. PMLR, 2021, pp. 1105–1117.
 - [26] K. Lowrey, A. Rajeswaran, S. Kakade, E. Todorov, and I. Mordatch, “Plan online, learn offline: Efficient learning and exploration via model-based control,” in *Int. Conf. on Learning Representations*, 2019.
 - [27] D. Hoeller, F. Farshidian, and M. Hutter, “Deep value model predictive control,” in *Conference on Robot Learning*, 2020.
 - [28] A. Ilyas, L. Engstrom, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry, “A closer look at deep policy gradients,” in *Int. Conf on Learning Representations*, 2020.
 - [29] D. Q. Mayne, “Differential dynamic programming—a unified approach to the optimization of dynamic systems,” in *Control and Dynamic Systems*. Elsevier, 1973, vol. 10, pp. 179–254.
 - [30] E. Todorov and W. Li, “A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems,” in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 300–306.
 - [31] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, “Crocodyl: An efficient and versatile framework for multi-contact optimal control,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2536–2542.
 - [32] H. Son, J. W. Jang, W. J. Han, and H. J. Hwang, “Sobolev training for the neural network solutions of pdes,” *arXiv preprint arXiv:2101.08932*, 2021.
 - [33] L.-z. Liao and C. A. Shoemaker, “Advantages of differential dynamic programming over newton’s method for discrete-time optimal control problems,” Cornell University, Tech. Rep., 1992.
 - [34] W. M. Czarnecki, S. Osindero, M. Jaderberg, G. Świrszcz, and R. Pascanu, “Sobolev training for neural networks,” in *NIPS*, 2017, p. 4281–4290.
 - [35] T. M. Mitchell, S. B. Thrun *et al.*, “Explanation-based neural network learning for robot control,” *Advances in neural information processing systems*, pp. 287–287, 1993.
 - [36] R. Masuoka, “Noise robustness of ebnn learning,” in *International Conference on Neural Networks (IJCNN-93-Nagoya, Japan)*, vol. 2, 1993, pp. 1665–1668.
 - [37] J.-W. Lee and J.-H. Oh, “Hybrid learning of mapping and its jacobian in multilayer neural networks,” *Neural computation*, vol. 9, no. 5, pp. 937–958, 1997.
 - [38] J. B. Witkoskie and D. J. Doren, “Neural network models of potential energy surfaces: Prototypical examples,” *Journal of chemical theory and computation*, vol. 1, no. 1, pp. 14–23, 2005.
 - [39] S. Fleury, P. Soueres, J.-P. Laumond, and R. Chatila, “Primitives for smoothing mobile robot trajectories,” *IEEE transactions on robotics and automation*, vol. 11, no. 3, pp. 441–448, 1995.
 - [40] R. V. Florian, “Correct equations for the dynamics of the cart-pole system,” *Center for Cognitive and Neural Studies (Coneural), Romania*, 2007.
 - [41] R. Bellman, “Dynamic programming, princeton univ,” *Pres, Princeton*, 1957, 1957.
 - [42] —, “On the theory of dynamic programming,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 38, no. 8, p. 716, 1952.