

LEOC: A Principled Method in Integrating Reinforcement Learning and Classical Control Theory

Naifu Zhang

Department of Computer Science and Technology, Tsinghua University, Beijing, China

NZ248@CANTAB.NET

Nicholas Capel

Amazon Studios, North Building, 1620 26th St, Santa Monica, USA

NCCAPEL@AMAZON.COM

Abstract

There have been attempts in reinforcement learning to exploit *a priori* knowledge about the structure of the system. This paper proposes a hybrid reinforcement learning controller which dynamically interpolates a model-based linear controller and an arbitrary differentiable policy. The linear controller is designed based on local linearised model knowledge, and stabilises the system in a neighbourhood about an operating point. The coefficients of interpolation between the two controllers are determined by a scaled distance function measuring the distance between the current state and the operating point. The overall hybrid controller is proven to maintain the stability guarantee around the neighborhood of the operating point and still possess the universal function approximation property of the arbitrary non-linear policy. Learning has been done on both model-based (PILCO) and model-free (DDPG) frameworks. Simulation experiments performed in OpenAI gym demonstrate stability and robustness of the proposed hybrid controller. This paper thus introduces a principled method allowing for the direct importing of control methodology into reinforcement learning.

Keywords: Reinforcement Learning, Robotics, Artificial Intelligence, Robustness

1. Introduction

In recent years, the rise of deep learning has unlocked a new class of function approximation and representation techniques that enable easy discovery of low-dimensional features in extremely high dimensional data (Bengio et al.). These techniques have naturally been applied to the field of Reinforcement Learning (RL) - a practice that allows agents to learn behaviours through interactions with an environment in order to achieve some desired outcome.

However, the output policies of the algorithm are opaque, and the relationship between the states and actions cannot be easily understood. Crucially, the policy learnt from optimising a reward function provides no performance guarantee at certain operating points. For instance, a popular deep RL algorithm Deep Deterministic Policy Gradient (DDPG) exhibits inconsistent performances (Lillicrap et al., 2016), often failing to deliver a truly stable controller (Gu et al., 2016a). When using these techniques for medical purposes (Ferrarese et al., 2016) or for controlling a chemical plant, control actions with stability and overshoot bounds in an operating point (e.g. a pH or pressure level) could be critical. The ability of users to understand and trust the outputs of these trained controller policies can be key to adoption (Ribeiro et al., 2016).

When we have a system with known local linear dynamics, we can use a powerful suite of tools from control theory to design linear policies with certain highly desired properties, such as better

system stability performance in terms of input responses and increased robustness to poor dynamics models. Our **Locally Enforced Optimal Controller (LEOC)** hybrid policy allows for the manual specification of such a linear controller within a desired region of control, and outside that region, smoothly switches to an arbitrary non-linear policy that is rich enough to model large classes of functions. We theoretically prove that our hybrid policy exhibits the desirable properties of both the linear controller and the learnt arbitrary policy, and verify this experimentally. This paper thus introduces a principled method for controller design that combines tools from control theory with reinforcement learning.

In this paper, LEOC is utilised within the model-based framework of PILCO (Deisenroth et al.) and the model-free framework of DDPG (Lillicrap et al., 2016), showing that it can in theory be used with other model-based or model-free systems with little to no modification.

2. Related Work

The idea of incorporating existing knowledge into a reinforcement learning model is well-studied. Such techniques include, but are not limited to - Meta Learning (Finn et al., 2017; Mishra et al., 2017; Duan et al.) and Hierarchical Reinforcement Learning (Nachum et al.; Vezhnevets et al.; Frans et al., 2017), wherein an agent is trained over a distribution of task and learns to quickly adapt its existing algorithms to a new task or learns a meta-policy over subgoals.

The idea of incorporating techniques from control theory into reinforcement learning has also been studied. Techniques such as gain scheduling (Leith and Leithead, 2000), and iLQG (Todorov and Li, 2005) have been employed. However, a major limitation of these techniques is that they require full access to the underlying physics model and its derivatives. In applications, there are many systems that are well studied about certain operating points, with unknown global behaviour outside of them. Rather than embedding prior knowledge into the environment or reward function, we deviate from existing approaches by enforcing the optimal control into the policy itself.

In recent times, the desire to use reinforcement learning in safety-critical settings has inspired some work in formal methods for reinforcement learning. These methods include shielding (Alshiekh et al., 2017), logically constrained learning (Bohrer et al., 2018), justified speculative control (Fulton and Platzer, 2018) and constrained Bayesian optimization (Ghosh et al.). These approaches provide formal safety guarantees for reinforcement learning by defining safety specifications, and constraining the learning to a known safe subset of the state space. Yet another subfield of interest is that of robust reinforcement learning, where a policy is considered to be optimal if it has the maximum worst case return (Garcia et al.). Strands of research involve augmenting the MDP model to create a robust MDP model that provides probabilistic guarantees (Wiesemann et al., 2013), or by training adversarial agents to provide conflicting inputs to maximize worst case performance (Pinto et al.).

Our technique is unique in that instead of editing the MDP framework and augmenting it with partial observations, probabilistic guarantees, formal methods or highly parameterised networks, we simply allow for the direct importation of control methodology focused on the local linear approximation about the operating point. This allows for very simple application of well-known control theory, and provides performance guarantees about operating points in which systems spend a majority of their time in.

We hypothesize that this hybrid controller will be very useful in two different kinds of systems: first, systems with general non-linear dynamics, but operating points about which stability


is desired; second, systems with very well-studied local behaviour, but unknown global behaviour. Examples include combustion (Fuerhapter et al., 2004), continuous non-isothermal stirred tank reactors (Alexandridis and Sarimveis, 2011), and catalytic crackers (Arbel et al., 1995).

3. Background

3.1. RL Set-up

We utilize the base reinforcement learning agent and environment setup in Figure 1. Notice that the agent can take any arbitrary non-linear policy. Our model interpolates the outputs of the non-linear controller with a linear controller that takes over at an operating point centered at some position in the state space. Our system functions as long as there exists an agent that accepts the state as input and the action as output.

Figure 1: Reinforcement Learning setup



3.2. Markov Decision Process (MDP)

The formal language that we adopt in this paper is one of the Markov Decision Process (MDP) (Bellman; Puterman, 1994; Ghavamzadeh et al., 2016) where:

- $\mathbf{x} \in X$ denotes a state from a set of measurable, observable and fully controllable possible states
- $\mathbf{u} \in U$ denotes an action drawn from a set of possible actions
- $P(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t) \in P(S)$ refers to the transition dynamics, which is the probability of distribution over the next states conditioned on the previous actions and states
- $P_0 \in P(S)$ denotes the distribution of the initial state
- $c(\mathbf{x}, \mathbf{u}) \in P(\mathbb{R})$ represents the reward obtained when an action \mathbf{u} is taken in state \mathbf{x}

We also introduce a controller or policy function, $\pi(\mathbf{x}_t)$ such that $\mathbf{u}_t = \pi(\mathbf{x}_t)$. The problem of trying to discover the ideal policy $\pi^*(\mathbf{x}_t)$ to achieve the best possible reward thus reduces to:

$$\pi^*(\mathbf{x}_t) = \arg \min_{\pi} \sum_{\forall t} \mathbb{E}[c(\mathbf{x}_t)] \quad (1)$$

3.3. Learning Frameworks

The optimisation of the baseline controller policies is done in two reinforcement learning frameworks, namely PILCO (Deisenroth et al.) and DDPG (Lillicrap et al., 2016). PILCO belongs to the class of model-based frameworks, whilst DDPG is a model-free and off-policy algorithm.

These models are chosen to represent two broad categories of reinforcement learning agents (model-based vs model-free). We hope to show in the following sections that our hybrid controller generalizes to a wide variety of agents.

Briefly, a key difference is that PILCO attempts to learn the unknown deterministic transition function $f(\cdot)$, where

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) \quad (2)$$

The probabilistic dynamics model is implemented using a Gaussian process (GP). The GP allows us to perform next-step prediction, propagating state distribution for model evaluation and optimisation.

DDPG, on the other hand, is a model-free algorithm based on the deterministic policy gradient and deep Q-learning that can operate over continuous action spaces. It consists of a critic learned using the Bellman equation as in Q-learning (Watkins and Dayan, 1992), and an actor that is optimised via policy gradient methods (Silver et al.).

A discussion of these algorithms at length is not within the purview of this paper, and we refer interested reader to the original papers - Deisenroth (2010) Deisenroth et al. and Lillicrap et al. (2016).

4. Overview of Hybrid Controller

This section gives a treatment of our hybrid controller policy which consists of a linear and a non-linear part.

At the operating point, we define a standard linear controller:

$$G(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b} \quad (3)$$

where $\mathbf{x} \in \mathbb{R}^D$ is the state, $\mathbf{W} \in \mathbb{R}^{F \times D}$ is a parameter matrix of weights and $\mathbf{b} \in \mathbb{R}^F$ is a bias vector. In each control dimension $d \in [1, \dots, D]$, the control action is given by a linear combination of states and bias. See subsection 5.2 for more detailed discussions on how to obtain \mathbf{W} and \mathbf{b} in the experiments.

Outside of the neighbourhood of operating point, we switch to a control policy that takes the form of an arbitrary (usually non-linear) function, represented by $H(x)$.

In the case of PILCO, we choose a radial basis function. In the case of the DDPG, we choose a deep neural network. The requirements of the chosen $H(x)$ are that it is a universal function approximator (i.e. $\exists H$ such that $\|H - f\|_1 < \epsilon$), that it is differentiable with respect to any of the model parameters (i.e. it is trainable), and it accepts the state as an input and outputs the control action.

The overall hybrid policy is thus,

$$\pi(\mathbf{x}) = r(\mathbf{x}) \underbrace{(\mathbf{W}\mathbf{x} + \mathbf{b})}_{G(\mathbf{x})} + (1 - r(\mathbf{x}))H(\mathbf{x}) \quad (4)$$

In the limit that the system is at the operating point, the control action should be that of the linear controller. In general, the closer the system is to the operating point, the more the hybrid controller is to approximate the linear controller. The hybrid controller could therefore be parameterized as a weighted average of the linear and non-linear controllers, with $r(\mathbf{x})$ determining the dependence on either.

Naturally, $d(\mathbf{x})$ (Equation 5b) can be interpreted as a Euclidean distance weighted in different dimensions by $\frac{1}{\lambda_i}$, where $\lambda_i > 0$, with $d(\mathbf{x})$ being computed between the current state \mathbf{x} and the operating point \mathbf{a} .

$$r(\mathbf{x}) = \frac{1}{(1 + d(\mathbf{x}))^2} \quad (5a)$$

$$d(\mathbf{x}) = (\mathbf{x} - \mathbf{a})^T \mathbf{\Lambda}^{-1} (\mathbf{x} - \mathbf{a}) \quad (5b)$$

$$\mathbf{\Lambda} = \text{diag}(\lambda_i) \quad (5c)$$

We also note that λ_i relates to the relative weight of each dimension. It could not be known *a priori*, and thus has to be learnt too during the optimisation procedure.

4.1. Stability Analysis

To prove that local asymptotic stability is maintained, first we consider the linearization of the hybrid controller

$$\pi(\mathbf{x}) = \pi(\mathbf{a}) + \nabla \pi(\mathbf{a})(\mathbf{x} - \mathbf{a}) + \dots \quad (6)$$

In order to evaluate the linearization of $\pi(x)$:

$$\pi(\mathbf{x}) = r(\mathbf{x})G(\mathbf{x}) + (1 - r(\mathbf{x}))H(\mathbf{x}) \quad (7a)$$

$$\nabla \pi(\mathbf{x}) = \nabla r(\mathbf{x})G(\mathbf{x}) + r(\mathbf{x})\nabla G(\mathbf{x}) + (1 - r(\mathbf{x}))\nabla H(\mathbf{x}) - \nabla r(\mathbf{x})H(\mathbf{x}) \quad (7b)$$

When $\mathbf{x} = \mathbf{a}$, $r(\mathbf{a}) = 1$ and $\nabla r(\mathbf{a}) = 0$, we can substitute \mathbf{x} with \mathbf{a} , obtaining the following

$$\pi(\mathbf{a}) = G(\mathbf{a}) \quad (8a)$$

$$\nabla \pi(\mathbf{a}) = \nabla G(\mathbf{a}) \quad (8b)$$

Equation 8 shows that the local linearization of the hybrid policy reduces to the linear controller $G(\mathbf{a})$. Therefore, if the linear controller is stable about the operating point, the hybrid controller is also stable about the operating point.


4.2. Proof of Universal Function Approximation

With $r(\mathbf{x})$ defined in Equation 5a, where $\lambda_i > 0$, $\mathbf{a} \in \mathbb{R}^D$, $\mathbf{W} \in \mathbb{R}^{F \times D}$, $\mathbf{b} \in \mathbb{R}^F$, $\mathbf{x} \in \mathbb{R}^D$ and the other parameters as defined in Equation 4, the following can be seen by applying Holder's inequality:

$$\|f - r(\mathbf{x})G(\mathbf{x}) - (1 - r(\mathbf{x}))H(\mathbf{x})\| \leq \|f - H(\mathbf{x})\| + \|r(\mathbf{x})H(\mathbf{x})\| + \|r(\mathbf{x})G(\mathbf{x})\| \quad (9a)$$

$$\leq \|f - H(\mathbf{x})\| + \|r(\mathbf{x})\| \|H(\mathbf{x})\|_\infty + \|r(\mathbf{x})G(\mathbf{x})\| \quad (9b)$$

Figure 2: Hybrid controller architecture: \mathbf{x} is passed into the non-linear controller $H(\mathbf{x})$, linear controller $G(\mathbf{x})$ and $r(\mathbf{x})$. Together, these make up the hybrid controller policy $\pi(\mathbf{x})$.



The rest of the analysis focuses on the case when $D = 1$, but can be trivially extended to arbitrary dimensions, D .

$$\|r(x)\| = \int_{-\infty}^{\infty} \left| \frac{1}{(\lambda^{-1}x^2 + 1)^2} \right| = \frac{\pi\sqrt{\lambda}}{2} \quad (10a)$$

$$\|xr(x)\| = \int_{-\infty}^{\infty} \left| \frac{x}{(\lambda^{-1}x^2 + 1)^2} \right| = \lambda \quad (10b)$$

We know that as $H(\mathbf{x}) \in \mathbb{L}^1$, $\|H(\mathbf{x})\|_{\infty} < \infty$. Moreover, $\|G(x)r(x)\| = \|(\mathbf{W}x + \mathbf{b})r(x)\| \leq \mathbf{W} \|xr(x)\| + \mathbf{b} \|r(x)\|$. Hence $\|r(\mathbf{x})\| \|H(\mathbf{x})\|_{\infty}$ and $\|r(\mathbf{x})G(\mathbf{x})\|$ can be made arbitrarily small by tuning the value of λ_i .

As H is a universal function approximator by construction, this implies $\|f - H(\mathbf{x})\| < \epsilon$, where $f \in \mathbb{L}^1$. Therefore:

$$\|f - H(\mathbf{x})\| + \|r(\mathbf{x})\| \|H(\mathbf{x})\|_{\infty} + \|r(\mathbf{x})G(\mathbf{x})\| < \epsilon + \gamma \quad (11)$$

where γ is an arbitrarily small positive number that can be achieved by choosing a suitable λ_i . We have thus proven that the universal function approximation property is maintained for the LEOC controller function.

5. Experiments

We evaluate our proposed approach on three RL tasks: Swing-up Pendulum, CartPole, and MountainCar.

5.1. Environments




Figure 3: Experimental environments.

In each of the environments, the agent is initialised off the operating point and the objective is to effect a continuous force on the agent such that it stabilises at the operating point.

Swing-up Pendulum OpenAI’s Pendulum-v0 (Figure 3a) (Brockman et al., 2016) is used. The objective is to swing the pendulum up and balance it in the state $(\theta, \dot{\theta}) = (0, 0)$.

CartPole A non-discrete force u is applied to the cart (Figure 3b). The objective is to swing the pendulum up in the middle of the track and balance it in the state $(x, \dot{x}, \theta, \dot{\theta}) = (0, 0, 0, 0)$.

MountainCar In our environment (Figure 3c), the cart is rewarded not only for reaching but also staying at the top of the mountain, i.e. $(x, \dot{x}) = (0, 0)$.

5.2. Linear Controller

Each of these systems follows the state-space equation in the canonical form at the operating point, where \mathbf{x} represents the observed states, and \mathbf{y} the output.

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \quad (12a)$$

$$\mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{u} \quad (12b)$$

The cost function $c(\mathbf{x}, \mathbf{u})$ in each environment takes the form of the squared distance function, i.e. $c(\mathbf{x}, \mathbf{u}) = -(\mathbf{x} - \mathbf{a})^T \mathbf{K}(\mathbf{x} - \mathbf{a}) - k\mathbf{u}^T \mathbf{u}$.

With access to this model knowledge (albeit limited to the operating point) and cost function, we can design linear controller $G(\mathbf{x})$ using methods developed in optimal control theory.

We derive the optimal state-feedback control gains to achieve closed-loop stable and high performance controller. The gain matrix \mathbf{K} is computed via the Linear Quadratic Regulator (LQR). The feedback action could thus be synthesised by the following

$$\mathbf{u} = -\mathbf{K}\mathbf{x} \quad (13)$$

Equation 13 is simply Equation 3 with $\mathbf{W} = -\mathbf{K}$ and $\mathbf{b} = 0$.

The rest of the LEOC hybrid controller, namely $H(\mathbf{x})$ and $r(\mathbf{x})$, are learnt from training data in an RL framework. For all our experiments we initialize \mathbf{A} to be the identity matrix. Optimisers used are Adam and BFGS respectively for DDPG and PILCO.

6. Results and Discussion

In this section, we focus on providing experimental verification some of the theoretical statements we made earlier. In particular, we focus on substantiating the claims in subsection 4.1 that the LEOC hybrid controller behaves very much like a linear controller around the operating point, thereby also inheriting the desirable properties of the linear controller, namely, **stability**, **transient responses**, and **robustness**. All results in this section are expected, and serve as verification of theory.

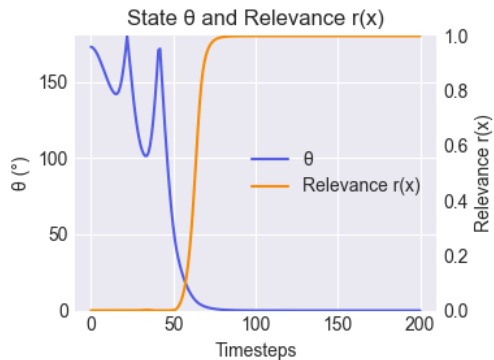
6.1. Training and General Comments

But first, we make a few general comments.

Figure 4 illustrates the internals of a trained hybrid controller in the Swing-up Pendulum environment. We can see the linear controller increasingly taking over when θ gets closer to the operating point. The hybrid controller is hence correctly interpolating between the linear and non-linear policies as designed in section 4.

With respect to Figure 5, we make a couple of observations. First, trivially, convergence in awards can be achieved in training LEOC hybrid controller. This supports the theory in subsection 4.2 that the hybrid controller could approximate the control function as well as the baselines.

Figure 4: The dynamics of a trained Pendulum hybrid controller are shown. The blue line plots θ whereas the orange line plots relevance $r(x)$.



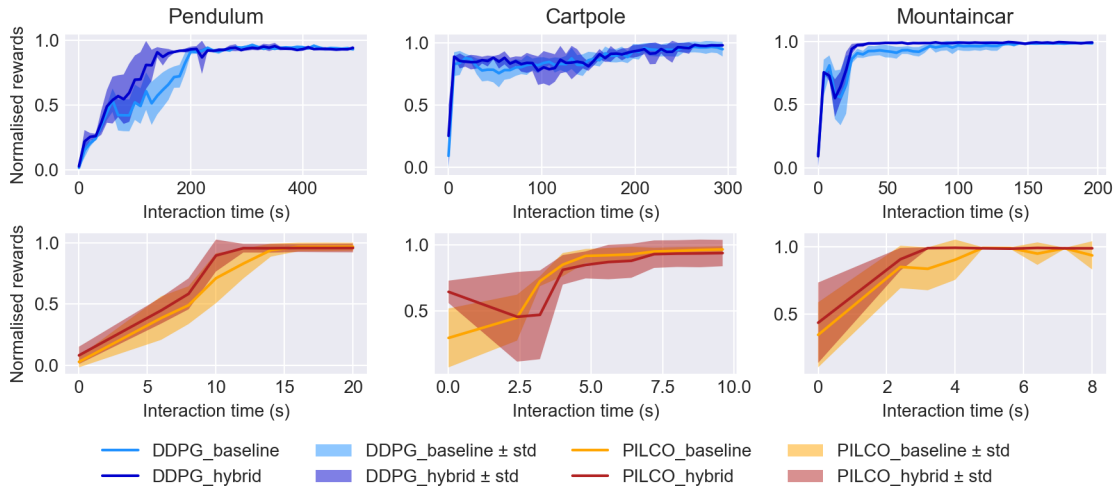


Figure 5: Rewards against interaction time taken to train. The top row compares DDPG baseline with LEOC hybrid implemented in the DDPG framework; the bottom row compares PILCO baseline with LEOC hybrid implemented in PILCO. PILCO is a vastly more data efficient approach than DDPG, and hence takes shorter interaction time.

Second, pre-specifying part of the hybrid controller in advance accelerates training, as measured by total interaction time. We note that our hybrid controller converges very quickly once the linear part is activated with $r \approx 1$. The baseline controllers, on the other hand, may require extra system interaction time to explicitly learn the dynamics about the operating point, especially in the Pendulum and MountainCar environments.

6.2. Stability and Transient Response

Next, we examine the system’s behavior to an impulse and step input when it is at the operating point. Figure 6 shows the impulse response and the step response of representative examples of the various policies.

The impulse response gives insights to the stability and steady state characteristics of the systems. We note that the steady state errors of both DDPG and PILCO baselines are non-zero. There is in fact also no guarantee on the stability of the DDPG and PILCO baselines. As shown in Figure 6 and observed by Gu et al. (2016b), DDPG baseline policy could oscillate around the operating point, therefore not being stable in the Lyapunov stability sense. This serves as empirical evidence to back up the claim made in section 1: that cost-function based approaches in end-to-end reinforcement learning can lead to less desirable properties locally.

The LEOC hybrid controllers, on the other hand, deliver better characteristics in a host of metrics, including steady state errors, peak overshoot and settling time, with smaller means and stds, and guaranteed performance (Table 1). The hybrid controllers essentially share the same transient responses as linear controllers, as proven in subsection 4.1. This is further evidence that control performance is preserved across the linear and LEOC hybrid controllers at the operating point.

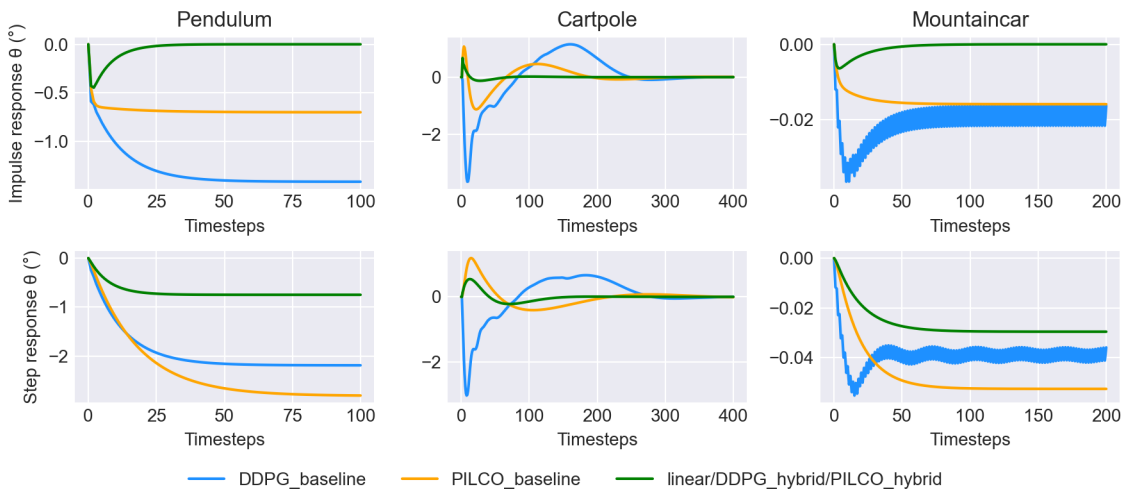


Figure 6: System transient responses in state θ or x , whichever is appropriate for the environment. The top row shows impulse response; the bottom row shows step response. The DDPG_hybrid and PILCO_hybrid lines may not be visible as they overlap with the line for linear controller.

Pendulum	Steady-state Error ($^{\circ}$)		Overshoot ($^{\circ}$)		Settling Time (timesteps)	
	mean	std	mean	std	mean	std
DDPG baseline	-1.79	0.66	-	-	29.33	16.13
DDPG hybrid	-2.92e-13	0.00	0.44	0.00	26.00	0.00
PILCO baseline	0.33	6.05	5.26	3.33	49.33	21.23
PILCO hybrid	-3.13e-13	-2.33e-16	0.47	2.09e-03	26.00	0.00
Linear	-2.95e-13	0.00	0.45	0.00	26.00	0.00

Table 1: Impulse response metrics in the Pendulum environment. The LEOC hybrid controllers are essentially the same as linear controllers (in bold). We show only one environment here as an illustration - the same point can be made in the other environments.

6.3. Robustness Analysis

Finally, we evaluate the robustness and generalization of the learned policy with respect to varying test conditions. We train the policy based on certain parameter (e.g. cart mass or gravitational acceleration constant g) values. At test time, however, we evaluate the learned policies by changing parameter values and estimating the average cumulative rewards of a system initialised in the operating point. Specifically, we vary the parameter values from 50% to 500% of the original values used in training.

The rewards are used as a proxy for stability, whereby an unstable system would return a worse reward. As seen in Figure 7, our LEOC hybrid policies return mostly 0 rewards despite parameter noise and are more robust than either PILCO or DDPG baselines. In fact, LEOC hybrid policies perform as well as linear controllers.

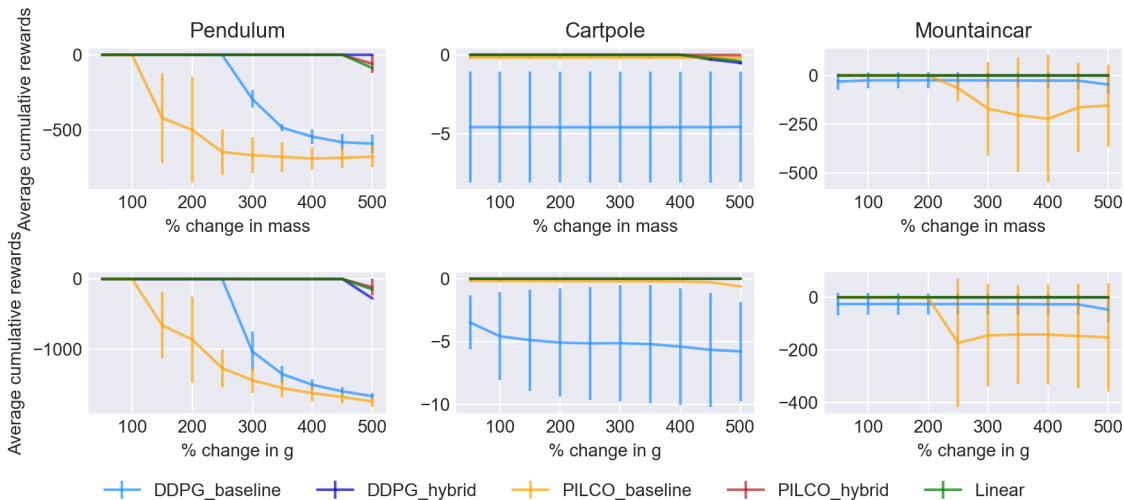


Figure 7: Mean cumulative rewards \pm std against parameter noise. The top row shows the rewards when mass parameters in the environments are changed; the bottom row shows the rewards when g in the environments is changed. The DDPG_hybrid and PILCO_hybrid lines may not be too visible as they overlap with the line for linear controller.

All in all, we observe that the baseline policies optimised through the traditional reinforcement learning route cannot guarantee **stability**, **transient response performance** or **robustness**. The LEOC hybrid controller, on the other hand, behaves as a linear controller near the operating point, and adopts all the associated desirable as theorised in subsection 4.1.

7. Conclusion

In this paper we introduce LEOC. We prove two desirable properties of our controller - that it is stable about a desired operating point for a given system, and that it retains the ability to serve as a universal function approximator. Via a series of experiments we establish the effectiveness of this new controller by showing similar or better system convergence with lower interaction time. We also demonstrate that it has high robustness to system modelling uncertainties, with performance similar to a pre-designed linear controller, with much better performance than the baseline. We thus demonstrate that our new controller is a practical method for incorporating existing model knowledge into reinforcement learning systems via manipulation of controller models about linearised operating points.

Full implementation of the system can be found at <https://github.com/nrjc/LEOC>.

Acknowledgments

We thank the authors of the GPflow, PILCO and Tensorflow Agents packages for part of our implementation of PILCO and DDPG respectively.

References

- Alex Alexandridis and Haralambos Sarimveis. *Control of processes with multiple steady states using MPC and RBF neural networks*, volume 29. Elsevier B.V., 2011. doi: 10.1016/B978-0-444-53711-9.50140-1. URL <https://linkinghub.elsevier.com/retrieve/pii/B9780444537119501401>.
- Mohammed Alshiekh, Roderick Bloem, Ruediger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe Reinforcement Learning via Shielding. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pages 2669–2678, 8 2017. URL <http://arxiv.org/abs/1708.08611>.
- Arnon Arbel, Irvan H. Rinard, Reuel Shinnar, and Ajit V. Sapre. Dynamics and Control of Fluidized Catalytic Crackers. 2. Multiple Steady States and Instabilities. *Industrial and Engineering Chemistry Research*, 34(9):3014–3026, 9 1995. ISSN 15205045. doi: 10.1021/ie00048a013. URL <https://pubs.acs.org/doi/pdf/10.1021/ie00048a013>.
- Richard Bellman. A Markovian decision process. Technical report. URL <https://www.jstor.org/stable/24900506>.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation Learning: A Review and New Perspectives. Technical report. URL <http://www.image-net.org/challenges/LSVRC/2012/results.html>.
- Brandon Bohrer, Yong Kiam Tan, Stefan Mitsch, Magnus O Myreen, André Platzer, and Yong Kiam Tan. VeriPhy: Verified Controller Executables from Verified Cyber-Physical System Models. *dl.acm.org*, 14(18):617–630, 6 2018. doi: 10.1145/3192366.3192406. URL <https://doi.org/10.1145/3192366.3192406>.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. 6 2016. URL <http://arxiv.org/abs/1606.01540>.
- M Deisenroth, CE Rasmussen of the 28th ..., and undefined 2011. PILCO: A model-based and data-efficient approach to policy search. *ausy.informatik.tu-darmstadt.de*. URL http://www.ausy.informatik.tu-darmstadt.de/uploads/Publications/Deisenroth_ICML_2011.pdf.
- MP Deisenroth. *Efficient reinforcement learning using Gaussian processes*. 2010. URL <https://books.google.com/books?hl=en&lr=&id=Vt0q--oGhZ4C&oi=fnd&pg=PR7&dq=E%3Fcient+Reinforcement+Learning+using+Gaussian+Processes&ots=syD3mSiJny&sig=zFBGgKHHhS7kaMZk1N9DCDJXKPk>.
- Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. Under review as a conference paper at ICLR 2017 RL 2 : FAST REINFORCEMENT LEARNING VIA

- SLOW REINFORCEMENT LEARNING. Technical report. URL <https://arxiv.org/abs/1611.02779>.
- Alessia Ferrarese, Giada Pozzi, Felice Borghi, Alessandra Marano, Paola Delbon, Bruno Amato, Michele Santangelo, Claudio Buccelli, Massimo Niola, Valter Martino, and Emanuele Caspasso. Malfunctions of robotic system in surgery: Role and responsibility of surgeon in legal point of view, 1 2016. ISSN 19112092. URL </pmc/articles/PMC5329842/?report=abstracthttps://www.ncbi.nlm.nih.gov/pmc/articles/PMC5329842/>.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. (2015):5–7, 3 2017. URL <http://arxiv.org/abs/1703.03400>.
- Kevin Frans, Jonathan Ho, Xi Chen, Pieter Abbeel, and John Schulman. Meta learning shared hierarchies, 10 2017.
- A. Fuerhapter, E. Unger, W. F. Piock, and G. K. Fraidl. The new AVL CSI engine - HCCI operation on a multi cylinder gasoline engine. In *SAE Technical Papers*. SAE International, 3 2004. ISBN 0768013194. doi: 10.4271/2004-01-0551. URL <https://www.sae.org/publications/technical-papers/content/2004-01-0551/>.
- Nathan Fulton and André Platzer. Safe Reinforcement Learning via Formal Methods Toward Safe Control Through Proof and Learning. Technical report, 2018. URL www.aaai.org.
- J Garcia, F Fernández Journal of Machine Learning, and undefined 2015. A comprehensive survey on safe reinforcement learning. *pdfs.semanticscholar.org*. URL <https://pdfs.semanticscholar.org/c0f2/c4104ef6e36bb67022001179887e6600d24d.pdf>.
- Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian Reinforcement Learning: A Survey. *Foundations and Trends in Machine Learning*, 8(5-6):359–483, 9 2016. doi: 10.1561/22000000049. URL <http://arxiv.org/abs/1609.04436http://dx.doi.org/10.1561/22000000049>.
- Shromona Ghosh, Felix Berkenkamp, Gireeja Ranade, Shaz Qadeer, and Ashish Kapoor. Verifying Controllers Against Adversarial Examples with Bayesian Optimization. Technical report. URL <https://ieeexplore.ieee.org/abstract/document/8460635/>.
- Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E. Turner, and Sergey Levine. Q-Prop: Sample-Efficient Policy Gradient with An Off-Policy Critic. 11 2016a. URL <http://arxiv.org/abs/1611.02247>.
- Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, Sergey Levine, and Slevine@google Com. Continuous Deep Q-Learning with Model-based Acceleration. Technical report, 2016b.
- Doug J. Leith and W. E. Leithead. Survey of gain-scheduling analysis and design, 7 2000. ISSN 00207179.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In

- 4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR, 9 2016. URL <https://goo.gl/J4PIAz>.
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A Simple Neural Attentive Meta-Learner. *arXiv*, abs/1707.0:14 pp.–14 pp., 7 2017. URL <http://arxiv.org/abs/1707.03141>.
- Ofir Nachum, Google Brain, Shixiang Gu, Honglak Lee, and Sergey Levine. Data-Efficient Hierarchical Reinforcement Learning. Technical report. URL <https://sites.google.com/view/efficient-hrl>.
- Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust Adversarial Reinforcement Learning. Technical report.
- ML Puterman. Markovian decision problems. 1994.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. 2 2016. URL <http://arxiv.org/abs/1602.04938>.
- David Silver, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic Policy Gradient Algorithms. Technical report.
- Emanuel Todorov and Weiwei Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *Proceedings of the American Control Conference*, volume 1, pages 300–306, 2005. doi: 10.1109/acc.2005.1469949.
- Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. FeUdal Networks for Hierarchical Reinforcement Learning. Technical report. URL <https://arxiv.org/abs/1703.01161>.
- Christopher J C H Watkins and Peter Dayan. Q-Learning. Technical report, 1992.
- Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust Markov Decision Processes. *Mathematics of Operations Research*, 38(1):153–183, 2 2013. ISSN 1526-5471. doi: 10.1287/moor.1120.0566. URL <http://pubsonline.informs.org><https://doi.org/10.1287/moor.1120.0566><http://www.informs.org><http://dx.doi.org/10.1287/moor.1120.0566>.