An Iterative Path Integral Reinforcement Learning Approach

Evangelos Theodorou, Jonas Buchli, Freek Stulp, and Stefan Schaal

Although significant progress has been made in learning control from trajectory roll-outs, performing Reinforcement Learning in high dimensional and continuous state action spaces remains a challenging problem. Especially in the domain of humanoid robot control, the high dimensionality, the lack of accurate dynamical models, and the hybrid and unstable nature of the underlying dynamics due to contact with the environment make the application of traditional model based and model free Reinforcement Learning methods rather difficult.

In this short paper we present a new probabilistic method [1] for reinforcement learning that is derived based from the Bellman principle for continuous state-action spaces. Under the constraint of nonlinear stochastic dynamics, the result of the Bellman principle is the stochastic Hamilton Jacobi Bellman (HJB) equation, which is a nonlinear and second order Partial Differential Equation (PDE). Exponentiation of the value function and the use of mild assumptions on the exploration noise and control cost transforms the nonlinear PDE into a linear and second order PDE. The resulting PDE is the so called Backward Chapman Kolmogorov PDE. The use of the first of the nonlinear Feynman Kac Lemmas provides a bridge between Forward Stochastic Differential Equations (FSDEs) and Partial Differential Equations (PDEs) [2] of the type of a Backward Chapman Kolmogorov PDE. The FSDEs are used for sampling parts of the state space which are relevant to the task goal. Moreover sampling is performed in an iterative fashion to localize exploration and avoid the curse of dimensionality. The resulting algorithm, called Policy Improvement with Path Integrals (PI^2), takes on a surprisingly simple form, has no algorithmic open tuning parameters besides the exploration noise, and performs numerically robustly in high dimensional learning problems. It also makes an interesting connection to previous work on RL based on probability matching [3], [4], [5] and explains why probability matching algorithms can be successful.

Policy Improvement with Path Integrals

Let us define a finite-horizon reward function for a trajectory τ_i starting at time t_i in state \mathbf{x}_{t_i} and ending at time t_N as $R(\tau_i) = \phi_{t_N} + \int_{t_i}^{t_N} r_t dt$ with $\phi_{t_N} = \phi(x_{t_N})$ denoting a terminal reward at time t_N and r_t denoting the immediate reward at time t. The goal is to find the controls θ_t that minimize the cost function $V(\mathbf{x}_{t_i}) = V_t =$ $\min_{\boldsymbol{\theta}_{1,i+1}} E_{\boldsymbol{\tau}_{i}}[R(\boldsymbol{\tau}_{i})]$ where the expectation E[.] is taken over all trajectories starting at \mathbf{x}_{t_i} . We consider the rather general control system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}_t, t) + \mathbf{G}(\mathbf{x}_t) (\boldsymbol{\theta}_t + \epsilon_t) =$ $\mathbf{f}_t + \mathbf{G}_t (\boldsymbol{\theta}_t + \epsilon_t)$ with $\mathbf{x}_t \in \Re^{n \times 1}$ denoting the state of the system, $\mathbf{G}_t = \mathbf{G}(\mathbf{x}_t) \in \Re^{n \times p}$ the control matrix, $\mathbf{f}_t =$

TABLE I

PSEUDOCODE OF THE \mathbf{PI}^2 ALGORITHM FOR A 1D PARAMETERIZED Policy (Note that the discrete time step dt was absorbed as a CONSTANT MULTIPLIER IN THE COST TERMS).

- Given:
 - An immediate cost function $r_t = q_t + \boldsymbol{\theta}_t^T \mathbf{R} \boldsymbol{\theta}_t$
 - A terminal cost term ϕ_{t_N}
 - A stochastic parameterized policy $\mathbf{a}_t = \mathbf{g}_t^T (\boldsymbol{\theta} + \boldsymbol{\epsilon}_t)$
 - The basis function \mathbf{g}_{t_i} from the system dynamics The variance Σ_{ϵ} of the mean-zero noise ϵ_t

 - The initial parameter vector $\boldsymbol{\theta}$
- **Repeat** until convergence of the trajectory cost *R*:
 - step 1: Create K roll-outs of the system from the same start state \mathbf{x}_0 using stochastic parameters $\boldsymbol{\theta} + \boldsymbol{\epsilon}_t$ at every time step step 2: For all K roll-outs, compute: * step 2.1: $\mathbf{M}_{t_j,k} = \frac{\mathbf{R}^{-1}\mathbf{g}_{t_j,k}}{\mathbf{g}_{t_j,k}^T \mathbf{R}^{-1}\mathbf{g}_{t_j,k}}$ * step 2.2: Compute the cost for each sampled trajectory: $S(\boldsymbol{\tau}_{i,k}) = \oint_{T=1}^{k} e_{t,k} + \sum_{j=i}^{N-1} q_{t_j,k} + \frac{1}{2} \sum_{j=i+1}^{N-1} (\boldsymbol{\theta} + \boldsymbol{\theta}_{i,k})$ $\mathbf{M}_{t_j,k} \boldsymbol{\epsilon}_{t_j,k} T \mathbf{R} (\boldsymbol{\theta} + \mathbf{M}_{t_j,k} \boldsymbol{\epsilon}_{t_j,k})$ * step 2.3: $P\left(\boldsymbol{\tau}_{i,k}\right) = \frac{e^{-\frac{1}{\lambda}S(\boldsymbol{\tau}_{i,k})}}{\sum_{k=1}^{K} [e^{-\frac{1}{\lambda}S(\boldsymbol{\tau}_{i,k})}]}$ - **step 3: For** all *i* time steps, compute: * step 3.1: $\delta \theta_{t_i} = \sum_{k=1}^{K} \left[P\left(\boldsymbol{\tau}_{i,k} \right) \mathbf{M}_{t_i,k} \, \boldsymbol{\epsilon}_{t_i,k} \right]$ - step 4: Compute $\delta \theta^m = \frac{\sum_{i=0}^{N-1} (N-i) \, w_{t_i}^m \, \delta \theta_{t_i}^m}{\sum_{i=0}^{N-1} (N-i) \, w_{t_i}^m}$ - step 5: Update $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \delta \boldsymbol{\theta}$ step 6: Create one noiseless roll-out to check the trajectory cost $R = \phi_{t_N} + \sum_{i=0}^{N-1} r_{t_i}$

 $\mathbf{f}(\mathbf{x}_t) \in \Re^{n \times 1}$ the passive dynamics, $\boldsymbol{\theta}_t \in \Re^{p \times 1}$ the control vector and $\epsilon_t \in \Re^{p \times 1}$ Gaussian noise with variance Σ_{ϵ} . As immediate reward we consider $r_t = r(\mathbf{x}_t, \boldsymbol{\theta}_t, t) = q_t +$ $\boldsymbol{\theta}_t^T \mathbf{R} \boldsymbol{\theta}_t$ where $q_t = q(\mathbf{x}_t, t)$ is an arbitrary state-dependent reward function, and R the positive definite weight matrix of the quadratic control cost.

The \mathbf{PI}^2 algorithm is given in I. In step (2.3), the term $P(\boldsymbol{\tau}_{i,k})$ is the discrete probability at time t_i of each trajectory roll-out that is computed with the help of the cost $S(\boldsymbol{\tau}_{i,k})$ step (2.2). For every time step of the trajectory, a parameter update $\delta \theta_{t_i}$ is computed in step (3.1) based on a probability weighted average cost over the sampled trajectories. The parameter updates at every time step are finally averaged in step (4) and the parameter update $\theta^{(new)}$ = $\boldsymbol{\theta}^{(old)} + \delta \boldsymbol{\theta}$ takes place in step (5).

Evaluations

We have evaluated \mathbf{PI}^2 in three different tasks and experimental platforms. In the first evaluation in Figure 1 we are comparing \mathbf{PI}^2 against Policy Gradient Methods in the task of reaching a goal state by passing via a intermediate target



Fig. 1. Comparison of learning multi-DOF movements (2,10, and 50 DOFs) with planar robot arms passing through a via-point G. a,c,e) illustrate the learning curves for different RL algorithms, while b,d,f) illustrate the end-effector movement after learning for all algorithms. **PI**² performs an order of magnitude better than other methods.

G for a planar robot manipulator of 2,10 and 50 DOF. In the second evaluation in Figure 2, the Little Dog robot learns the task of how to jump over a large gap. In the third evaluation we use a Phantom robot (Figure 3) and we are applying \mathbf{PI}^2 not only for learning optimal state trajectories but also for learning the optimal gain schedule. The task for the phantom robot is to reach a goal state by accurately passing via an intermediate target, while maximizing compliance during control. The resulting behavior is illustrated in Figure 4.

REFERENCES

- E.A. Theodorou, J. Buchli, and S. Schaal. Reinforcement learning of motor skills in high dimensions: A path integral approach. In *International Conference Of Robotics and Automation (ICRA2010)*, 2010. To Appear.
- [2] Jiongmin Yong. Relations among ODEs, PDEs, FSDEs, BSDEs, and FBSDEs. In Proceedings of the 36th IEEE Conference on Decision and Control, 1997, volume 3, pages 2779–2784, Dec 1997.
- [3] P. Dayan and G. Hinton. Using EM for reinforcement learning. *Neural Computation*, 9, 1997.
- [4] J. Peters and S. Schaal. Learning to control in operational space. International Journal of Robotics Research, 27:197–212, 2008.
- [5] J. Koeber and J. Peters. Learning motor primitives in robotics. In D. Schuurmans, J. Benigio, and D. Koller, editors, *Advances in Neural Information Processing Systems 21 (NIPS 2008)*, Vancouver, BC, Dec. 8-11, 2009. MIT Press, Cambridge, MA.

Topics: Learning Algorithms; Preference:Oral



(a) Real & Simulated Robot Dog



(b) Learning curve for Dog Jump with $\mathbf{PI}^2 \pm 1std$





Fig. 3. Illustration of the three DOF kinematics of the Phantom robot.



Fig. 4. The plots in the first row illustrate the joint movements of the robot. In the second row the time varying gains schedule is illustrated. The dashed line is the initial trajectory and the blue line corresponds to the trajectory after learning. At the intermediate target states the gains increase so that the task of passing through the intermediate goal can be accomplished, while the gains are kept low otherwise.