
GPU based Path Integral Control with Learned Dynamics

Grady Williams

Autonomous Control and
Decision Systems Laboratory
Georgia Institute of Technology
gradyrw@gatech.edu

Eric Rombokas

Mechanical Engineering
University of Washington
eric@rombokas.com

Tom Daniel

Department of Biology
University of Washington
danielt@uw.edu

Abstract

We present an algorithm which combines recent advances in model based path integral control with machine learning approaches to learning forward dynamics models. We take advantage of the parallel computing power of a GPU to quickly take a massive number of samples from a learned probabilistic dynamics model, which we use to approximate the path integral form of the optimal control. The resulting algorithm runs in a receding-horizon fashion in realtime, and is subject to no restrictive assumptions about costs, constraints, or dynamics. A simple change to the path integral control formulation allows the algorithm to take model uncertainty into account during planning, and we demonstrate its performance on a quadrotor navigation task. In addition to this novel adaptation of path integral control, this is the first time that a receding-horizon implementation of iterative path integral control has been run on a real system.

1 Introduction

General motion planning is one of the core problems in robotics and there have been recent successes in this field through the use of model-based optimal control. Some recent examples are [1] and [2] for legged robot locomotion, and [3] for aggressive control of autonomous quadrotors. Despite these successes, a major drawback with this approach is the reliance on accurate forward or inverse dynamics models, which can be difficult or impossible to develop through physical analysis. A promising method of dealing with this problem is to learn a dynamics model through interaction with the environment instead of relying on an expert provided physics based model.

An important issue in this approach is that the learned dynamics only accurately describe a small region of the state space, where the training examples lie. Without taking any extra precautions, the controller assumes that the model accurately describes the dynamics for the entire space. This effect, termed model bias in [4], is potentially catastrophic and must be dealt with for consistent and safe performance.

Various approaches have been successful in dealing with model bias and have achieved good performance on simple tasks, the most well known of these is perhaps PILCO [4] which demonstrated remarkable learning performance on a cart-double-pole swing-up task while learning dynamics as a Gaussian Process. Another recent algorithm is due to Kuindersma et al. [5] who considers risk control for stochastic optimization and demonstrates good performance on stabilization tasks. In

contrast to these works, we consider control for tasks which require more general cost descriptions, such as obstacle navigation.

1.1 Contributions

We present an algorithm and computational framework for combining model based receding-horizon path integral control with machine learning algorithms that are capable of learning probabilistic forward dynamics models. Specifically we modify the recently developed receding-horizon formulation of Policy Improvement with Path Integrals (PI^2 -RH) [6] to take into account model uncertainty, and then use Locally Weighted Projection Regression (LWPR) [7] to learn a probabilistic dynamics model of a nano-quadrotor. The new algorithm performs unprecedented massive sampling of the probabilistic dynamics, but runs in real time due to its implementation on a graphics processing unit (GPU). We report the algorithm’s performance controlling the nano-quadrotor on an obstacle avoidance navigation task. Additionally this is the first implementation of receding-horizon PI^2 on a real machine.

2 Learning forward models with LWPR

Many machine learning algorithms are capable of acting as black box predictors for dynamical systems (see [8] for an overview), however only a subset are able to give probabilistic predictions at a cheap computational cost. LWPR is an algorithm based on a locally weighted incremental version of partial least squares, and under the classic probabilistic interpretation of weighted least squares it’s able to give uncertainty estimates along with a mean prediction. The algorithm learns a set of local linear models and equips each model with a center point in the state space and a distance metric. These values are used to compute predictions and uncertainties as a weighted average of all the models, let $y_i(x)$, $\sigma_i(x)$, and c_i be the mean, standard deviation, and center of the i_{th} local model and let D_i be the distance metric. If there are L local models then the LWPR mean $y(x)$ and variance $\sigma(x)^2$ predictions are given by:

$$y(x) = \sum_{i=1}^L w_i y_i, \quad \sigma(x)^2 = \sum_{j=1}^L w_j ((y - y_j)^2 + \sigma_k^2) \quad \text{where} \quad w_j = \frac{e^{-\frac{1}{2}(x-c_j)^T D_j (x-c_j)}}{\sum_{i=1}^L e^{-\frac{1}{2}(x-c_i)^T D_j (x-c_i)}} \quad (1)$$

The key advantage of LWPR is that the computational complexity of updating the local models and making predictions is linear in the dimension of the state space, as opposed to scaling with the number of training inputs like many other algorithms (e.g Gaussian Processes).

3 Control

In this section we briefly review Policy Improvement with Path Integrals (PI^2)[9] and its receding horizon implementation (PI^2 -RH)[6]. We then propose a modification to the algorithm which allows it to take into account model uncertainty.

3.1 Receding-Horizon path integral control

Consider a stochastic dynamical system of the form $dx = f(x, t)dt + G(x)u(t)dt + B(x)dw$, this system has non-linear passive dynamics, but is linear in the control and noise terms. We further assume that the noise is control dependent and that $BB^T = \lambda GR^{-1}G^T$, for some $\lambda > 0$. Now assume that we’re given a cost function of the form $J(x) = \phi(T) + \int_0^T (q(x, t) + u(t)^T R u(t))dt$, under these assumptions the stochastic Hamilton-Jacobi-Bellman equation for $V(x) = \min_u \mathbb{E}[J(x)]$ can be related to a path integral (See [9] for the derivation), and the optimal controls take the form:

$$u_{t_i} = \int P(\tau_i) u_L(\tau_i) d\tau_i \quad \text{and} \quad P(\tau_i) = \frac{\exp(-\frac{1}{\lambda} \tilde{S}(\tau_i))}{\int \exp(-\frac{1}{\lambda} \tilde{S}(\tau_i)) d\tau_i} \quad (2)$$

Where the cost-to-go function is: $\tilde{S}(\tau_i) = S(\tau_i) + \sum_{i=1}^N (u_t^T G_c^T H^{-1} G_c u_t dt + u_t^T G_c^T H^{-1} B dw_t^k)$.

And the portion of the cost-to-go depending on the passive dynamics is: $S(\tau_i) = \phi_{t_N} + \sum_{j=i}^{N-1} q(x, t)dt$ with $H = G_c R^{-1} G_c^T$ and $\mathbf{u}_L(\tau_i) = R^{-1} G_c^T (G_c R^{-1} G_c^T)^{-1} B_c(x)dw$. In these equations $\tau = \{x_0, x_1, x_2, \dots, x_{N-1}\}$ denotes a discretized trajectory, and τ_i is the remaining piece of the trajectory at the i^{th} timestep, N is the number of timesteps in a trajectory (ie. the number of discrete steps between 0 and T), and dt is the size of a single step.

In the PI^2 algorithm the optimal controls are computed by sampling from the stochastic dynamics in order to approximate (2). It's usually necessary to perform this approximation in an iterative fashion: start with an initial guess, sample possible trajectories from the dynamics (termed "rollouts") and compute an approximation to (2), the resulting approximation will be an improvement over the initial controls but may not be sufficient. If this is the case the process is repeated until the algorithm converges to a solution, and once the solution is computed the control is executed in an open loop fashion.

The receding-horizon implementation of PI^2 modifies this approach by constantly re-planning while simultaneously executing. Instead of iterating to convergence as in traditional PI^2 only a small, pre-defined, number of iterations are performed and then only a single timestep of the resulting trajectory is executed. The algorithm then receives state-feedback and warm-starts its optimization for the next timestep with the un-executed portion of the previous timestep's control plan. This modification transforms PI^2 into an implicit feedback controller, and allows it to perform difficult navigation tasks with a simulated quadrotor, which the original PI^2 algorithm has difficulty with [6].

3.2 Accounting for model uncertainty

In order to combat model bias we modify PI^2 in order to take into account the uncertainty of dynamics predictions from a learned probabilistic model. The role of the passive dynamics in evaluating rollouts is in the $S(\tau_i)$ portion of the cost-to-go function. In all previous model-based applications of PI^2 , $S(\tau_i)$ is evaluated based on a known passive dynamics $f(x, u)$ model. In order to evaluate $S(\tau_i)$ under the uncertain dynamics it makes sense to evaluate $\mathbb{E}[S(\tau_i)|F(x, u)]$, the expectation of the cost-to-go with respect to the uncertain dynamics $F(x, u)$, instead of computing a deterministic evaluation of $S(\tau_i)$ based on the mean prediction. This expectation is analytically intractable, but we can compute a crude Monte-Carlo approximation by taking multiple draws from the probability distribution for each rollout. A similar method for evaluating candidate control plans is considered by Bagnell and Schneider in [10], however we employ this strategy on a much larger scale. Under this approach a single rollout becomes a set of samples from a conditional probability distribution (termed sub-rollouts), and the cost-to-go of a rollout is evaluated as the average of the sub-rollouts.

4 System

We implemented the modified algorithm on a crazyflie [11] nano-quadrotor. The crazyflie weighs 19 grams and measures 9 centimeters from rotor to rotor. An onboard radio communicates with the ground station and onboard sensors and processors compute euler angles, euler angle rates, and use a PID controller to achieve commanded attitude rate targets. We used a Vicon Motion Capture system to calculate quadrotor position and velocity. The algorithms runs in realtime on the ground station computer which had an Nvidia GTX 780 Ti with 2880 CUDA cores. This computer was able to perform an iteration using only the mean prediction for 1000 50-step rollouts in about 6.5 milliseconds, and for 1000 rollouts with 32 sub-rollouts an iteration runs in about 18 milliseconds.

4.1 Models

We used a standard quadrotor model based on rigid body dynamics to run PI^2 -RH in order to compare it with the LWPR algorithm. Controls took the form of desired attitude rates (ϕ, θ, ψ) , and desired total thrust (\mathcal{F}) where the desired rates affect the actual rates according to the equation $\dot{r} = 25(r_{desired} - r_{actual})$.

We used flight examples taken from flights with PI^2 -RH and the analytic model to train 3 LWPR models with inputs $\phi, \theta, \psi, \mathcal{F}$ and with each output being one of the three cartesian accelerations $(\ddot{x}, \ddot{y}, \ddot{z})$. Then, we replaced the rigid-body formulas for those derivatives with the LWPR models

to obtain the learned probabilistic dynamics model. The learned models consisted of 15, 15, and 24 locally linear models for \dot{x} , \dot{y} , and \dot{z} respectively. In general the learned models were superior to the analytic model, the propagation errors over 1 second of both the analytic and LWPR model for a representative trial run of the quadrotor are given in Table 1.

Position/Velocity Component	Analytic Dynamics			LWPR Dynamics		
	x	y	z	x	y	z
Position Error (m)	.38	.63	.52	.27	.37	.20
Velocity Error (m/s)	1.01	1.71	1.25	.78	1.05	.36

Table 1: Average error of euler integration simulation after 1 second over a typical quadrotor run.

5 Experiments

In our experiments the quadrotor attempted to navigate between three points while avoiding three obstacles placed in the flight arena. The quadrotor performed 4 laps around these points during a trial, and performed 5 trials in total for each parameter setting. We tested the algorithm using the rigid body model and LWPR models with 1 (just the mean prediction), 4, 8, 16, and 32 sub-rollouts. For an instantaneous state cost we used: $q(x) = (x - W_x)^2 + (y - W_y)^2 + 10(z - W_z)^2 + \frac{1}{5}(\phi^2 + \theta^2 + \psi^2) + \frac{1}{10}(\dot{x}^2 + \dot{y}^2 + \dot{z}^2) + 100 \sum_{i=1}^3 \exp(-10(d_{x,i}^2 + d_{y,i}^2)) + 10C$. Here W_i denotes the i th component of the point that the quadrotor is moving towards, once the quadrotor comes within a quarter meter of the point the goal is changed to another point. The algorithm is not given advance knowledge that this switch will occur. $d_{x,i}$ is the x distance between the quadrotor and the i th obstacle, and similarly for $d_{y,i}$. The variable C is an indicator variable determining whether a crash has occurred or not.

The terminal cost was set to 0, and the cost term: $u_t^T G_c^T H^{-1} G_c u_t dt + u_t^T G_c^T H^{-1} B dw_t^k$ is considered 0 as well. This is because the term $u_t^T G_c^T H^{-1} G_c u_t dt$ is state independent, so the corresponding term in the denominator can be pushed out of the integral and cancels with the term in the numerator. The second term $u_t^T G_c^T H^{-1} B dw_t^k$ is not 0, but since control costs are very small it is close to zero and computational effort can be saved by not computing it, without suffering adverse affects.

We ran the algorithm with the maximal settings that allowed real time computation under the given parameters. This meant that the analytic model, and LWPR models with $M=1$ and $M=4$ were allowed to perform two iterations of optimization per timestep while all other settings could only perform 1 iteration in real time. Every trial used a time horizon of 1 second and 50 timesteps per second. The parameters for each setting are given in Table 2.

	RBD Model	M = 1	M = 4	M = 8	M = 16	M = 32
Rollouts	1000	1000	1000	1000	970	950
Optimizations per Timestep	2	2	2	1	1	1

Table 2: Number of rollouts and optimization iterations per timestep possible for the sub-rollout setting.

5.1 Results

The quadrotor was able to consistently complete the navigation task for all of the settings, except for $M = 1$ when only the mean prediction of LWPR was used. For this setting the quadrotor was overly-aggressive and unable to consistently stay in the flight volume while avoiding obstacles, out of 7 total trials the quadrotor successfully completed the task 4 times, flew into an obstacles once, and flew out of the flight arena twice.

We report the data from the 4 completed trials. Interestingly, when the quadrotor does complete the task for $M = 1$ the scores are very good. In fact the lowest average for total cost belongs to the successfully completed trajectories with $M = 1$. The issue is not poor performance, rather it is overly aggressive and risky performance which comes from not accounting for bias in the learned model.

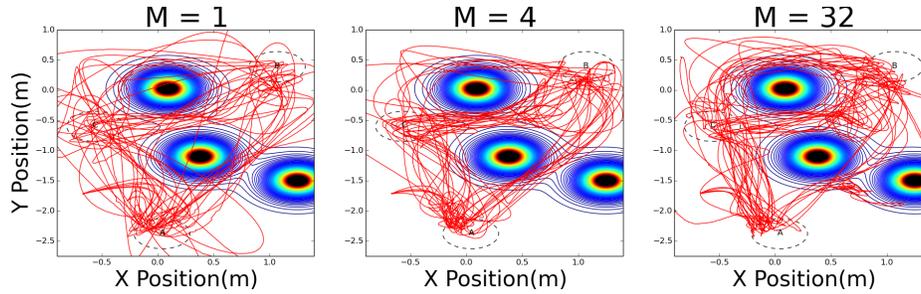


Figure 1: Trajectories of all the trials for $M = 1$ (left), $M = 4$ (center), and $M = 32$ (right). The trajectories are overlaid with a contour plot of the portion of the cost due to obstacles.

Sub-Rollouts	Avg. Time	Avg. Total Cost	Avg. Cost/Sec.	Avg. of 8 Closest Passes (m)
Analytic Model	37.05	6943.27	186.93	.23
$M = 1$	35.39	5733.55	162.01	.19
$M = 4$	34.91	5868.72	168.22	.30
$M = 8$	47.14	7273.80	154.27	.29
$M = 16$	49.40	7491.60	152.0	.32
$M = 32$	52.13	8523.40	162.99	.33

Table 3: Chart of performance metrics for all of the parameter settings. Note the the Cost/Sec is the cost over the optimization horizon.

Our method of taking into account the model uncertainty significantly improves the overall performance and consistency of the vehicle, just using a small number of sub-rollouts ($M = 4$) allows the vehicle to safely complete the task, and outperforms the analytic model in every metric. Using a greater number of sub-rollouts results in the vehicle choosing more deliberate plans, which have a lower cost over the 1 second planning horizon and stay further away from the obstacles.

The time to completion, total cost, cost over the time horizon, and the average distance of the 8 closest passes to obstacles is given in Table 3. We take the average of the 8 closest passes because there are 8 times that the quadrotor has to cross between (or go very far around) the obstacles during a trial.

5.2 Implicit variance minimization

The average variance of the LWPR predictions is given in Table 4. There's a clear trend of decreasing variance in the predictions as the number of sub-rollouts increases, this data in conjunction with the more deliberative actions taken by the parameter settings with higher numbers of sub-rollouts is evidence that the algorithm chooses safer, more certain actions when it has more knowledge about the effects of uncertainty on the system.

Number of Sub-Rollouts	\ddot{x}	\ddot{y}	\ddot{z}
$M = 1$.27	.39	1.39
$M = 4$.25	.38	1.32
$M = 8$.22	.34	1.18
$M = 16$.21	.33	1.17
$M = 32$.20	.32	1.14

Table 4: Average variance of acceleration predictions.

It's important to note that this variance minimization is happening without any variance term in the cost function. Minimization is happening implicitly because of the dangerous effects of executing uncertain trajectories in an obstacle filled environment.

6 Discussion

We have developed and tested an algorithm which successfully combines a modified version of PI^2 -RH, and LWPR. This is the first time that PI^2 -RH has been implemented on a real system, and also the first time that it has been combined with a learned dynamics model. Our results show that the PI^2 -RH algorithm successfully completes a difficult navigation task using an analytic model, and is outperformed by the new algorithm which uses a learned model and takes into account model bias. The naive implementation of the learned model, where just the mean prediction is used, results in un-reliable and dangerous performance. The new algorithm also minimizes the variance of its trajectories when obstacles are present, despite there being no variance term in the cost function. PI^2 -RH makes no restrictive assumptions about the form of the system dynamics or cost function, so given an accurate model it can theoretically perform a very wide class of tasks. Used in conjunction with model learning the potential exists for autonomous agents to learn complex tasks completely from scratch in a data efficient manner.

Acknowledgments

This research was supported by ONR grant No. N000141010952 for K Morgansen and T Daniel Komen Endowed Chair (TLD). We would also like to thank the Veterans Affairs (VA), Rehabilitation Research and Development Center of Excellence for Limb Loss Prevention and Prosthetic Engineering who lent us the use of their Motion Analysis Laboratory.

References

- [1] M. Kalakrishnan et al. “STOMP: Stochastic trajectory optimization for motion planning”. In: *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 4569–4574.
- [2] Y. Tassa, T. Erez, and E. Todorov. “Synthesis and stabilization of complex behaviors through online trajectory optimization”. In: *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 4906–4913.
- [3] D. Mellinger and V. Kumar. “Minimum snap trajectory generation and control for quadrotors”. In: *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2520–2525.
- [4] M. Deisenroth and C. Rasmussen. “PILCO: A Model-Based and Data-Efficient Approach to Policy Search”. In: *In Proceedings of the International Conference on Machine Learning*. 2011.
- [5] S. Kuindersma, R. Grupen, and A. Barto. “Variable Risk Control via Stochastic Optimization”. In: *International Journal of Robotics Research* (2013).
- [6] E. Rombokas, G. Williams, and T. Daniel. “Receding Horizon Trajectory Optimization on a GPU”. In: *Systems, Man, and Cybernetics, Part B* (Currently Under Review).
- [7] S. Vijayakumar and S. Schaal. “Locally Weighted Projection Regression: An $O(n)$ Algorithm for Incremental Real Time Learning in High Dimensional Space”. In: *in Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, pp. 1079–1086.
- [8] D. Nguyen-Tuong and J. Peters. “Model learning for robot control: a survey”. English. In: *Cognitive Processing* 12.4 (2011), pp. 319–340. ISSN: 1612-4782. DOI: 10.1007/s10339-011-0404-1. URL: <http://dx.doi.org/10.1007/s10339-011-0404-1>.
- [9] E. Theodorou, J. Buchli, and S. Schaal. “A generalized path integral control approach to reinforcement learning”. In: *The Journal of Machine Learning Research* 11 (2010), pp. 3137–3181.
- [10] J. Bagnell and J. Schneider. “Autonomous Helicopter Control using Reinforcement Learning Policy Search Methods”. In: *Proceedings of the International Conference on Robotics and Automation 2001*. IEEE, 2001.
- [11] *Bitcraze*. <http://www.bitcraze.se/>. [Online].