

Stochastic processes and feedback-linearisation for online identification and Bayesian adaptive control of fully-actuated mechanical systems

Jan-P. Calliess, Antonis Papachristodoulou and Stephen J. Roberts
Department of Engineering Science, Oxford University, UK

November 19, 2013

Abstract

This work proposes a new method for simultaneous probabilistic identification and control of an observable, fully-actuated mechanical system. Identification is achieved by conditioning stochastic process priors on observations of configurations and noisy estimates of configuration derivatives. In contrast to previous work that has used stochastic processes for identification, we leverage the structural knowledge afforded by Lagrangian mechanics and learn the drift and control input matrix functions of the control-affine system separately. We utilise feedback-linearisation to reduce, in expectation, the uncertain nonlinear control problem to one that is easy to regulate in a desired manner. Thereby, our method combines the flexibility of nonparametric Bayesian learning with epistemological guarantees on the expected closed-loop trajectory. We illustrate our method in the context of torque-actuated pendula where the dynamics are learned with a combination of normal and log-normal processes.

1 Introduction

Control may be regarded as decision making in a dynamic environment. Decisions have to be based on beliefs over the consequences of actions encoded by a model. Dealing with uncertain or changing dynamics is the realm of adaptive control. In its classical form, parametric approaches are considered (e.g. [20]) and, typically, uncertainties are modelled by Brownian motion (yielding stochastic adaptive control [6, 11]) or via set-based considerations (an approach followed by robust adaptive control [15]). In contrast, we adopt an epistemological take on probabilistic control and bring to bear Bayesian nonparametric learning methods whose introspective qualities [7] can aide in addressing the exploration-exploitation trade-offs relative to one's subjective beliefs in a principled man-

ner [1]. Based on these Bayesian learning methods, it is our ambition to develop adaptive controllers with probabilistic guarantees (interpreted in an epistemological sense) on control success.

In contrast to classical adaptive control where inference has to be restricted to finite-dimensional parameter space, the nonparametric approach affords the learning algorithms with greater flexibility to identify and control systems with very few model assumptions. This is possible because these methods grant the flexibility to perform Bayesian inference over rich, infinite-dimensional function spaces that could encode the dynamics. This property has led to a surge of interest in Bayesian nonparametrics; particularly benefiting their algorithmic advancement and application to a plethora of learning problems. Due to their favourable analytic properties, *normal* or *Gaussian* processes (GPs) [2, 16] have been the main choice of method in recent years. Among other domains, GPs have been applied to learning discrete-time dynamic systems in the context of model-predictive control [9, 10, 12, 17], learning the error of inverse kinematics models [13, 14], dual control [1] as well as reinforcement learning and dynamic programming [4, 5, 8, 18].

On the flip side, the extent of flexibility can lead to the temptation to use the approach in a black-box fashion, disregarding most structural knowledge of the underlying dynamics [8–10, 12, 18]. This can result in unnecessarily high-dimensional learning problems, slow convergence rates and often necessitates large training corpora, typically to be collected offline. In the extreme, the latter requirement can cause slow prediction and conditioning times. Moreover, they have been used in combination with computationally intensive planning methods such as dynamic programming [4, 5, 18] rendering real-time applicability difficult.

In contrast to all this work, we will incorporate structural a-priori knowledge of the dynamics afforded by Lagrangian mechanics (without sacrificing the flexibility afforded by the nonparametric nature). This requires, in some instances, partial departure from Gaussianity (e.g. if the sign of a function component of the dynamics is known) but improves the detail with which the system is identified and can reduce the dimensionality of the identification problem. Furthermore, our method will use the uncertainties of the models to decide upon training example incorporation and decision making.

Aside from learning, our method employs feedback-linearisation [19] in an outer-loop control law to reduce the complexity of the control problem. Thereby, in expectation, the problem is reduced to controlling a double-integrator via an inner-loop control law. If we combine the outer-loop controller with an inner-loop controller that has desirable guarantees (e.g. stability) for the double-integrator, these properties extend to the expected given non-linear closed-loop dynamics. The resulting approach enables rapid decision making and can be deployed online.

Our work is presented at the AMLSC Workshop at NIPS, 2013. During the review process, we were made aware of GP-MRAC [3]. The authors utilise a Gaussian process on joint state-control space to learn the error of an inversion controller in model-reference adaptive control. Under the assumption that the

GP could be stated as an SDE of time, they prove stability. In contrast to this work, our method is capable of identifying the drift and control input vector fields constituting the underlying control-affine system individually, yielding a more fine-grained identification result. While this benefit requires the introduction of probing signals to the control during online learning, each of the coupled learning problems has state space dimensionality only. Moreover, our method and stability results are not limited to Gaussian processes. If the control-input vector fields are identified with a log-normal process, our controller will automatically be cautious in scarcely explored regions.

2 Method

2.1 Model

Dynamics. Let $I \subset \mathbb{R}$ be a (usually continuous) set of times, \mathcal{Q} denote the configuration space, \mathcal{X} the state space and \mathcal{U} the control space. Via the principle of least action and the resulting Euler-Lagrange equation, Lagrangian mechanics leads to the conclusion that controllable mechanical systems are of second order and can be written in *control-affine* form:

$$\ddot{q} = a(q, \dot{q}) + b(q, \dot{q})u. \quad (1)$$

Here, $q \in \mathcal{Q}$ is a generalized coordinate of the configuration and $u \in \mathcal{U}$ is the control input. Functions a, b are called drift and input functions, respectively. In the pendulum control domain we consider below, q will encode joint angles and u is a torque \ddot{q} is proportional to.

Defining $x_1 := q, x_2 := \dot{q} \in \mathcal{Q}$, we can write the state as $x := [x_1, x_2]$. The dynamics can be restated as the system of equations

$$\dot{x}_1 = x_2 \quad (2)$$

$$\dot{x}_2 = a(x_1, x_2) + b(x_1, x_2)u \quad (3)$$

$$= a(x_1, x_2) + \sum_{j=1}^m u_j b_j(x_1, x_2) \quad (4)$$

where $m = \dim \mathcal{U}$ and $b_j(x_1, x_2)$ is the j th row of matrix $b(x_1, x_2) \in \mathbb{R}^{n \times m}$. In this work, we assume the system is *fully actuated*. That is, we assume that $b(q, \dot{q})$ always is full-rank: $\text{rank } b(q, \dot{q}) = \dim \mathcal{Q} =: n, \forall q$. That is, full-actuation enables us to instantaneously set the acceleration in all dimensions of \mathcal{Q} . However, we do not have immediate control over joint-angle velocities. Incorporating this kind of knowledge afforded by Lagrangian mechanics is beneficial both from a principled Bayesian vantage point and in order to decompose the dimensionality of the learning task.

Epistemic uncertainty and learning. Both dynamics functions a and b can be uncertain a priori. That is, a priori our uncertainty is modelled by

the assumption that $a \sim \Pi^a, b \sim \Pi^b$ where Π^a, Π^b are stochastic processes. The processes reflect our epistemic uncertainty about the true underlying (deterministic) dynamics functions a and b . If data becomes available over the course of the state evolution, we can update our beliefs over the dynamics in a Bayesian fashion. That is, at time $t \in I$ we assume $a \sim \Pi^a | \mathcal{D}_t, b \sim \Pi^b | \mathcal{D}_t$ where \mathcal{D}_t is the data recorded up to time t . The process of conditioning is often referred to as (Bayesian) *learning*.

Data collection. We assume our controller can be called at an ordered set of times $I_u \subset I$. At each time $t \in I_u$, the controller is able to observe the state $x_t = x(t)$ ¹ and to set the control input $u_t = u(t, x_t)$. The controller may choose to evoke learning at an ordered subset $I_\lambda \subset I_u$ of times. To this end, at each time $\tau \in I_\lambda$, the controller evokes a procedure explicated in Sec. 2.2 if it decides to incorporate an additional data point (t, x_t, u_t) into data set \mathcal{D}_t ($t > \tau$). The decision on whether to update the data will be based on the belief over the data point’s anticipated informativeness as approximated by its variance.²

For simplicity, we assume that learning can occur every Δ_λ seconds and the controller is called every $\Delta_u \leq \Delta_\lambda$ seconds. A continuous control takes place in the limit of infinitesimal Δ_u .

2.2 Learning procedure

To enable learning, we will require derivatives of the state (that is estimates of \ddot{q} and \dot{q}). If we do not have physical means to measure velocities and accelerations, obtaining numerical estimates becomes necessary based on observations of $q(t) = x_1(t)$. To estimate derivatives, we chose a second-order method. That is, our state derivative estimates are $\dot{y}(t_i + \Delta_o) := \frac{x(t_i + 2\Delta_o) - x(t_i)}{2\Delta_o}$ where Δ_o is a period length with which we can observe states. In this work, we assume $\Delta_o = \Delta_u$.

Assuming online learning, the data sets \mathcal{D}_t are found incrementally. Since it is hard to use the data to infer a and b simultaneously, we will have to actively decide which one we desire to learn about (and set the control accordingly – which we will then refer to as a *probing control*). To this end, we distinguish between the following learning components:

- Learning $a(x)$: Assume we are at time $t \in I_\lambda$ and that we decide to learn about a . This decision is made, whenever our uncertainty about $a_t := a(x_t)$, encoded by $\text{var}[a(x_t)]$, is above a certain threshold θ_{var}^a . When learning is initiated, we keep the control constant for two more time steps $t + \Delta_u, t + 2\Delta_u$ to obtain a good derivative estimate as described above. To remove additional uncertainty due to ignorance about b , we set probing control $u_t = u_{t+\Delta_u} = u_{t+2\Delta_u} = 0$ yielding dynamics $\dot{x} = a(x)$ during time interval $[t, t + 2\Delta_u)$. On the basis of a derivative estimate \dot{y}_t , we

¹In fact, we can only observe q and have to obtain noisy observations of \dot{q} as we will describe below.

²Variance is known to approximate entropic measures of uncertainty (cf. [1]) and often easier to compute than entropy.

can determine a noisy estimate $\tilde{a}_{t+\Delta_u}$ of unknown function value $a_{t+\Delta_u}$ at time t as per

$$\tilde{a}_{t+\Delta_u} = \dot{y}(t + \Delta_u).$$

So, $(t + \Delta_u, \tilde{a}_{t+\Delta_u}, 0)$ is added to the data after time $t + 2\Delta_u$.

- Learning $b_j(x)$: At time $t \in I_\lambda$, we choose to learn about function b_j whenever our uncertainty about $a(x(t_i))$ is sufficiently small (i.e. $\text{var}[a](x_i) \leq \theta_a$) and our uncertainty about b_j is sufficiently large ($\text{var}[b]_j(x_i) > \theta_b$). When learning is initiated, we keep the control constant for two more time steps $t + \Delta_u$, $t + 2\Delta_u$ to obtain a good derivative estimate as described above.

Let $e_j \in \mathbb{R}^m$ be the j th unit vector. To learn about $b_j(x)$ at state x we apply a control action $u = u_j e_j$ where $u_j \in \mathbb{R} \setminus \{0\}$. Inspecting Eq. 4 we can then see that $b_j(x) = \frac{\dot{x}_2 - a(x)}{u_j}$. Since $a(x)$ will generally be a random variable, so is $b_j(x)$ having mean $\langle b(x) \rangle = \frac{\dot{x}_2 - \langle a(x) \rangle}{u_j}$ and variance $\text{var}[b(x)] = \frac{1}{u_j^2} \text{var}[a(x)]$. We obtain a noisy estimate \dot{y} of its derivative analogously to above. Modelling \dot{x}_2 as a random variable with mean \dot{y}_2 , $b_j(x)$ becomes a random variable with mean

$$\langle b_j(x) \rangle = \frac{\dot{y}_2 - \langle a(x) \rangle}{u_j} \quad (5)$$

and variance

$$\text{var}[b_j(x)] = \frac{\text{var}[\dot{x}_2] + \text{var}[a(x)]}{u_j^2} \leq \frac{\text{var}[\dot{x}_2] + \theta_a}{u_j^2}. \quad (6)$$

Therefore, after time $t+2\Delta_u$, we add training point $(x_{t+\Delta_u}, \langle b_j(x_{t+\Delta_u}) \rangle, u_t)$ to the data set. The additional variance (as per Eq. 6) is captured by setting observational noise levels for Π^b accordingly.

2.3 Control law

Unless the control actions are chosen to aid system identification (as described above), we will want to base our control actions on our probabilistic belief model over the dynamics. Given such an uncertain model, it remains to define an (outer-loop) control policy u with desirable properties. In this work, we propose to define a control law that, when not learning, uses the probabilistic model to guarantee arbitrary desired behaviour in expectation. Such behaviour can include, but is not limited to global asymptotic convergence to a goal state.

Let $a_t := a(x(t))$, $b_t := b(x(t))$ and $q_t := q(t)$. Acceleration $\ddot{q}_t = a_t + b_t u(t)$ is a random variable with mean $\langle \ddot{q}_t \rangle = \langle a_t | \mathcal{D}_t \rangle + \langle b_t | \mathcal{D}_t \rangle u$.

Hence, when applying inversion control law

$$u(t, x; u') := \langle b(x) | \mathcal{D}_t \rangle^{-1} (-\langle a(x) | \mathcal{D}_t \rangle + u') \quad (7)$$

we get an unexpected closed-loop dynamics of

$$\langle \dot{q}_t | \mathcal{D}_t \rangle = \langle x_2(t) | \mathcal{D}_t \rangle = \langle \dot{x}_1(t) | \mathcal{D}_t \rangle = \dot{y}_1 \quad (8)$$

$$\langle \ddot{q}_t | \mathcal{D}_t \rangle = \langle \dot{x}_2(t) | \mathcal{D}_t \rangle = \langle a_t | \mathcal{D}_t \rangle + \langle b_t | \mathcal{D}_t \rangle \langle b_t | \mathcal{D}_t \rangle^{-1} (-\langle a_t | \mathcal{D}_t \rangle + u') = u'. \quad (9)$$

Consequently, our control law guarantees feedback-linearisation inexpectation (and of the dynamics of the mean trajectory). That is, by choosing u' to impose desired behaviour for the double integrator problem $\ddot{q} = u'$ (which is easy), we can re-shape the dynamics such that the closed-loop dynamics exhibits that behaviour inexpectation in the actual system $\ddot{q} = a(q, \dot{q}) + b(q, \dot{q})u$.

For instance, a simple method of guaranteeing global asymptotic convergence of the state towards a goal state $\xi = [\xi_1, \xi_2]$ would be to set the inner-most control law to the proportional feedback law

$$u'(t, x; w) := w_1(\xi_1 - x_1) + w_2(\xi_2 - x_2) \quad (10)$$

where $w_1, w_2 > 0$.

Theorem 2.1. *Assume we are not performing probing actions anymore. That is, we are at time t_0 such that $t_0 > \sup I_\lambda$. Let $u'(t, x)$ be any control that ensures the double-integrator dynamics of the form*

$$\dot{z}_1 = z_2, \dot{z}_2 = u'(t, z)$$

to have ξ as a globally asymptotically stable equilibrium point. Then, our control law as per Eq. 7, with inner control law $u'(x, t)$, ensures ξ is a globally asymptotically stable equilibrium of the expected dynamics. In particular, $\lim_{t \rightarrow \infty} \|\langle q_t - \xi_1 | \mathcal{D}_{t_0} \rangle\|^2 = 0 \wedge \lim_{t \rightarrow \infty} \|\langle \dot{q}_t - \xi_2 | \mathcal{D}_{t_0} \rangle\|^2 = 0$.

Proof. (Sketch) Let ∇_t denote the differential operator with respect to time. Leveraging the linearity of the differential operator, we can exchange it with the expectation operator. Thereby, we conclude from Eq. 8 and Eq. 9 that $\nabla_t \langle x_1(t) | \mathcal{D}_{t_0} \rangle = \langle x_2(t) | \mathcal{D}_{t_0} \rangle, \nabla_t \langle x_2(t) | \mathcal{D}_{t_0} \rangle = u'$. Defining $z_i := \langle x_i(t) | \mathcal{D}_{t_0} \rangle$ yields the quadratic regulator problem: $\nabla_t z_1 = z_2, \nabla_t z_2 = u'$. By assumption, we know that u' ensures that ξ is a globally asymptotic equilibrium point of this dynamic system. Hence, in particular, $\lim_{t \rightarrow \infty} \|z_1(t) - \xi_1\|^2 = 0 \wedge \lim_{t \rightarrow \infty} \|z_2(t) - \xi_2\|^2 = 0$. Resubstituting the definitions of $\langle x_i(t) | \mathcal{D}_{t_0} \rangle$ for z_i and subsequently, of $q = x_1, \dot{q} = x_2$, yields the desired statement. \square

3 Experiments – Learning to control a torque-controlled damped pendulum with a combination of normal and log-normal processes

We explored our method's properties in simulations of a rigid pendulum with (a priori unknown) drift $a(x) := -\frac{g}{l} \sin(x_1) - \frac{r(x_1)}{ml^2} x_2$ and constant input function

$b(x) = \frac{1}{ml^2}$. Here, $x_1 = q, x_2 = \dot{q} \in \mathbb{R}$ are joint angle position and velocity, r denotes a friction coefficient, g is acceleration due to gravity l is the length and m the mass of the pendulum. The control input $u \in \mathbb{R}$ applies a torque to the joint that corresponds to joint-angle acceleration. The pendulum could be controlled by application of a torque u to its pivotal point. $q = 0$ encode the pendulum pointing downward and $q = \pi$ denoted the position in which the pendulum is upward. Given an initial configuration $x_0 = [q_0, \dot{q}_0]$ we desired to steer the state to a terminal configuration $\xi = [q_f, 0]$.

For learning, we assumed that $a \sim \mathcal{GP}(0, K_a)$ and $b \sim \log \mathcal{GP}(0, K_b)$ had been drawn from a log-normal process.³ The latter assumption encodes a priori knowledge that control input function b can only assume positive values (but, to demonstrate the idea of cascading processes, we had discarded the information that b was a constant). During learning, the latter process was based on a standard normal process conditioned on log-observations of \hat{b} . To compute the control as per Eq. 7, we need to convert the posterior mean over $\log b$ into the expected value over b . The required relationship is known to be as follows:

$$\langle b(x)|\mathcal{D}_t \rangle = \exp\left(\langle \log b(x)|\mathcal{D}_t \rangle + \frac{1}{2}\text{var}[\log b(x)|\mathcal{D}_t]\right). \quad (11)$$

If required the posterior variance can be obtained as

$$\text{var}[b(x)|\mathcal{D}_t] = \exp\left(2\langle \log b(x)|\mathcal{D}_t \rangle + \text{var}[\log b(x)|\mathcal{D}_t]\right) \exp\left(\text{var}[\log b(x)|\mathcal{D}_t] - 1\right).$$

Note, the posterior mean over b increases with the variance of our normal process in log-space, and, the control law as per Eq. 7 is inversely proportional to the magnitude of this mean. Hence, the resulting controller is *cautious*, in the sense that control output magnitude is damped exponentially in regions of high uncertainty (variance).

To simulate a discrete 0th order sample-and-hold controller in a continuous environment, we simulated the dynamics between two consecutive controller calls (occurring every Δ_u seconds) employing standard ODE-solving packages (i.e. Matlab’s ode45 routine).

We illustrated the behaviour of our controllers in a sequence of four experiments. The parameter settings are provided in Tab. 1. Recorded control energies and errors (in comparison to continuous proportional controllers) are provided in Tab. 2.

Our Bayesian controller maintains an epistemic beliefs over the dynamics. These beliefs govern our control decisions (including those when to learn). Furthermore, to keep prediction times low, beliefs are only updated when the current variance indicated a sufficient of uncertainty. Therefore, one would expect to observe three properties of our controller:

(i) When the priors are chosen sensibly (could be indicated by the dynamic functions’ likelihood under the probabilistic models), we expect good control performance.

³For details on normal processes see [16].

(ii) Prior training improves control performance and, reduces learning, but is not necessary to reach the goal. Both properties can be observed in Exp.1 and Exp. 3.

(iii) When the controller is ignorant of the inaccuracy of its beliefs over the dynamics (i.e. the actual dynamics are unlikely but the variances are low), control may fail since the false beliefs are not updated. An example of this is provided in exp. 2.

(iv) We can overcome such problems practically, by employing the standard technique (see [16]) of choose the prior hyper-parameters that maximise marginal likelihood. In Exp. 3, this approach was successfully applied to the control problem of Exp. 2.

Parameter(s) :	(l,r,m)	Δ_u	Δ_l	$(\theta_{var}^a, \theta_{var}^{\log b})$	x_0	ξ	(w_1, w_2)	t_f
Exp. 1	(1,1,0.5)	.01	.5	(.001, .005)	(0,-2)	$(\pi, 0)$	(1,1)	20
Exp. 2	(1,0.5,4)	.01	1	(.001, .005)	(0,-2)	$(\pi, 0)$	(2,2)	15
Exp. 3	(1,0.5,4)	.01	1	(.001, .005)	(0,-2)	$(\pi, 0)$	(2,2)	20

Table 1: Parameter settings.

Controller :	$\int_I u_{adapt}^2(t)dt$				$\int_I (x(t) - \xi)^2 dt$				$(\mathcal{D}_{t_f}^a , \mathcal{D}_{t_f}^b)$	
	P1	P100	SP1	SP2	P1	P100	SP1	SP2	SP1	SP2
Exp. 1	134	644	139	57	137	10	59	25	(18, 20)	(23, 53)
Exp. 2	552	11942	14759	17029	139	10	82	72	(2,1)	(2,1)
Exp. 3	730	11942	2370	1559	184	10	83	17	(11,2)	(11,2)

Table 2: Cumulative control energies, squared errors and data sizes (rounded to integer values). Pk : P-controller with feedback gain k . P1 failed to reach the goal state in all experiments. High-gain controller P100 succeeded in reaching the goal in all experiments but required a lot of energy. SP1: stochastic process - based controller with empty data set to start with. SP2: reset SP1 with training data collected from the first run.

Experiment 1. We started with a zero-mean normal process prior over $a(\cdot)$ endowed with a rational quadratic kernel with automated relevance detection (RQ-ARD) [16]. The kernel hyper-parameters were fixed. Observational noise variance was set to 0.01. The log-normal process over $b(\cdot)$ was implemented by placing a normal over $\log b(\cdot)$ with zero mean and RQ-ARD kernel with fixed hyper-parameters and observational noise level 0.1. Note, the latter was set higher to reflect the uncertainty due to Π^a . In the future, we will consider incorporating heteroscedastic observational noise based on $\text{var}[a]$ and the sampling rate. Also, one could incorporate knowledge about periodicity in the kernel.

Results are depicted in Fig. 1 and 2. We see that the system was accurately identified by the stochastic processes. When restarting the control task with

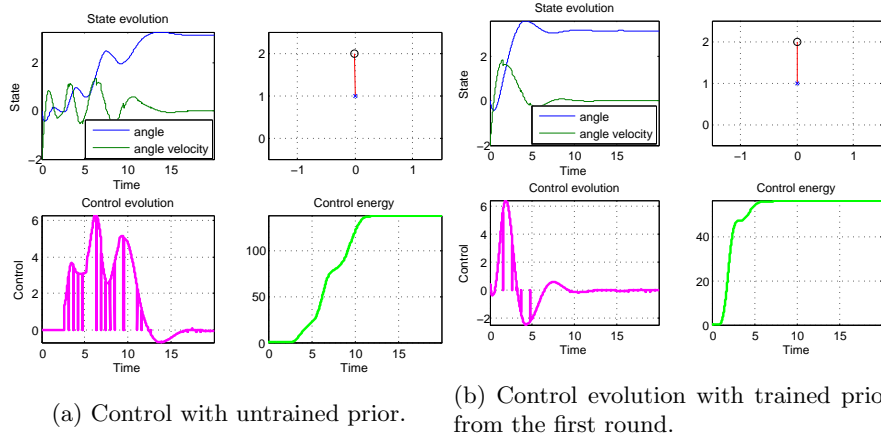


Figure 1: Experiment 1. Comparison of runs with untrained and pre-trained processes. The top-right image shows the final position of the pendulum having successfully reached the target angle $\xi_1 = \pi$. The dips in the control signal represent probing control actions arising during online learning.

stochastic processes pre-trained from the first round, the task was solved with less learning, more swiftly and with less control energy.

Experiment 2. We investigated the impact of inappropriate magnitudes of confidence in a wrong model. We endowed the controller’s priors with zero mean functions and SE-ARD kernels [16]. Length scales of kernel K_a were set to 20 and the output scale to 0.5. In addition to the low output-scale, we set observational noise variance to a low value of 0.0001 suggesting (ill-founded) high confidence in the prior. The length scale of kernel K_b was set to 50 with low output scales and observational noise variance of 0.5 and 0.001, respectively.

The results, depicted in Fig. 3. As to be expected, the controller fails to realise the inadequateness its beliefs. This results in a failure to update its beliefs and consequently, in a failure to converge to the target state.

Of course, this could be overcome with an actor-critic approach. Such solutions will be investigated in the context of future work.

Experiment 3.

Exp. 2 was repeated. This time, however, the kernel hyper-parameters were found by maximizing the marginal likelihood of the data (using Matlab’s *fmincon*). The automated identification of hyper-parameters is beneficial in practical scenarios where definition of a good prior for the underlying dynamics may be hard to conceive.

The optimiser succeeded in finding sensible parameters that allowed good control performance. As before, the method benefited from prior training yielding faster convergence and lower control effort. Both untrained and pre-trained methods outperformed the P -controllers either in terms of control energy or convergence. Finally, the SP controllers with hyper-parameter optimisation

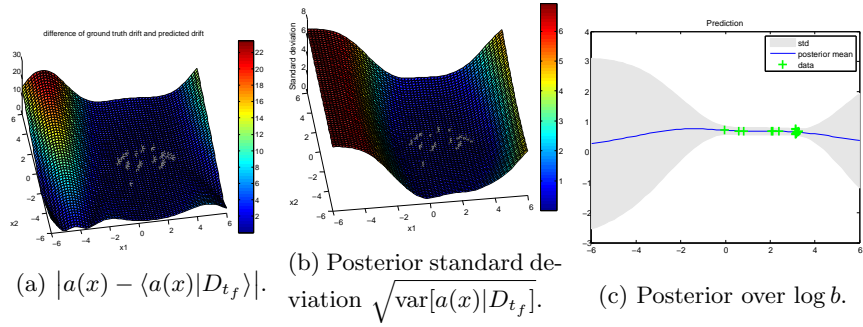


Figure 2: Experiment 1. Posterior models of SP1. Stars indicate training examples. The stochastic process has learned the dynamics functions in explored state space with sufficient accuracy.

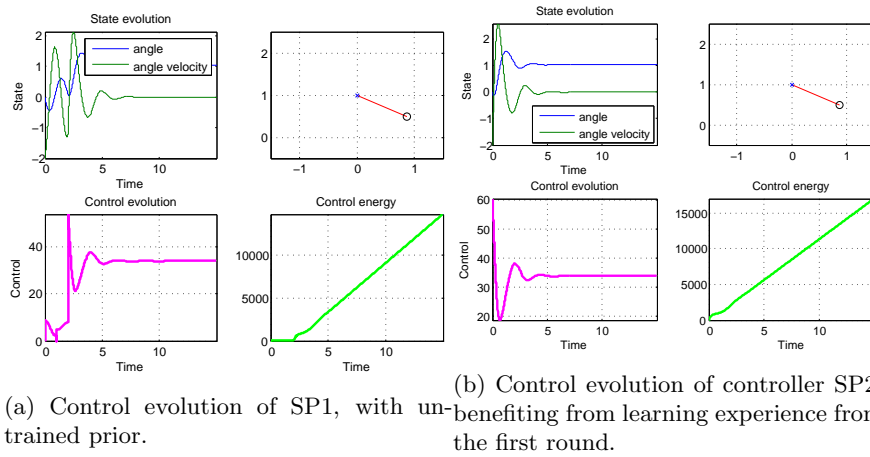


Figure 3: Experiment 2. Comparison of runs with untrained and pre-trained processes. Neither run succeeds in arriving at the target state due to being overly confident.

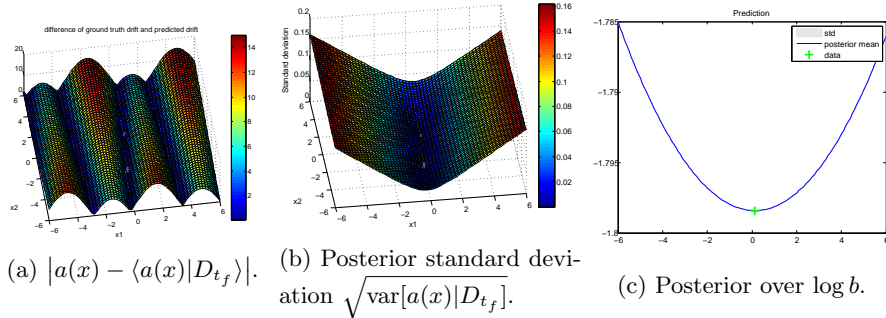


Figure 4: Experiment 2. Posterior models of SP1. Stars indicate training examples. Note, the low posterior variance suggests misleading confidence in an inaccurate model.

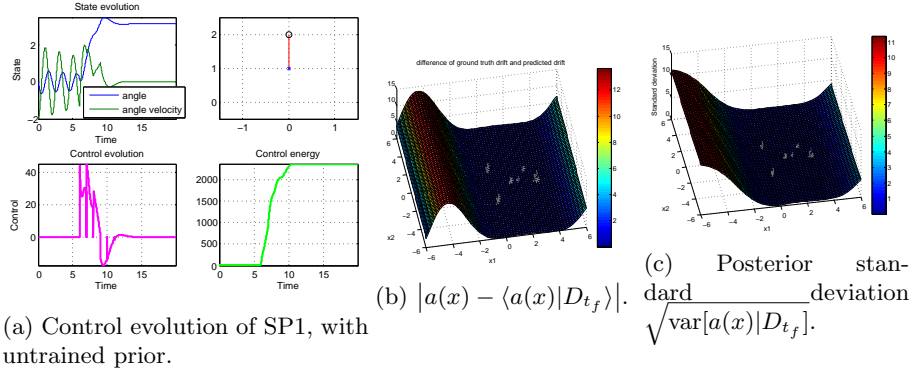


Figure 5: Experiment 3. Posterior models of SP1. Stars indicate training examples. The optimisation process succeeded in finding a sufficiently appropriate model.

outperformed the SP controllers with fixed hyper-parameters set in exp. 2 (c.f. Tab. 2).

4 Conclusions

We have applied Bayesian nonparametric methods to learn online the drift and control input functions of a fully-actuated control-affine second-order dynamical system. Paired with the idea of feedback-linearisation we devised a control law that switches between probing actions for learning and control signals that drive the expected trajectory towards a given setpoint. Our simulations have illustrated our controller's behaviour in the context of a pendulum regulator problem and that it can successfully solve the identification and control problems. They

have also served as an illustration of the inherent pitfalls of Bayesian control – that is, guarantees are stated relative to epistemological beliefs (encoded by a posterior) over the dynamical system in question. Therefore, the controller’s performance may be undermined by ignorance over the potential falsity of prior beliefs (cf. exp. 3). However, as illustrated in Exp. 3, even the most simple model selection methods can alleviate the burden of having to conceive a good fixed prior.

In *future work*, we will explore how to employ the actor-critique approach to uncover over-confidence of our models and to initiate learning. At present, our control law achieves desired performance of the expected trajectory. We will investigate how to extend the guarantees to achieve performance guarantees in expectation and within probability bounds. Other theoretical questions under investigation are analysis of the trade-offs between the impact of probing actions (to learn), the desire to keep prediction time low, information gain and control refresh cycle length Δ_u . Finally, we will assess our methods’ performance in higher-dimensional systems.

References

- [1] Tansu Alpcan. Dual Control with Active Learning using Gaussian Process Regression. *Arxiv preprint arXiv:1105.2211*, pages 1–29, 2011.
- [2] H. Bauer. *Wahrscheinlichkeitstheorie*. deGruyter, 2001.
- [3] Girish Chowdhary, HA Kingravi, JP How, and PA Vela. Bayesian Nonparametric Adaptive Control using Gaussian Processes. Technical report, MIT, 2013.
- [4] MP Deisenroth, J. Peters, and C. E. Rasmussen. Approximate dynamic programming with gaussian processes. *ACC*, June 2008.
- [5] M.P. Deisenroth, C. E. Rasmussen, and J. Peters. Gaussian process dynamic programming. *Neurocomputing*, 2009.
- [6] T.E. Duncan and B.Pasik-Duncan. Adaptive control of a scalar linear stochastic system with a fractional brownian motion. In *FAC World Congress*, 2008.
- [7] H. Grimmer, R. Paul, R. Triebel, and I. Posner. Knowing when we dont know: Introspective classification for mission-critical decision making. In *ICRA*, 2013.
- [8] J. Ko, D. Klein, D. Fox, and D. Haehnel. Gaussian Processes and Reinforcement Learning for Identification and Control of an Autonomous Blimp. In *ICRA*, 2007.
- [9] J. Kocijan and R. Murray-Smith. Nonlinear Predictive Control with a Gaussian. *Lecture Notes in Computer Science 3355*, Springer, pages 185–200, 2005.
- [10] J. Kocijan, R. Murray-Smith, C.E. Rasmussen, and B. Likar. Predictive control with Gaussian process models. In *The IEEE Region 8 EUROCON 2003. Computer as a Tool.*, volume 1, pages 352–356. Ieee, 2003.
- [11] P. R. Kumar. A survey of some results in stochastic adaptive control. *Siam J. Control and Optimization*, 23, 1985.
- [12] Roderick Murray-smith, Carl Edward Rasmussen, and Agathe Girard. Gaussian Process Model Based Predictive Control. In *IEEE Eurocon 2003: The International Conference on Computer as a Tool*, 2003.

- [13] D. Nguyen-Tuong and J. Peters. Using model knowledge for learning inverse dynamics. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2010.
- [14] D. Nguyen-Tuong, J. Peters, M. Seeger, and B. Schölkopf. Learning inverse dynamics: a comparison. In *Europ. Symp. on Artif. Neural Netw.*, 2008.
- [15] Ioannou P. and J. Sun. *Robust Adaptive Control*. Prentice Hall, 1995.
- [16] C.E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [17] Alex Rogers, Sasan Maleki, Siddhartha Ghosh, and N.R. Jennings. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. In *2nd Int. Workshop on Agent Technology for Energy Systems (ATES 2011)*, 2011.
- [18] A. Rottmann and W. Burgard. Adaptive Autonomous Control using Online Value Iteration with Gaussian Processes. In *ICRA*, 2009.
- [19] M. W. Spong. Partial feedback linearization of underactuated mechanical systems. In *Proc. IEEE Int. Conf. on Intel. Robots and Sys. (IROS)*, 1994.
- [20] K.Y. Volyanskyy, M.M. Haddad, and A.J. Calise. A new neuroadaptive control architecture for nonlinear uncertain dynamical systems: Beyond sigma- and e-modifications. In *CDC*, 2008.