

# Nonparametric Dynamics Estimation for Time Periodic Systems

Edgar D. Klenske<sup>1</sup>, Melanie N. Zeilinger<sup>1,2</sup>, Bernhard Schölkopf<sup>1</sup>, Philipp Hennig<sup>1</sup>

**Abstract**—Screws and gears are a source of periodically recurring nonlinear effects in mechanical dynamical systems. Unless the state sampling frequency is much higher than the periodic effect, model-free controllers cannot always compensate these effects, and good physical models for such periodic dynamics are challenging to construct. We investigate nonparametric system identification with an explicit focus on periodically recurring nonlinear effects. Within a Gaussian process regression framework, we design a locally periodic covariance function to shape the hypothesis space, which allows for a structured extrapolation that is not possible with more widely used covariance functions. These predictions are then used in model predictive control to construct a control signal correcting for the predicted external effect. We show that this approach is beneficial for state sampling times that are smaller than, but comparable to, the period length of the external effect. In experiments on a physical system, an electrically actuated telescope mount, this approach achieves a reduction of about 20 % in root mean square error.

## I. INTRODUCTION

Many controlled systems suffer from additive effects that recur periodically over time, for example arising from mechanical imprecision in gears and drives or from periodic movements in biological systems. Since these effects are often small in comparison to the required control precision, they are usually considered as errors and are neglected in the controller design. For high precision control systems, such errors can be the dominant source of problems. Correcting errors only after they are measurable necessarily leads to a delay in the error correction. If these errors can be anticipated, the control performance can be improved. While errors arising stochastically or from unknown effects cannot be preempted in this way, periodic effects are amenable for prediction: Since their future behaviour resembles their past, extrapolation (prediction) is easier and more structured. Based on this idea, we present a framework for identification and control of unmodelled periodic effects. Our framework continually performs identification at runtime, and is thus applicable to stochastic time varying systems.

The correction of periodic errors has repeatedly been studied. [1] investigated this issue in the “Very Large Telescope”, using an internal parametric model for the known error sources, and a Kalman filter as an observer for the model parameters. High precision tracking of spacecrafts on periodic trajectories was improved by [2] with the help

of predictive filtering using an extended Kalman filter. Not only mechanical systems suffer from periodic errors, but also systems that compensate the motion of biologic systems, like a beating heart. [3] also employ an extended Kalman filter for state estimation in that setting, but allow the nonlinear model to change over time, because some parameters of the model vary. [4] used model predictive control in a learning controller setting to adapt to model mismatch.

In contrast to previous methods, the approach presented here does not rely on a pre-specific finite-dimensional model class. Instead, we propose a nonparametric system identification framework based on Gaussian process (GP) regression. It is closely related to least-squares regression used in system identification, but based on a probabilistic interpretation, which can be used to guide exploration during identification [5]. There is recent work on using GPs for state filtering [6] and on modelling and control of nonlinear systems [7]. While in [7] GPs are used in a reinforcement learning setting, we are incorporating the GP prediction in an optimal control framework. The idea of using the learned model in predictive control is conceptually similar to [4], with the key difference that we use a GP to predict time varying effects. The textbook by [8] provides a general introduction to Gaussian processes.

A Gaussian process model is parametrised by two objects – mean and covariance function. When used in system identification, in particular the choice of covariance function has a strong effect on performance, and requires consideration of the particular dynamics to be identified. The notion of a “universal kernel” [9], [10] is misleading, because it only applies in the limit of large datasets, and can be subject to extremely slow (logarithmic) convergence [11], [12]. In any case, the notion does not apply to extrapolation, which is the relevant setting for control (put simply, predicting the future is harder than interpolating between past events). Even nonparametric regression methods therefore require careful modelling decisions.

In this work, we focus on a particular class of models involving periodicity (§ II), and suggest a family of regression models for the identification of such systems (§ III-A). In the case where the system modelled up to the periodic effect is linear, these predictions can then be used in a modified Kalman filter to improve state estimation (§ III-C), and in a linear model predictive control setting to achieve optimal closed-loop performance (§ III-D). We study qualitative properties of this framework on a toy problem in the form of a double integrator with periodic errors acting on both states (§ IV-A). As the development of this method was driven

<sup>1</sup> Max Planck Institute for Intelligent Systems, Spemannstraße 38, 72076 Tübingen, Germany {eklenske, bs, phennig}@tue.mpg.de

<sup>2</sup> Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, CA 94720, USA melanie.zeilinger@eecs.berkeley.edu

by a real problem in astronomy, the method is evaluated on this problem both in simulation (§ IV-B.1) and hardware experiments (§ IV-B.2).

## II. PROBLEM STATEMENT

We consider linear systems with a nonlinear and time varying additive function  $g(t)$ .

$$\dot{x}(t) = Ax(t) + Bu(t) + g(t) \quad (1a)$$

$$y(t) = x(t) + v(t) \quad (1b)$$

with state  $x \in \mathbb{R}^E$ , input  $u \in \mathbb{R}^F$ , output  $y \in \mathbb{R}^E$  and system matrices  $A$  and  $B$  of appropriate size. For simplicity, we assume full state measurement, where  $v$  denotes an independently and identically Gaussian distributed measurement noise. The approach can be extended to the case of partial state measurement (i.e.  $y = Cx$  with rank-deficient matrix  $C$ ) by using an observer to estimate the remaining states. The function  $g : \mathbb{R}_+ \rightarrow \mathbb{R}^D$  captures nonlinear time dependent effects. In particular, we are concerned with systems exhibiting some form of periodic behaviour in  $g(t)$ . For a strictly periodic function, there exists a constant period  $\omega$ , such that  $g(t + n\omega) = g(t)$  for  $n \in \mathbb{N}$ . However, not all error sources in real systems are perfectly periodic in this sense, they show various forms of phase-shift and desynchronisation. To address this issue, we generalise our framework to consider “locally-periodic” functions. Intuitively speaking, these are functions for which  $g(t) \approx g(t + n\omega)$  for  $n\omega \ll \ell$  and  $g(t) \not\approx g(t + n\omega)$  for  $n\omega \gg \ell$ , where  $\ell$  is some measure of locality.

A more technically precise formulation of this notion is: We consider a hypothesis class of functions  $g(t)$  to which we assign a prior probability measure such that the covariance between two function values at  $t$  and  $t'$  is

$$\text{cov}(g(t), g(t')) = k_{\text{periodic}}(|t - t'|) \cdot k_{\text{decay}}(|t - t'|), \quad (2)$$

where  $k_{\text{periodic}}$  is a periodic function with  $k_{\text{periodic}}(n\omega) = 1$  and  $k_{\text{decay}}$  is a monotonically decreasing, positive function. (See also Figure 1 for an illustration.)

We consider the case where the linear parts of the model (matrices  $A$  and  $B$ ) are known. Based on this, the discrete-time model is

$$x_{k+1} = A_d x_k + B_d u_k + a_k, \quad (3)$$

where  $A_d$  and  $B_d$  are obtained from a zero-order-hold discretisation

$$A_d = e^{A\Delta t} \quad G_d = \int_0^{\Delta t} e^{A\tau} d\tau \quad B_d = G_d B \quad (4)$$

with sampling time  $\Delta t$ .  $a_k$  is the convolution of the periodic component  $g(t)$  with the state transition matrix  $e^{At}$ . The regression problem is to infer a model for  $g(t)$  from observations of the change in state measurement,  $\Delta y = y(t) - y(t - \Delta t)$ . From these we construct observations for the regressor by local linearisation and correcting with the linear dynamics on the estimated state  $\hat{x}_k$ :

$$g(t) \approx G_d^{-1} (\Delta y - (A_d - I)\hat{x}_k - B_d u_k). \quad (5)$$

Before we proceed, it should be intuitively clear that the performance gain one can expect from the use of a periodic model for nonlinear effects depends on the sampling rate of the control system: If states are measured, and the control signal adjusted, at a frequency much higher than the periodicity of the modelling error, a locally linear model is sufficient and a periodic model offers almost negligible benefit. But if the error’s frequency is comparable to that of the state measurement, explicitly predicting the periodic dynamics in the “dark” phases between measurements becomes promising. See Section IV-B.1 for a more detailed discussion.

## III. METHODS

### A. Gaussian process regression

A Gaussian process  $\mathcal{GP}(g; \mu, k)$  is an (infinite-dimensional) probability distribution over the space of real-valued functions  $g : \mathbb{R}^D \rightarrow \mathbb{R}^D$ , such that every finite,  $D$ -dimensional linear restriction to function values  $g(T) \in \mathbb{R}^D$  at times  $T \in \mathbb{R}^D$  is a  $D$ -variate Gaussian distribution  $\mathcal{N}(g(T); \mu(T), k(T, T))$ . It is parametrised by two objects: A mean function  $\mu : \mathbb{R}^D \rightarrow \mathbb{R}^D$ , and a covariance function  $k : \mathbb{R}^{D_1} \times \mathbb{R}^{D_2} \rightarrow \mathbb{R}^{D_1 \times D_2}$ . The mean has a relatively straightforward role; it simply shifts predictions. The covariance function’s responsibility is more intricate. It can be interpreted as a similarity measure over function values over  $\mathbb{R}$  and controls the shape of the Gaussian process belief in the space of functions. It has to be chosen such that, for any  $T, T' \in \mathbb{R}^D$ , the matrix  $k(T, T') \in \mathbb{R}^{D \times D}$ , also known as the *kernel*, is positive semidefinite.

The utility of Gaussian process priors arises from the closure of the Gaussian exponential family under multiplication: If data points  $z_i$  are observed at times  $t_i$  with Gaussian noise  $\varepsilon$

$$z_i = g(t_i) + \varepsilon \quad \varepsilon \sim \mathcal{N}(0, \sigma^2), \quad (6)$$

then the posterior distribution under the Gaussian process prior and the Gaussian likelihood encoding this observation is also a Gaussian process [8, §2.2], with the posterior mean and covariance

$$\mu(t_*) = k_{t_* T}^\top K^{-1} \mathbf{z} \quad (7)$$

$$\Sigma(t^*, t_*) = k_{t^* t_*} - k_{t^* T}^\top K^{-1} k_{T t_*}, \quad (8)$$

where  $K = k(T, T) + \sigma^2 I$  is the kernel Gram matrix with observation noise  $\sigma^2$ ,  $\mathbf{z}$  is the set of observations and  $t^*, t_*$  are prediction times. For notational brevity, we are using the short-hand  $k_{tT} := k(t, T)$ . GP regression is a very general framework for nonlinear regression. In the context of our particular setup, it may in fact also be used to construct probabilistic models for fully nonlinear systems  $g = g(x, t)$ , without any major changes. For the purposes of this paper, however, we focus on the simpler case of  $g = g(t)$ , allowing for the use of linear control techniques, which can be efficiently implemented in practice.

GP regression is related to least-squares regression as the maximum of the Gaussian process posterior under a GP prior with zero mean and kernel  $k$  is identical to the least-squares estimator regularised in the reproducing kernel Hilbert space (RKHS) reproduced by  $k$ . The difference in the frameworks

is chiefly conceptual. Gaussian process regression interprets the exponential of the RKHS regulariser  $\|g\|_k^2$  as a generative (Gaussian process) model (prior) for  $g$ . The added value of this probabilistic interpretation is manifold. For example, it provides, after finitely many observations, a description of the system to be identified in terms of a stochastic differential equation. Here, we are making use of this description to reason about the structure of the hypothesis class, but it can also be used for more advanced ends, for example to guide active system identification [5].

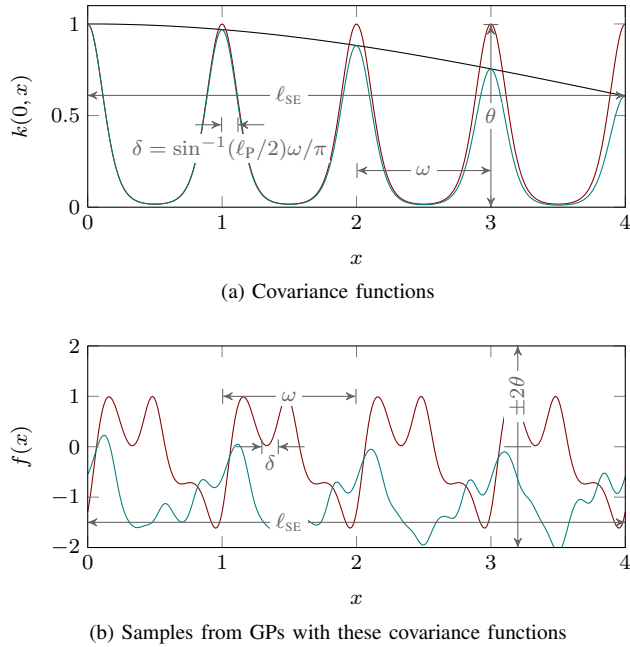


Fig. 1. (a): The compound kernel  $k_C$  (green) of Equation (11) is the product of  $k_{SE}$  (black, Eq. (9)) and  $k_P$  (red, Eq. (10)). (b): Samples drawn from Gaussian process priors using these covariance functions (same colours). Samples using the periodic kernel are perfectly periodic, while samples using the compound kernel are only periodically similar on a scale controlled by the parameter  $\ell_{SE}$  of Eq. (11). This “local periodicity” can be used to increase modelling flexibility.

The way to construct a periodic hypothesis class, and the central idea of this work, is to construct a covariance function that focuses prior probability mass on periodic functions (or, in least-squares parlance, whose RKHS consists of functions that are locally periodic): Among the most popular kernels for regression purposes is the square-exponential (aka. radial basis function, Gaussian) kernel

$$k_{SE}(t, t'; \ell_{SE}) = \exp\left(-\frac{(t - t')^2}{2\ell_{SE}^2}\right), \quad (9)$$

with length-scale  $\ell_{SE}$ . This kernel gives a stationary model which does not allow for structured extrapolation. [13] proposed constructing periodic priors, through a sine-transformation of the input:

$$k_P(t, t'; \ell_P, \omega) = \exp\left(-\frac{2 \sin^2\left(\frac{\pi}{\omega}(t - t')\right)}{\ell_P^2}\right), \quad (10)$$

with length-scale  $\ell_P$  and period-length  $\omega$ . Function values  $g(t), g(t')$  jointly sampled from Gaussian process priors with

this covariance function are perfectly correlated if  $t - t' = \omega$ , thus sampled functions are perfectly periodic with period  $\omega$ . Within this period length, samples vary on a typical regularity length scale of  $\delta = \sin^{-1}(\ell_P/2)\omega/\pi$ , over a range with standard deviation  $\theta$  (see Figure 1). For many systems, strict periodicity is too strong an assumption. To weaken the perfect correlation, we use the fact that the kernel property is closed under multiplication and addition (i.e. kernels form a semiring, [8, Section 4.2.4]): Although the product of two Gaussian processes is not a Gaussian process, the product of two kernel functions is also a kernel and therefore a valid covariance function for a Gaussian process. The same holds true for the summation of kernels.

A suitable kernel can thus be constructed in a qualitative manner, using physical knowledge about the system. Multiplying the periodic kernel with a broader square-exponential gives another kernel whose corresponding Gaussian process condenses mass at functions that change over time-scales:

$$k_C(t, t'; \theta^2, \ell_{SE}, \ell_P, \omega) = \theta^2 \cdot k_{SE}(t, t'; \ell_{SE}) \cdot k_P(t, t'; \ell_P, \omega), \quad (11)$$

with signal variance  $\theta^2$  and the other parameters as stated above. This kernel considers two input times similar if they are similar under both the square-exponential *and* the periodic kernel. If  $\ell_{SE} \gg \omega$ , this allows encoding a decay in the periodic behaviour over several oscillations. The different covariance functions are shown in Figure 1a, exemplary randomly sampled functions from Gaussian processes with those covariance functions are shown in Figure 1b. The posterior mean of GPs with nonperiodic and periodic covariance, trained on periodic data, is shown in Figure 2. Predictions far away from datapoints are equal, whereas close predictions show more structure with the locally periodic kernel.




Fig. 2. Comparison of Gaussian process posteriors (posterior mean function as thick line, shaded region covers two marginal standard deviations) arising from the same, periodic data (black) for the square-exponential kernel (top) and a product of periodic and square-exponential kernel (bottom). The locally periodic kernel provides a richer extrapolation, which can be used for improved model predictive control.

The combined kernel from Equation (11) has four hyperparameters (which we will subsume in the vector  $\eta :=$

$(\theta^2, \ell_{SE}, \ell_P, \omega)$ ), including the frequency  $\omega$  of the periodicity. Inferring these hyperparameters  $\eta$  is important for good modelling performance. The fundamental framework for such inference is provided by Bayes' theorem. However, the likelihood for  $\eta$  under  $\mathbf{z}$  is a nontrivial (and non-Gaussian) function of  $\eta$ , given by

$$p(\mathbf{z}|T, \eta) = \int p(\mathbf{z}|g, T, \eta)p(g)dg = \mathcal{N}(\mathbf{z}; 0, K(\eta)). \quad (12)$$

A fast, popular, and theoretically well-founded way of achieving at least a point estimate for  $\eta$  (as opposed to a full posterior) is to find the maximum of the posterior, an approach known as *evidence maximisation*, or as the regularised form of “type-II maximum likelihood” [8].




Fig. 3. Slice through the hypothesis space of hyperparameters along the dimension of the hyperparameter  $\omega$  defining the period length of  $g$ . logarithm of the type-II (marginal) likelihood in thin blue, logarithm of posterior distribution (sum of log likelihood and vague log prior, not shown) in green. The shape of the log posterior is dominated by the likelihood, indicating that most prior assumptions are dominated by the observed data.

The optimisation of the hyperparameters is easier when performed on the logarithm of the likelihood, which is

$$\log p(\mathbf{z}|T, \eta) = -\frac{1}{2}\mathbf{z}^T K(\eta)^{-1}\mathbf{z} - \frac{1}{2}\log |K(\eta)| - \frac{D}{2}\log 2\pi. \quad (13)$$

Even in the log domain, this objective function is not likely convex. Regularisation, by adding the logarithm of a prior, makes the optimisation problem identifiable, and more regular (albeit not necessarily convex)

$$\begin{aligned} \arg \max_{\eta} p(\eta|\mathbf{z}, T) &= \arg \max_{\eta} p(\mathbf{z}|T, \eta)p(\eta) \\ &= \arg \max_{\eta} (\log p(\mathbf{z}|T, \eta) + \log p(\eta)). \end{aligned} \quad (14)$$

The computational cost of adding a prior is negligible compared to the matrix inversion needed for the calculation of the likelihood in Equation (13). An advantage of using this kind of hierarchical (“type II”) inference over other, ad-hoc estimators is that it automatically enjoys theoretical guarantees: Since the set of hyperparameters is finite, the shape of the posterior is more and more dominated by the likelihood as the number of data points increases. If the likelihood is identifiable and sufficiently regular, the prior eventually becomes irrelevant and, if the other modelling

assumptions are correct, converges to correct values [14], [15].

As priors for the hyperparameters, we use Gamma distributions of the following form, which provide a comparably flexible family of priors for non-negative variables

$$\log p(\eta_i) = (k-1)\log(\eta_i) - \frac{\eta_i}{t} - \Gamma(k) - t^k, \quad (15)$$

where  $\eta_i$  is one single hyperparameter,  $t$  and  $k$  are the parameters of the Gamma prior and  $\Gamma$  is the Gamma function. To illustrate the process, the negative log likelihood and posterior are shown for one hyperparameter, the period length  $\omega$ , in Figure 3.

### B. The discrete-time model

To utilise the identified error dynamics in control, a time-varying affine discrete-time model is derived, which can then be directly incorporated in standard linear control and estimation techniques. Since the periodic function  $g(t)$  cannot be modelled, an approximation is derived by using the mean prediction  $\tilde{g}(t)$  of the previously described Gaussian process, evaluated at time  $t$

$$\dot{\hat{x}} = Ax + Bu + \tilde{g}(t). \quad (16)$$

With the discretisation of Equation (4), the approximate discrete-time model results to

$$x_{k+1} = A_d x_k + B_d u_k + \tilde{a}_k, \quad (17)$$

where  $\tilde{a}_k$  is obtained by integrating system (16) (with a numerical ODE-solver) and subtracting the linear terms:

$$\tilde{a}_k = \int_{k\Delta t}^{(k+1)\Delta t} \dot{\hat{x}}(x, u, \tau) d\tau - A_d \hat{x}_k - B_d u_k. \quad (18)$$

The discrete model (17) can in principle be applied to all linear state estimation or control methods. We show how to use the advanced model prediction in a Kalman filter and in model predictive control in the next sections.

### C. State filtering

Since the state measurements are subject to noise, a Kalman filter is introduced to estimate the state. The Kalman filter consists of two basic steps: The prediction of the state through the dynamic equations and the correction of this prediction with the current measurement. To incorporate the nonlinear prediction from the GP without complicating the Kalman filter equations, the Kalman filter utilises Equation (3) where the approximation  $\tilde{a}_k$  from Equation (18) is used. The measurement update of the Kalman filter remains unchanged.

Without the incorporation of the predicted affine part  $\tilde{a}_k$ , the periodic error would be classified as process noise, leading to an increase in the Kalman gain. This increase then puts more weight on the noisy measurements in the update equation, which in turn deteriorates estimation performance.

#### D. Control

Model predictive control (MPC) is a widely used technique, which has become standard in many industrial applications, particularly the process industry. See [16] for an introduction. Because MPC requires the solution of an optimal control problem at every sampling time, it was classically used for rather slow dynamic systems. In recent years, research has focused on expanding the scope of MPC and enable it for systems with fast dynamics.

MPC uses a prediction of the dynamics in order to compute an optimal sequence of control inputs according to a given cost function. The first control input is applied to the system and the optimisation is repeated for the estimated (or measured) state at the next sampling time in a so called receding horizon fashion. This framework allows to directly incorporate the derived discrete-time model

$$\min J = \sum_{n=0}^{N-1} l(x_n, u_n) \quad (19a)$$

$$\text{s.t. } x_{n+1} = A_d x_n + B_d u_n + \tilde{a}_n \quad (19b)$$

$$x_0 = \hat{x}_k \quad x_n \in \mathbb{X} \quad u_n \in \mathbb{U}, \quad (19c)$$

where  $\hat{x}_k$  is the current estimated state at the  $k$ 'th time step,  $l : \mathbb{R}^E \times \mathbb{R}^F \rightarrow \mathbb{R}_+$  is the cost, which is assumed to be a convex, positive definite function and  $\mathbb{X}, \mathbb{U}$  are polytopic sets.  $N$  is the length of the prediction horizon. The resulting problem is a convex optimisation problem that can be solved efficiently using available optimisation software.

#### IV. EXPERIMENTS

##### A. Toy problem

As a simple problem for providing intuition, consider the following linear dynamic (double integrator) system with an additive time-periodic component  $g$ :

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t) + g(t) \quad (20)$$

$$g(t) = \begin{bmatrix} \sin(t) \\ \cos(1.3t) \end{bmatrix}. \quad (21)$$

The goal is to control the first state of the system to the origin. We use a quadratic cost function

$$l(x_n, u_n) = \frac{1}{2} x_n^T Q x_n + \frac{1}{2} u_n^T R u_n. \quad (22)$$

We set the state cost to a diagonal matrix  $Q = \text{diag}(1000, 1)$ , which is mainly penalising the first state, while keeping the second state bounded. The weight on the control inputs is set to  $R = 10$ , allowing for aggressive control behaviour. The horizon length of the MPC is set to  $N = 15$ . State and input constraints are omitted for simplicity.

This system was simulated numerically with a sampling rate of 1 Hz. Figure 4 shows control inputs and resulting state trajectories, for a model predictive controller without information about the nonlinearity  $g$ , and a model predictive controller using the posterior mean function of a periodic Gaussian process regressor as a model for  $g$ . The behaviour

shown in these plots is relatively robust to changes in the cost functions  $Q$  and  $R$ , because, by Equation (20), this system can only correct for errors in the first state.

After an identification phase in the first 5s of the experiment, the GP based controller shows a drastic performance improvement. Omitting the first 5 seconds of identification phase, RMS error drops by almost 90%, from 0.94 for the linear model to 0.099 for the GP based controller. Speaking more qualitatively, Figure 4 also shows less residual structure in the controlled state  $x_1$ .

While the GP based controller is effective at removing the periodic structure from the first controlled state, the regression model itself remains able to predict the periodic error correctly into the future, even when trained exclusively on controlled states. This is possible because the regression model explicitly models the controlled dynamics, so it can account for the shift of periodicity from the states to the control input. This feature of the framework is crucial for identifying controlled systems (see next section). In contrast, a fixed parametric model of the dynamics could over time result in misaligned predictions and, in the worst case, even lead to resonant amplification of error.




Fig. 4. Closed-loop state and input trajectories under the MPC controller. MPC using a linear model is shown as thin lines, the GP based MPC controller as thick lines.

##### B. Periodic error correction for telescope mounts

The original motivation for the work presented here is the control of periodic errors in astrophotography systems. Telescope mounts correct for Earth's rotation relative to the sky by a linear motion at the sidereal velocity. This motion is typically produced by mechanical devices using worm gears, which gives rise to periodic deviations. Because contemporary telescopes, even those used by amateurs, have high optical resolution, and images are taken with long exposure times, these mechanical imprecisions are frequently the dominant source of error on astronomical photographs.

Existing *periodic error correction* systems require careful system identification by the user of the telescope. The corresponding measurements need to be repeated after every repositioning of the telescope (that is, multiple times per night), and still regularly lead to unsatisfactory performance.

A problem specific to this astronomical application is that state measurement is performed by taking photographs of the night sky, which requires relatively long exposure times, so that the measurement frequency can reach the order of  $1/10$  of that of the periodic error. This is precisely the domain in which we expect to see utility from a periodic model.

We use the presented framework consisting of a GP to learn the error input, a Kalman filter to estimate the state and an MPC controller.

1) *Simulation*: The time constant of the periodic error in telescopes is relatively slow. To allow rapid prototyping, we designed a simulated telescope system. This also showed that the speed is not a relevant state, only the position is relevant and it can be directly influenced by the input of the system. The linear component of the system is well controlled, and can thus effectively be ignored. The telescope system reduces to

$$\dot{x} = u + g(t), \quad (23)$$

with an unknown function  $g(t)$  that is periodic, but changes slowly over time.

Figure 5 shows simulation results that empirically confirm the intuition from Section II that the benefit of periodic prediction in control depends on the sampling rate: On the numerical simulation of the telescope, we compare, for various sampling rates of the state

- a simple Kalman-filter based MPC controller using the linear model without explicit model for  $g$
- two MPC controllers both using a nonparametric, but fully stationary (i.e. not periodic) models for  $g$  using a GP model with the square-exponential covariance function. One of these models uses a length scale smaller than the periodicity (so it can extrapolate periodic swings locally, but not beyond one period), the other a scale longer than the periodicity (so it averages over the periodic variations)
- two MPC controllers using instances of the periodic model for  $g$  described in this paper; one in which the hyperparameters are fixed to a good value a priori (amounting to the assumption that the period of  $g$  is known), the other using the full setup described before, in which hyperparameters are learnt by type-II maximum likelihood during identification.

Since we are only interested in the limit performance in this particular experiment, all the controllers were run for an extended identification phase of 10 period lengths to avoid artefacts from the identification. Figure 5 then shows control RMS error after this phase as a function of the sampling time. Between measurements, the MPC controllers are operated in an open-loop mode, the control inputs are obtained from the last MPC optimisation.

The results are as expected: For sampling rates much higher than  $\omega$ , the dynamics are locally linear, and all models perform equally well. For sampling rates between 10% and 80% of  $\omega$ , the periodic model offers considerable benefit. When the sampling rates are lower than the periodicity, all models perform some kind of averaging (each of them a




Fig. 5. Comparison of the error at different sampling rates (in simulation). MPC control inputs are computed at the indicated sampling times, shown as a fraction of the period length. The control rate was 1 Hz for every plot. Between measurements, the MPCs are operated in an open-loop mode. For this figure only, the velocity is assumed measurable and  $g(t)$  is set to be a sine with fixed frequency  $\omega = 100$  s.

different one). In this regime, the Nyquist rate imposes limits on identifiability of the system, which adversely affects the performance of the periodic nonparametric model. If  $\omega$  is known precisely, very good control is possible even for sampling rates lower than  $\omega$ . The green line in Figure 5 represents the performance of a system fully ignorant of  $\omega$  in the beginning. One can expect prior information about  $\omega$  of varying vagueness to give performance somewhere in the region between the green and blue curves in Figure 5. Of course, the case where sampling rate and  $\omega$  are equal is special, since then  $g$  appears constant, and even the informed periodic model can only ever learn the behaviour of  $g$  at one unique point during the period.

2) *Hardware*: We also tested our implementation on the real physical system. We use a commercially available and unmodified Vixen Sphinx<sup>1</sup> telescope mount for the tests (see Figure 6). Without closed-loop control, this mount shows about 20 arcs of RMS error. Because outdoor measurements are time-consuming and subject to difficult to control external effects, we use an experimental setup with a second, high precision gearless ASA DDM60Pro<sup>2</sup> telescope mount equipped with a laser “star” as tracking reference. It shows a typical pointing error of about 0.1 arcs. The measurement is done with a Canon EF400DO lens<sup>3</sup> on a The Imaging Source DMK 41AU02.AS<sup>4</sup> camera.

For the hardware interaction, an open source available telescope guiding software (“PhD Guiding”<sup>5</sup>) is used. In the original implementation this software uses a deadbeat controller. The telescope is connected to the computer with a

<sup>1</sup><http://www.vixenoptics.com/mounts/sphinx.htm>

<sup>2</sup>[http://www.astrosysteme.at/eng/mount\\_ddm60.html](http://www.astrosysteme.at/eng/mount_ddm60.html)

<sup>3</sup>[http://www.canon-europe.com/For\\_Home/Product\\_Finder/Cameras/EF\\_Lenses/Telephoto/EF\\_400mm\\_f4\\_DO\\_IS\\_USM/](http://www.canon-europe.com/For_Home/Product_Finder/Cameras/EF_Lenses/Telephoto/EF_400mm_f4_DO_IS_USM/)

<sup>4</sup><http://www.astronomycameras.com/products/usb/dmk41au02as/>

<sup>5</sup><http://www.stark-labs.com/phdguiding.html>



Fig. 6. The telescope mount used for the tracking experiments. On the right side is the camera lens used as guiding telescope. A main telescope is not used for the tests.

Shoestring GPUSB<sup>6</sup>, a device that sends puls-width modulated signals to telescopes over a proprietary 4-wire interface.

We altered the software to gain access to the measured displacement of the camera image. The value is sent through network socket to MATLAB, where the controller developed in this paper calculates the optimal control signal which is in turn sent back to the guiding software. The software then sends the control signal to the telescope hardware. For plotting and calculation of the RMS error, the measured displacement is converted from pixels into arc-seconds ( $1/3600$  of a degree) with an empirically determined conversion factor.

For real-time implementations, algorithmic complexity is relevant. The computational cost of the GP prediction scales with  $\mathcal{O}(D^3)$  in the number of datapoints. In order to prevent too much computational effort, we limited the number of stored data points to 90, which gives sufficient coverage over 90s, or about 3 periods of the main periodic component, which suffices for predictive control of this component.

For the prediction of the dynamics in the model predictive control, an ODE-solver is employed to predict the affine term from the mean of the GP prediction. This has manageable computational cost because the inference cost of a Gaussian process is dominated by the initial one-time operation of inverting the Gram matrix  $K$  in  $\mathcal{O}(D^3)$  time, while subsequent evaluations of the mean function at  $M$  times only has cost  $\mathcal{O}(MD)$  (see Equation (7)).

The optimisation of the hyperparameters is also an expensive part of this algorithm. As the kernel Gram matrix has to be built and inverted at every iteration of the optimisation, we have to keep the number of steps small at every sampling time. We use a BFGS optimiser, and pre-condition the optimiser used in time-step  $k$  by the Hessian estimate constructed by the optimiser in the preceding time step  $k - 1$ . With this modification, it is possible to run one single linesearch of the optimiser during the time the camera needs to take the next image.

<sup>6</sup><http://www.store.shoestringastronomy.com/gpusb.htm>

For both setups, one with a MPC based on the linear model without  $g$  and one with the GP prediction for the periodic error, the algorithm was run 3 times for 25 min each. Both the sampling and the control time were set to 3s. The results of these runs is shown in Table I. The RMS error drops by 22.64% through the use of GP predictions in this hardware setup.

	Run 1	Run 2	Run 3	Mean
Plain MPC	0.9839	1.0234	0.9353	<b>0.9809</b>
GP-MPC	0.7365	0.7792	0.7605	<b>0.7587</b>

TABLE I  
EXPERIMENTAL RESULTS (RMS ERROR, IN ARCS)





Fig. 7. Measurements of state for the physical experiments with the telescope, for plain MPC using only a linear model (a), and periodic Gaussian process model based MPC (b). Results from run 2 of 3 in our experiments, which gave the highest (worst) RMS error for both models.

## V. CONCLUSION

High precision control of dynamical systems requires precise models of even minor nonlinearities. Where analytical models are not available, they need to be constructed numerically from interaction with the system. Periodic nonlinearities are an especially promising domain in this regard, as they can be extrapolated well into the future. We have studied a nonparametric modelling framework based on a carefully crafted Gaussian process prior exhibiting a weak, localised form of periodicity. Because Gaussian regression returns models in the form of stochastic differential equations, they can be combined directly with existing control frameworks. Numerical and physical experiments confirm the intuitive result that the benefit of periodic models depends on the relative size of state sampling and disturbance frequencies. They also show considerable increases in control performance, confirming the practical utility of this framework.

## REFERENCES

- [1] T. Erm and S. Sandrock, "Adaptive periodic error correction for the VLT," in *Proceedings of SPIE*, vol. 4837, 2003, p. 900.
- [2] J. Crassidis and F. Markley, "Predictive filtering for nonlinear systems," *Journal of Guidance, Control, and Dynamics*, vol. 20, no. 3, pp. 566–572, 1997.
- [3] S. Yuen, P. Novotny, and R. Howe, "Quasiperiodic predictive filtering for robot-assisted beating heart surgery." in *ICRA*, 2008, pp. 3875–3880.
- [4] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, to appear.
- [5] P. Hennig, "Optimal reinforcement learning for Gaussian systems," in *Advances in NIPS 24*, 2011, pp. 325–333.
- [6] J. Ko and D. Fox, "GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models," *Autonomous Robots*, vol. 27, no. 1, pp. 75–90, 2009.
- [7] J. Hall, C. E. Rasmussen, and J. M. Maciejowski, "Modelling and control of nonlinear systems using Gaussian processes with partial model information," in *CDC*, 2012, pp. 5266–5271.
- [8] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [9] I. Steinwart, "On the influence of the kernel on the consistency of support vector machines," *The Journal of Machine Learning Research*, vol. 2, pp. 67–93, 2002.
- [10] C. A. Micchelli, Y. Xu, and H. Zhang, "Universal kernels," *The Journal of Machine Learning Research*, vol. 7, pp. 2651–2667, 2006.
- [11] A. Van der Vaart and J. Van Zanten, "Rates of contraction of posterior distributions based on gaussian process priors," *The Annals of Statistics*, pp. 1435–1463, 2008.
- [12] A. van der Vaart and H. van Zanten, "Information rates of nonparametric gaussian process methods," *J. Mach. Learn. Res.*, vol. 12, pp. 2095–2119, July 2011.
- [13] D. MacKay, "Introduction to Gaussian processes," *NATO ASI Series F Computer and Systems Sciences*, vol. 168, pp. 133–166, 1998.
- [14] L. LeCam, "Convergence of estimates under dimensionality restrictions," *The Annals of Statistics*, pp. 38–53, 1973.
- [15] I. A. Ibragimov and R. Z. Hasminskii, *Statistical Estimation: Asymptotic Theory*. Berlin: Springer-Verlag, 1981.
- [16] J. Rawlings and D. Mayne, *Model Predictive Control: Theory and Design*. Nob Hill Publishing, 2009.