Compensating torque ripples in a coarse pointing mechanism for free-space optical communication

A Gaussian process repetitive control approach

Mooren, Noud; van Meer, Max; Witvoet, Gert; Oomen, Tom

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Compensating torque ripples in a coarse pointing mechanism for free-space optical communication: A Gaussian process repetitive control approach ☆

Noud Mooren [a], Max van Meer [a], Gert Witvoet [a,b,*], Tom Oomen [a,c]

[a] *Eindhoven University of Technology, Department of Mechanical Engineering, Eindhoven, The Netherlands*
[b] *TNO, Optomechatronics department, Delft, The Netherlands*
[c] *Delft University of Technology, Delft Center for Systems and Control, Delft, The Netherlands*

A B S T R A C T

Actuators that require commutation algorithms, such as the switched reluctance motor (SRM) considered in this paper and employed in the coarse pointing assembly (CPA) for free-space optical communication, often have torque-ripple disturbances that are periodic in the commutation-angle domain that deteriorate the positioning performance. The aim of this paper is to model the torque ripple as a Gaussian Process (GP) in the commutation-angle domain and consequently compensate for it at arbitrary velocity. The approach employs repetitive control (RC) at a constant velocity. A spatial GP with a periodic kernel is trained using data that is obtained from the RC step resulting in a static non-linear function for compensation at arbitrary velocity. Stability conditions are provided for both steps. The approach is successfully applied to a CPA prototype to improve the tracking performance for laser communication, where the torque ripple is compensated at arbitrary velocity.

1. Introduction

Many future technological advances, including scientific earth observations, 5G and 6G connectivity, and the Internet-of-Things, rely heavily on high-throughput and secure data communication with worldwide coverage. This calls for new advances in communication technology between satellites and from satellites to ground stations, well beyond the possibilities of the current radio-frequency (RF) standard. Recent developments in (inter-)satellite communication exploit free-space optical communication (FSOC), where a transmitter points a laser beam towards a receiver over a very long distance. Distances range from several hundreds of kilometers from lower-earth orbit (LEO) to ground, to thousands of kilometers between LEO satellites, up to tens-of-thousands of kilometers for links between ground stations and satellites in a geostationary orbit (GEO). FSOC terminals can be much faster and lighter than their RF counterparts, partly because all the transmitted power can be condensed in a single beam with very small divergence, which on the other hand requires pointing accuracies in the order of micro-radians or even sub-micro-radians to cover such large distances. At the same time, FSOC terminals need to point their lasers over very large angular ranges due to the relative motion between transmitter and receiver.

To enable a large-stroke rotation for tracking moving objects, a so-called coarse pointing assembly (CPA) is commonly used, in which its size, weight, and cost are critical aspects for commercial (large series) space applications. Traditional CPAs are quite often relatively large, heavy, and expensive compared to the rest of the FSOC terminal, therefore, in [1] a novel CPA actuation concept has been developed, see Fig. 1, which offers a low-cost, low-power alternative for accurate long-stroke laser pointing for satellite-based FSOC. This CPA concept, developed by the Netherlands Organisation of Applied Scientific Research (TNO), uses a switched reluctance motor concept [2] consisting of three motor phases and a commutation strategy to rotate a toothed disk and is relatively cost-effective to manufacture.

The positioning performance of the CPA depends on the interplay between the commutation functions and the mechanical tolerances; imperfect commutation easily leads to a torque ripple, e.g., due to misalignment of the three motor phases. Compensating this torque ripple disturbance can thus significantly improve the CPA performance.
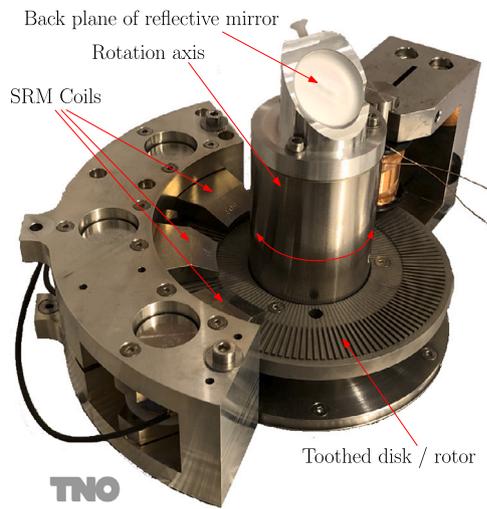
Fig. 1. Picture of the single-axis coarse pointing assembly (CPA) prototype, developed and tested at the Netherlands Organisation of Applied Scientific Research (TNO).

The commutation strategy is typically a trade-off between minimizing the torque ripple on one hand, and improving the power efficiency on the other hand [2,3], which are both key aspects for commercial space-based FSOC. Moreover, the torque ripple that appears is repeating in the commutation-angle domain. In particular, if the CPA operates with a constant velocity, then this disturbance is periodic in the time domain, and it is non-periodic otherwise. The optimal torque-ripple compensation will also differ from CPA to CPA, due to manufacturing and assembly tolerances, which calls for automated routines to keep this CPA a cost-effective alternative.

Data-driven learning of a torque ripple compensation function can significantly improve control performance, see, e.g., [4], alternative model-based approaches require more modeling effort [5]. Repetitive control (RC) enables asymptotic rejection of time-domain periodic disturbances with a known period, by learning an internal model from data [6–8]. In [9], a Gaussian process (GP)-based internal model is developed to incorporate prior knowledge for RC, which can be used to enable torque ripple compensation at constant velocity. Moreover, a spatial GP RC approach is presented in [10,11] for spatially-periodic disturbances that can be non-periodic in the time domain such as the torque ripple disturbance at varying velocity. However, this latter spatial approach updates continuously, requiring a substantial computational load, whereas an automated routine to compensate for the commutation-induced torque ripple does not necessarily require continuous updating, since it principally does not change over time. To avoid computational complexity in a low-power space-based environment, it is therefore preferred to learn a static compensation function that compensates for the torque ripple for arbitrary velocities.

Although recent developments in temporal RC can improve the positioning performance of the CPA for constant velocities, an approach that is computationally attractive and that allows compensating for the torque ripple at arbitrary velocity is not yet developed. The aim of this paper is to compensate for the torque ripple at arbitrary velocities by learning a static GP-based function in the position domain. The approach consists of two steps; first, time-domain GP-based RC [9] is applied for compensation at a constant velocity, and second, a spatial GP with a periodic kernel and data from step one is trained, resulting in a continuous compensation function that is periodic in the commutation-angle domain for compensation at arbitrary velocities. For the first step, identifying a correct compensation signal that

correlates with the commutation position, alternative approaches such as iterative learning control (ILC) [12] or adaptive robust control (ARC) [13,14] could in principle also be considered, but the unified and versatile GP-based RC approach of [9] is considered to be most effective, due to its simplicity in tuning and implementation, and the non-resetting nature of the commutation disturbances. Moreover, it easily allows to only compensate for disturbances with the correct periodicity; as desired, non-periodic disturbances are not learned, allowing for extrapolation to arbitrary velocities in the next step. For this second step, generating the compensation function for arbitrary velocities, the GP approach provides a non-parametric way to create a continuous function, which allows for accurate evaluation at any position. As such it avoids discretization and interpolation effects in traditional approaches such as look-up tables [11]. Summarizing, the following contributions are identified:

C1 analysis of the torque-ripple disturbance in the CPA to show that it is periodic in the commutation-angle domain, and modeling of the CPA as a linear system with a static non-linearity that represents the torque ripple (Sections 2 and 3),

C2 a two-step approach for learning of a static GP for torque ripple compensation at arbitrary velocity, including stability analysis for both steps (Section 4); and

C3 experimental validation of the two-step approach for arbitrary velocity, including experimental validation of GP-based RC in [9] for constant velocity, on the CPA (Section 5).

The application of this method enables over 300% performance improvement in the CPA by suppressing a substantial part of the torque ripple at arbitrary velocity. Finally, conclusions are presented in Section 6.

2. Problem definition

The CPA depicted in Fig. 1 is the main motivation for the work in this paper and is used for experimental validation. In this section, the setup is further introduced and the formal control problem is outlined.

2.1. Application motivation

The purpose of a CPA, such as the single-axis prototype depicted in Fig. 1, is to accurately rotate the last mirror in the optical path of a moving terminal, so as to accurately point the exiting laser towards a receiving moving object. In control terms, this means that the rotary stage carrying the mirror needs to track a reference signal that depends on the relative position, orientation, and velocity between the two moving objects; in the general case the rotational velocity of this stage thus varies over time. Fig. 2 shows an example of such satellite-level angular velocities for a 500 km LEO-to-ground link; the right plot is indicative for typical CPA velocities, although the exact numbers per axis will depend on the geometry and optical layout in a particular CPA design. The CPA presented here uses a switched reluctance motor (SRM) containing three motor phases to drive a toothed disk with 131 teeth [1]. A cross-section of the coils and the toothed disk is schematically shown in Fig. 3; by alternating the current between the three coils (through commutation functions) a continuous rotating motion can be generated. Important benefits of this actuator concept are the absence of physical contact between the coils and the rotating disk to reduce wear and the associated contamination, and the absence of cogging disturbances. However, due to possible imperfect commutation between the different coil currents (see Section 3.1), there is a torque ripple present that could hamper performance.
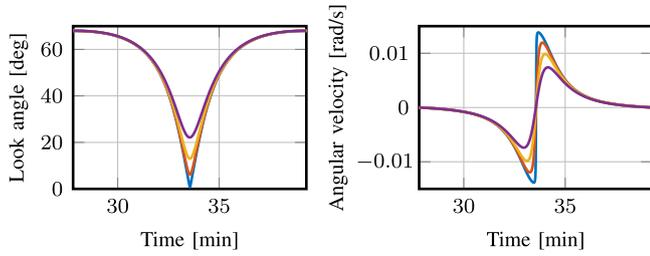
Fig. 2. Illustration of the look angle evolution between a LEO satellite at 500 km height and a ground station during a satellite overpass (left plot) and its time derivative (right plot), ranging from overpasses nearly directly above the ground station (——) to overpasses 200 km away (——).
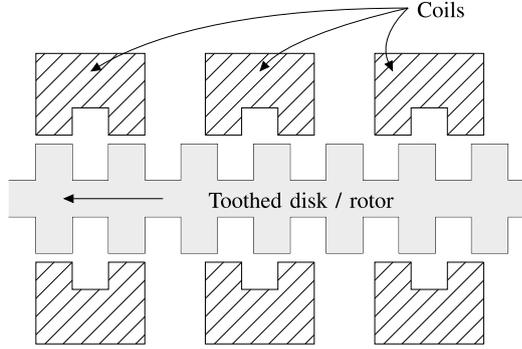
Fig. 3. CPA actuation principle with three coils to drive the toothed disk (rotor). For the depicted rotor position, a current through the first coil will give no torque, through the second coil a negative torque (to the right), and through the third coil a positive torque (to the left). This torque distribution will change as the rotor moves along the coils. Note that this SRM actuation concept is fully non-contact, hence mechanical abrasion and wear of the teeth or coils are highly unlikely.
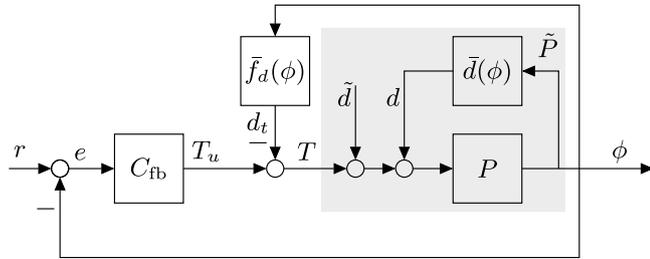
Fig. 4. Torque ripple compensation strategy where the aim is to learn the compensation function $\bar{f}_d(\phi)$ from data. The CPA is modeled as $\tilde{P}$ with non-linearity $\bar{d}(\phi)$ that is spatially periodic such that $d(k)$ represents a torque ripple or commutation-induced disturbance, and additional non-repeating disturbances $\tilde{d}$.

2.2. Control setting

The relevant control setting for the CPA is depicted in Fig. 4, where the goal in this paper is to design the static function $\bar{f}_d(\phi)$ to compensate for the torque ripple disturbance $\bar{d}(\phi)$, with ϕ the rotor position. The non-linear system $\tilde{P}$ represents the CPA including an existing commutation strategy and C_{fb} is a stabilizing feedback controller. It is assumed that the non-linear system $\tilde{P}$ is composed of a linear time-invariant (LTI) part P in conjunction with the static non-linearity $\bar{d}(\phi) : \mathbb{R} \mapsto \mathbb{R}$ that is a function of the rotor position. The resulting signal that acts on the input to P

$$d(k) = \bar{d}(\phi(k)), \tag{1}$$

represents the torque ripple or commutation-induced disturbance to be compensated. Additional non-repeating disturbances, denoted by $\tilde{d}$, may appear in the CPA. Since it is reasonable to assume that these do

not correlate with the position, they have very little influence on the presented approach and are therefore considered out of the scope of this paper. The notation $\bar{\cdot}$ refers to spatial functions.

Ideally, the compensation function to be designed $\bar{f}_d(\phi)$ is an exact model of $\bar{d}(\phi)$, i.e., if

$$\bar{f}_d(\phi) = \bar{d}(\phi), \tag{2}$$

then the commutation-induced disturbance $d(k)$ is perfectly compensated. To show this, assume that the reference is suppressed, e.g., with an additive feedforward $T_{\text{ff}} \approx P^{-1}r$ on T_u, then the commutation-induced error is given by

$$e = -PS(d - d_t), \tag{3}$$

where $S = (I + PC_{\text{fb}})^{-1}$ is the sensitivity function. If $\bar{f}_d$ satisfies (2), then $d(k) = d_t(k)$ and the error (3) is zero for arbitrary velocity variations.

For the remainder of this paper the following properties of the non-linear function $\bar{d}$ are considered.

Definition 1. The unknown static commutation-induced non-linearity $\bar{d}(\phi)$ is periodic, i.e.,

$$\bar{d}(\phi) = \bar{d}(\phi + \beta), \tag{4}$$

with spatial period β that is equal to the tooth pitch in the CPA.

Assumption 1. The unknown static non-linearity $\bar{d}(\phi)$ is sector bounded, i.e.,

$$\gamma_1 |\phi| \le \bar{d}(\phi) - \xi \le \gamma_2 |\phi| \tag{5}$$

for some $\gamma_1, \gamma_2 \in \mathbb{R}$ with $\gamma_1 < 0 < \gamma_2$ and $\xi \in \mathbb{R}$.

Assumption 1 is required to provide stability conditions in the remainder of this paper and is experimentally validated in Section 3. Definition 1 implies that $d(k)$ is inherently non-periodic in the time-domain, unless the angular position $\phi(k)$ increases linearly, i.e., when the rotational velocity is constant. It is therefore similar to a torque ripple disturbance that appears periodic in time with a constant rotational velocity, and non-periodic otherwise. The approach presented in this paper utilizes this to identify $\bar{d}(\phi)$ at constant velocity, to derive a compensation $d(k)$ for arbitrary and non-constant velocities.

2.3. Two step control approach

A two-step design method is developed to learn the spatially periodic function $\bar{f}_d$ from data. In the first step, data is collected to learn $\bar{f}_d(\phi)$ in the second step. To collect data for learning, recap that $d(k)$ is periodic in the time-domain with a known period if the rotational velocity is constant. However, due to imperfect commutation, among others, the velocity might vary slightly in practice. In this situation, asymptotic rejection of $d(k)$ can be obtained through a robust repetitive control (RC) approach [7,9,15] in the setting in Fig. 5, i.e., R enables to asymptotically learn a model $d_t^c(k)$ of the periodic disturbance $d(k)$ in a time-domain buffer, even in the presence of a non-perfect model.

Compensation at an arbitrary velocity is done in the second step in the setting in Fig. 4. It requires that the function $\bar{f}_d(\phi)$ is defined for all $\phi \in \mathbb{R}$, i.e., as a continuous function. Hence, the continuous and periodic function $\bar{f}_d(\phi)$ is learned in the commutation domain ϕ based on the time-domain data from step one, containing discrete points at various positions $\phi(k)$. Hence, interpolation and extrapolation of the data are required which is done by Gaussian process (GP) regression with a smooth and periodic prior as outlined in the remainder of this paper.

3. CPA disturbance analysis

In this section, the existing commutation strategy that is designed for the CPA is briefly discussed, which serves as a starting point for the work presented in this paper. Moreover, it is shown that the CPA can indeed be modeled as in Fig. 4, i.e., that a disturbance is present that is periodic in the commutation-angle domain.
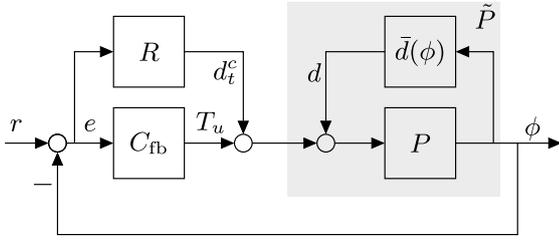
Fig. 5. Step 1 of the developed approach: the repetitive controller R is employed to learn the time-domain model d_t^c of the torque ripple d with constant velocity. The data d_t^c is consequently used to learn $\bar{f}_d(\phi)$ for compensation at arbitrary velocity in Fig. 4.
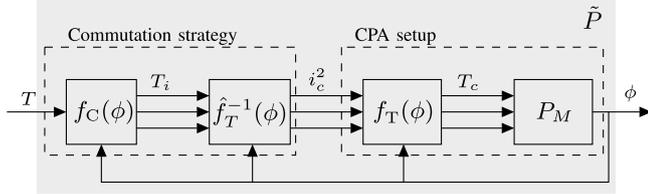
Fig. 6. Commutation strategy and the CPA setup.

3.1. CPA commutation strategy

The CPA setup with the commutation strategy for the three motor phases is schematically depicted in Fig. 6. A brief explanation of the commutation strategy is shortly summarized below for background, details can be found in [1]. The system $\tilde{P}$ in Fig. 4 contains the commutation strategy and CPA dynamics; the physical CPA setup without commutation, i.e., from the squared current of each of the coils i_c^2 to the rotor position ϕ is given by $P_M f_T(\phi)$ with $f_T(\phi) = \mathrm{diag}(f_{T,1}, f_{T,2}, f_{T,3})$ a position dependent function from coil current to torque output, and $P_M \in \mathcal{RH}_\infty$ the mechanical part that is proper and stable. Moreover, $f_{T,c}(\phi) : \mathbb{R} \mapsto \mathbb{R}$ for $c = 1, 2, 3$ is a non-linear function of ϕ that maps i_c^2 to the torque generated by coil c

$$T_c = \frac{1}{2}\frac{\partial \ell}{\partial \phi} i_c^2 = f_{T,c}(\phi) i_c^2, \tag{6}$$

where $\ell \in \mathbb{R}$ is the inductance and while neglecting magnetic saturation, see, e.g., [16, Chapter 5]. These $f_{T,c}$ are principally unknown, but have been estimated via dedicated torque measurements as described in [1]. The commutation strategy contains an approximate inverse of f_T denoted by $\hat{f}_T^{-1}$ such that ideally $f_T(\phi)\left(\hat{f}_T^{-1}(\phi)\right) = I$ for all $\phi \in [0, \beta)$ with β the tooth pitch, and $f_C(\phi) : \mathbb{R} \mapsto \mathbb{R}^3$ are the functions that distribute the required torque $T \in \mathbb{R}$ over the three individual motor phases. The functions f_C and $\hat{f}_T^{-1}$ for this specific setup are developed in [1] and used throughout the paper.

Remark 1. The functions $f_C(\phi)$ and $\hat{f}_T^{-1}(\phi)$ are periodic with a single tooth and defined for $\phi \in [0, \beta)$ with $\beta = \frac{2\pi}{131}$ the tooth pitch, and extend periodically outside this interval, i.e., $f_i(\phi) = f_i(\phi + \beta)$.

Due to imperfections in the CPA setup and the commutation strategy, e.g., the requested torque per coil T_i is not equal to the actual torque T_c, there is a torque ripple disturbance present that depends on the position output ϕ as in Fig. 4. Moreover, it is assumed that this torque ripple is repeating in the commutation-angle domain with a single tooth as in Definition 1.

3.2. Analysis of the disturbance

In this section, Definition 1 and Assumption 1 are experimentally validated to show that the torque ripple disturbance indeed repeats
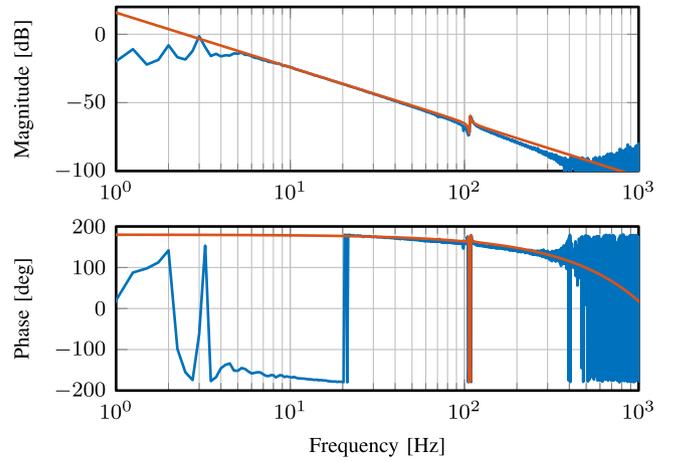
Fig. 7. Measured frequency response function (FRF) (——) of the CPA setup with a model fit (——). At frequencies $f < 4$ Hz the friction in the CPA bearings leads to a measurement error.
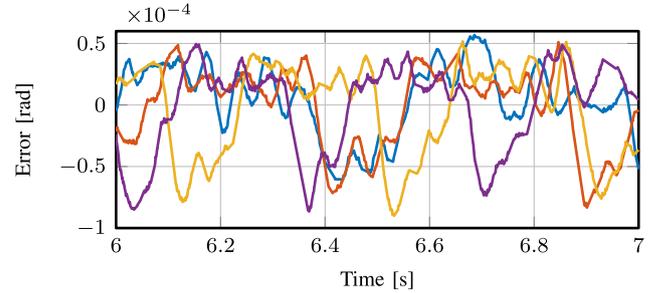
Fig. 8. Measurement of the error as a function of time for different velocities 0.08 rad/s (——), 0.1 rad/s (——), 0.12 rad/s (——) and 0.14 rad/s (——), which are all periodic but with different period.

with a single tooth. Experiments are performed with a feedback controller at several constant rotational velocities ranging from 0.08 rad/s up to 0.14 rad/s with 0.02 rad/s increments. Note that this is roughly a factor 10 faster than the satellite-level angular velocities shown in Fig. 2; this is partly to incorporate sufficient margin in the translation of these indicative velocities to single-axis CPA velocities, and partly to speed up the data collection process. The resulting error is shown in Fig. 8 as a function of time.

Under the assumption that the CPA can be modeled as in Fig. 4 with error (3), an estimate of the disturbance $\hat{d}$ can be obtained from the measured error as

$$\hat{d} = \widehat{PS}^{-1} e, \tag{7}$$

where $\widehat{PS}^{-1}$ is an approximate inverse model of the process sensitivity PS. The first transient part of the measurement data is omitted. A model $\widehat{PS}$ of PS is obtained with the measured frequency response function of the CPA in Fig. 7, using the existing commutation functions discussed in Section 3.1. Consequently, a stable inverse $\widehat{PS}^{-1}$ is computed with zero-phase-error-tracking-control (ZPETC) [17,18]. The resulting disturbance estimate is shown in Fig. 9 as a function of ϕ. The following observations can be made.

- The dominant part of the error in Fig. 8 is periodic in time for each velocity but the time-domain periods vary as the velocity changes.
- The disturbance estimates, i.e., filtered errors, as a function of the position in Fig. 9 are all periodic with a single tooth $\frac{2\pi}{131} = 0.048$ rad, independent of the velocity. Moreover, the disturbance estimates are similar for each velocity which motivates that the
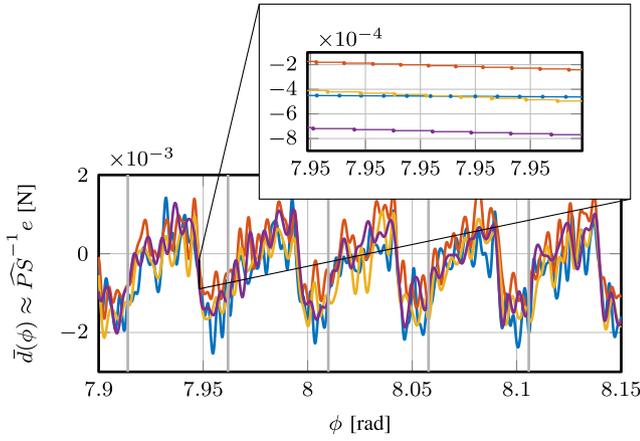
Fig. 9. Representation of the same error in Fig. 8, but filtered with an inverse model of PS and plotted as a function of the position ϕ [rad]. Now all measurements $\bar{d}(\phi)$ can be observed to have the same periodicity. Moreover, the zoomed inset shows that the data is sampled differently at each velocity.
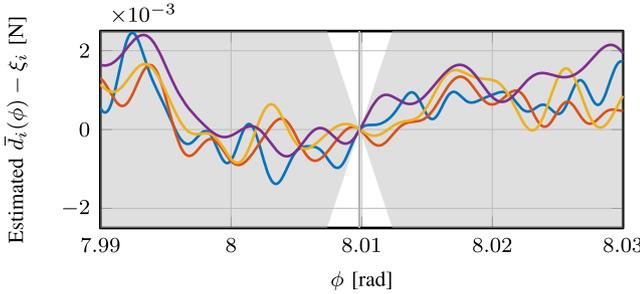
Fig. 10. Estimates of $\bar{d}(\phi)$ obtained from measured error data at different velocities, 0.08 rad/s (——), 0.1 rad/s (——), 0.12 rad/s (——) and 0.14 rad/s (——), to validate the sector bound (▨) condition in Assumption 1 for $-\gamma_1 = \gamma_2 = 1$ rad/N.

CPA can be modeled as in Fig. 4 and validates Definition 1. There are however additional disturbances visible in Fig. 9 that are not repetitive with ϕ (originating from e.g. the bearings or small differences in tooth geometry); this paper only addresses the compensation of the common commutation-induced or tooth periodic disturbance.

- Based on the representation of the measurement data in Fig. 10, which is a zoom of Fig. 9 around $\phi = 8.01$ rad, it can be concluded that the sector bound condition in Assumption 1 is valid for $\gamma_2 = -\gamma_1 = 1$, since all estimated $\bar{d}(\phi)$ remain within the gray sector-bounded areas.

Moreover, the zoomed inset in Fig. 9 shows that the data points are sampled differently for each velocity, i.e., non-equidistant in the rotor-position domain for a varying velocity. This yields that indeed the function $\bar{f}_d(\phi)$ for compensation must be defined for all ϕ to compensate for arbitrary velocity. The design of $\bar{f}_d(\phi)$ as a GP is outlined in the remainder of this paper.

Remark 2. The measured error can be used directly to approximate $\bar{d}$ using a model $\widehat{PS}$, however, the equality of this estimate can be poor depending on the model quality. For this reason, in the developed two-step approach, RC [9] is used that is robust for modeling errors up to 100% and for slight velocity variations [15].

4. Gaussian process torque ripple compensation

In this section, the two design steps to learn the compensation function $\bar{f}_d(\phi)$ in Fig. 4 are presented. It is shown first, how the continuous compensation function $\bar{f}_d(\phi)$ is learned given a data set with discrete observations of the disturbance through GP regression and suitable kernel design, which is the main contribution. This allows to evaluate $\bar{f}_d(\phi_i)$ at arbitrary $\phi_i \in \mathbb{R}$, i.e., also outside of the discrete training points. Thereafter, an RC approach is presented to generate the data set at constant velocity in the setting in Fig. 5, i.e., including the non-linearity. Stability conditions are provided for both design steps.

4.1. Torque ripple compensation through GP regression

For compensation at arbitrary velocity, the compensation function $\bar{d}(\phi)$ must be a continuous function, i.e., known for all ϕ. This section outlines how a continuous function is generated based on a set of discrete observations of the disturbance through GP regression.

Suppose that a data set $\mathcal{D}_N(\phi_N, \hat{d}_N)$ is available where $\phi_N \in \mathbb{R}^N$ and $\hat{d}_N \in \mathbb{R}^N$ contain $N \in \mathbb{N}$ noisy data points $(\hat{d}(\phi(k)), \phi(k))$ that give information about the true spatial disturbance $\bar{d}(\phi)$, i.e.,

$$\hat{d}_N(\phi_N) = \bar{d}(\phi_N) + \epsilon_N, \tag{8}$$

and assume that $\epsilon_N \sim \mathcal{N}\left(0_N, \sigma_n^2 I_N\right)$ is additive independent identically distributed (i.i.d.) Gaussian noise. The data set $\mathcal{D}_N$ is generated in a separate RC experiment at constant velocity as outlined in Section 4.4. The aim is to learn a continuous and periodic function based on $\mathcal{D}_N$ that is also defined in-between the positions ϕ_N.

Many approaches exist to interpolate between discrete data points by assuming an underlying structure or basis functions, e.g., linear interpolation or least-squares regression. In this paper, $\bar{f}_d$ is obtained using GP regression, which does not require to specifically define basis functions. GP regression combines the data $\mathcal{D}_N$ with prior knowledge, e.g., smoothness and periodicity, to determine a posterior distribution over functions [19,20]. The mean of the posterior distribution, i.e., the most likely sample, is the continuous function $\bar{f}_d(\phi)$ that is used for compensation.

Suppose that the spatial disturbance function is a Gaussian process

$$\bar{d}(\phi) \sim \mathcal{GP}\left(\mu(\phi), \kappa(\phi, \phi')\right), \tag{9}$$

which is completely characterized by the mean function $\mu(\phi)$ that is assumed to be zero, and the covariance function $\kappa(\phi, \phi')$. The design of $\kappa(\phi, \phi')$ is outlined in Section 4.2. By conditioning the GP, i.e., performing inference on the basis of $\mathcal{D}_N$ and the prior (9), a posterior distribution is obtained. The posterior distribution at any position $\phi_* \in \mathbb{R}$ can be computed by defining the correlation between ϕ^* and the data $\mathcal{D}_N$ which is generated in a separate experiment, i.e.,

$$\begin{bmatrix} \hat{d}_N \\ \bar{d}(\phi_*) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} K + \sigma_n^2 I_N & K_* \\ K_*^\top & K_{**} \end{bmatrix} \right), \tag{10}$$

where $K \in \mathbb{R}^{N \times N}$ is the covariance function $\kappa(\phi_N, \phi_N)$ evaluated at all combinations of (ϕ_N, ϕ_N), and similar for $K_* = \kappa(\phi_N, \phi_*) \in \mathbb{R}^{N \times 1}$ and $K_{**} = \kappa(\phi_*, \phi_*) \in \mathbb{R}$ [19]. As a result, the spatial disturbance model $\bar{f}_d(\phi_*)$ is given by the mean of the posterior distribution

$$\bar{f}_d(\phi_*) = \kappa(\phi_*, \phi_N)\alpha \tag{11}$$

$$\alpha = \left(K(\phi_N, \phi_N) + \sigma_n^2 I_N \right)^{-1} \hat{d}_N \in \mathbb{R}^N, \tag{12}$$

for the full posterior distribution, see, e.g., [10,19]. Note that the vector α is completely determined by the data $\mathcal{D}_N(\phi_N, \hat{d}_N)$, thus the implementation of (11) can be done efficiently in step 2 since the inversion in (12) can be computed a priori in step 1.

The compensation signal $d_t(k)$ is obtained by evaluating the GP-based spatial disturbance model (11) at the current position $\phi_* = \phi(k)$, i.e.

$$d_t(k) = \bar{f}_d(\phi(k)) = \kappa\left(\phi(k), \phi_N\right)\alpha. \tag{13}$$

The following section presents the design of a suitable covariance function κ.
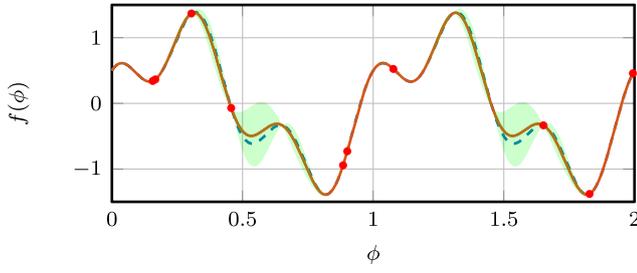
Fig. 11. Example 1: Gaussian process regression is performed with 10 random samples (⊙)[mred] denoted by $(f(\phi), \phi)$ taken from the true periodic function $\bar{d}(\phi)$ (- -). The resulting GP $f(\phi)$ has posterior mean (—) and posterior variance (▧), which is a good estimate of $\bar{d}(\phi)$ also in-between training points by selecting appropriate smooth and periodic prior knowledge.
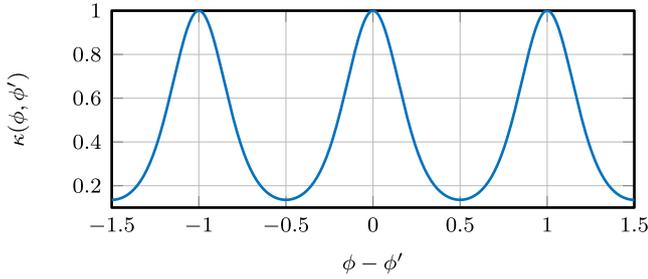
Fig. 12. Example 1: periodic kernel function with period $T = 1$, smoothness $l = 1$ and noise variance gain $\sigma_f^2 = 1$.

4.2. GP prior design for torque ripple disturbances

Since the spatial non-linearity is periodic as in Definition 1, the prior on $\bar{d}(\phi)$ must reflect this periodicity. Therefore, the covariance function in (9) is a periodic covariance function

$$\kappa(\phi, \phi') = \sigma_f^2 \exp\left(\frac{-2\sin^2\left(\pi(\phi - \phi')/T\right)}{l^2}\right), \tag{14}$$

with hyperparameters $\Theta_i = \{T, l, \sigma_f\}$. The hyperparameters can be interpreted as follows,

- $T \in \mathbb{R}$ is the spatial period of $\bar{d}(\phi)$ equal to the tooth pitch β;
- $l \in \mathbb{R}$ is the smoothness; and
- $\sigma_f \in \mathbb{R}$ is a gain relative to the noise variance σ_n^2.

Now consider the following illustrating example.

Example 1. A periodic function $\bar{d}(\phi) = \sin(2\pi\phi) + \frac{1}{2}\sin(6\pi\phi + \frac{1}{2}\pi)$ with period 1 is used as true function, from which 10 random samples $(\bar{d}(\phi_i), \phi_i)$ for $i = 1, 2, \ldots, 10$ are taken, see Fig. 11. A periodic covariance function is used for GP regression, with the hyperparameters: $T = 1$ equal to the period, $\sigma_f = 1$ and $l = 1$ as in Fig. 12. This shows that the covariance $\kappa(\phi, \phi')$ is large when $\phi - \phi'$ is either small (smoothness) or approximately equal to T (periodicity). GP regression is performed with the 10 data samples and periodic prior. The posterior mean and variance are shown in Fig. 11 together with the true function. The posterior mean, i.e., the most likely sample from the posterior, is a continuous function that corresponds well with the true function also in-between the data points. Moreover, the variance increases where the estimate is less reliable.

Example 1 shows that with smooth and periodic prior in GP regression, an accurate estimate of the true periodic function can be obtained based on a limited number of discrete data points. This is used to interpolate and extrapolate the time-domain disturbance model for compensation at any velocity.
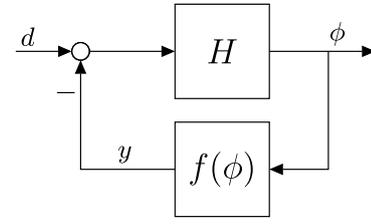
Fig. 13. Standard feedback interconnection with the linear system H and a non-linear memoryless function $f(\phi)$.

Remark 3. The periodic kernel (14) is also used in GP-based RC to learn a time-domain periodic function [9], or spatially periodic function [11].

4.3. Stability

To analyze stability of the interconnection in Fig. 4 with compensation function $\bar{f}_d$, it is reformulated in the standard feedback interconnection of a linear system in feedback with a static non-linearity as in Fig. 13. Consequently, a passivity-based stability analysis is used to provide stability conditions.

Lemma 1 (*Circle Criterion*). *Consider the standard feedback interconnection in* Fig. 13 *with the linear system* $H \in \mathcal{RH}_\infty^{p \times p}$ *and the memoryless function* $f(\phi)$ *as in [21, Theorem 7.2] and [22, p. 140]. If* H *is asymptotically stable and* $f(\phi)$ *satisfies the sector bound condition*

$$\gamma_1|\phi| \le f(|\phi|) \le \gamma_2|\phi| \tag{15}$$

for some $\gamma_1, \gamma_2 \in \mathbb{R}$ *with* $\gamma_1 < 0 < \gamma_2$, *then the interconnection is uniformly asymptotically stable, i.e., the output converges asymptotically to a stable steady-state trajectory, for all* $f(\phi)$ *that satisfy (15) if the Nyquist plot of* H *lies in the interior of the circular disk* $D(\gamma_1, \gamma_2)$, *with* $1/\gamma_1$ *and* $1/\gamma_2$ *the intersections of* $D(\gamma_1, \gamma_2)$ *with the real axis.*

The proof of Lemma 1 can be found in [21, Theorem 7.2, Definition 6.2], and [22, p. 141].

Theorem 1. *Consider the setting in* Fig. 4 *with linear system* P, *static unknown non-linearity* $\bar{d}(\phi)$, *non-linear function* $\bar{f}_d(\phi)$ *in (11), and suppose that* Assumption 1 *is satisfied. If*

$$\|PS(e^{j\omega})\|_\infty < \frac{1}{\delta}. \tag{16}$$

with $\delta \in \mathbb{R}$, *and* $\bar{f}_d(\phi)$ *is sector bounded,*

$$-\rho|\phi| \le \bar{f}_d(\phi) \le \rho|\phi| \quad \forall \quad \phi \in \left[-\frac{1}{2}\beta, \frac{1}{2}\beta\right), \tag{17}$$

where ϕ *can be shifted by* $n\beta$ *where* $n \in \mathbb{Z}$, *and for* $\rho \ge 0$, *such that*

$$\rho + \gamma \le \delta \tag{18}$$

with $\gamma = -\gamma_1 = \gamma_2 \ge 0$ *in* Assumption 1, *then the interconnection is asymptotically stable.*

The proof of Theorem 1 is given in the appendix.

Remark 4. In Lemma 1 with $\lambda_1 = -\lambda_2$, the condition (17) in Theorem 1 is similar to a small gain condition [21, Theorem 5.6]. For $\lambda_1 \ne -\lambda_2$ more general results can be obtained by following the same reasoning as in Theorem 1.

Remark 5. The sector-bound condition (A.2) in Theorem 1 requires that $\bar{f}_d(0) = 0$. If $\bar{f}_d(0) \ne 0$ then it is of the form $\bar{f}_d(\phi) = \bar{f}_d^0(\phi) + c$ for some constant $c \in \mathbb{R}$ such that $\bar{f}_d^0(0) = 0$. In this case replace $\bar{f}_d(\phi)$ by $\bar{f}_d^0(\phi)$ in (A.2) and the results in Theorem 1 remain valid, as c can be interpreted as a constant disturbance.
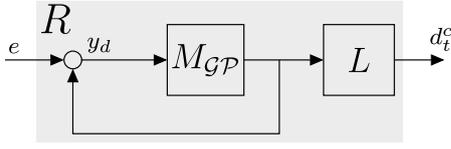
Fig. 14. Implementation of the repetitive controller R with Gaussian-process buffer $M_{\mathcal{GP}}$ and learning filter L.

Fig. 15. Modifying sensitivity function S_R with R designed for the CPA with rotational velocity 0.14 rad/s, such that the disturbance fundamental period equals $T = \frac{2\pi}{0.14 \cdot 131} = 0.343$ s. The periodic kernel has equal period, such that R suppresses disturbances at $\frac{1}{0.343} \approx 2.9$ Hz and its harmonics.

Theorem 1 implies that the reciprocal of the $\mathcal{H}_\infty$-norm of PS imposes the sector bound condition (A.2) and thereby an upper bound on the slope of $\bar{f}_d(\phi)$ around the origin. Moreover, the conditions in Theorem 1 can be checked on the basis of measurement data and the pre-designed function $\bar{f}_d(\phi)$ as shown in the following application on the CPA.

4.4. Repetitive control for constant velocity

The data set $\mathcal{D}_N$ that is used to train the GP in Section 4.1 contains observations of the disturbance which are obtained through an a priori performed RC experiment at constant velocity. The RC approach in [9] with a GP-based buffer is used, as it can easily be tuned to be robust for small period variations and does not impose restrictions on the buffer size (i.e. it does not require the periodicity to be an integer amount of samples); moreover, even in the presence of modeling errors, RC can asymptotically learn the periodic part of the error.

The repetitive controller R in the setting in Fig. 5 is shown in Fig. 14 and is given by

$$R = \frac{LM_{\mathcal{GP}}}{1 - M_{\mathcal{GP}}}, \tag{19}$$

where $M_{\mathcal{GP}} = \mu^{\mathcal{GP}} \Gamma$ is a GP-based buffer that is a FIR filter, Γ is a delay line of size N, and $\mu_{\mathcal{GP}} \in \mathbb{R}^{1 \times N}$ are coefficients that are computed as follows

$$\mu_{\mathcal{GP}} = \kappa(X_*, X) \left(\kappa(X, X') + \sigma_n^2 I_N \right)^{-1}, \tag{20}$$

with $X = T_s \cdot [N \; N-1 \; \dots \; 1]^\top \in \mathbb{R}^N$, $X^* = N \cdot T_s \in \mathbb{R}$ and T_s is the sample time. The same periodic kernel function (14) is used, where T is the expected time-domain period.

To incorporate period-time robustness, e.g., if the velocity varies slightly, a larger buffer size N can be used in combination with a locally-periodic kernel, for details, see [9]. Moreover, L is a stable learning filter to ensure stability.

Remark 6. Conceptually, alternative RC approaches such as [15] that are robust to period-time variations can also be used for this step, however, the presented GP-based approach yields more design freedom. Moreover, the buffer size N in the presented approach does not have to be exactly equal to the number of samples in a single period, as the period is provided in the kernel function and can be any real number. Furthermore, the presented approach does not require an explicit robustness filter Q, since robustness for e.g. modeling errors in L is covered by tuning the smoothness l in the kernel (14), as explained in [9].

Theorem 2. *The system in Fig. 5 with repetitive controller R in Fig. 14 and unknown non-linearity $\bar{d}(\phi)$ that is sector bounded as in Assumption 1 is asymptotically stable if the following conditions hold,*

$$|\mu_{\mathcal{GP}} \Gamma(e^{j\omega})(1 - S(e^{j\omega})P(e^{j\omega})L(e^{j\omega}))| < 1 \tag{21}$$

$$|P(e^{j\omega})S(e^{j\omega})S_R(e^{j\omega})| < \frac{1}{\gamma}, \tag{22}$$

for all $\omega \in [0, \pi]$, with $S_\text{R} = (I + PSR)^{-1}$ the modifying sensitivity.

The proof of Theorem 2 is given in the appendix.

The condition in Theorem 2 can also be checked a priori on the basis of a measured frequency response function of P. Moreover, by designing the learning filter L in (19) as

$$L \approx (PS)^{-1}, \tag{23}$$

the stability condition in (21) is satisfied. Many approaches exist to invert the process sensitivity PS, e.g., using stable inversion or zero-phase-error-tracking-control (ZPETC). Note that this choice of L is naturally in agreement with (7), i.e. its output d_t^c directly approximates the periodic disturbance that needs to be attenuated by the compensation $\bar{f}_d$. As such the learning filter L, being an approximate inverse of Fig. 7, incorporates the dynamic characteristics of the closed loop, including the dynamics of the CPA itself, in the presented approach.

Remark 7. If PS contains non-minimum phase zeros or delays, then the inverse can be non-causal, i.e., in the case of ZPETC $L = z^{-n_l} L_c$ with $n_l \in \mathbb{N}$ the number of samples preview and L_c the causal part of L. If L is non-causal then GP-based RC is also applicable, for implementation details see [9].

To show the attenuation properties of R in the setting in Fig. 14, the disturbance induced error e_d can be written as

$$e_d = -PSS_\text{R}d, \tag{24}$$

where S_R represents the closed-loop suppression of R in addition to the feedback controller. Consider the following example where S_R is shown for the CPA.

Example 2. A repetitive controller R is designed for the CPA that rotates with a constant velocity. The periodic kernel (14) is used with the period T seconds that corresponds with the time-domain period of the error, e.g., for 0.14 [rad/s] the period equals $T = \frac{2\pi}{0.14 \cdot 131} = 0.343$.

The resulting modifying sensitivity is shown in Fig. 15. This shows that R has high disturbance attenuation properties at $\frac{1}{T} \approx 2.9$ Hz and its harmonics corresponding with the period of the torque ripple at this specific velocity, in-between these frequencies the repetitive controller can slightly amplify disturbances.

Example 2 shows that indeed S_R is small at frequencies where the disturbance is expected and corresponds with the kernel period T. This is used to asymptotically reject the disturbance d for constant rotational velocity, even if the model $\widehat{PS}$ is non-perfect, resulting in the signal d_t^c that equals d, i.e., a time-domain model of the disturbance. Together with ϕ during the same experiment, this gives the data set $\mathcal{D}_N$ for GP regression.

4.5. Procedure

The following procedure summarizes implementation of the two-step approach that is introduced in this paper.
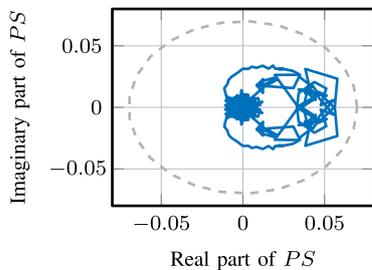
Fig. 16. Nyquist plot of the process sensitivity PS that is enclosed in a circle $D(-\delta, \delta)$ around the origin with radius $1/\delta = 0.07$.

Fig. 17. Cumulative amplitude spectrum (CAS) of the error with three velocities 0.08 rad/s (——), 0.1 rad/s (——), and 0.12 rad/s(——), for feedback (dotted), GP RC (solid) and the spatial GP compensation (dashed).

Procedure 1

Step 1) Constant velocity learning:

    (a) Identify a model $\widehat{PS}$ of PS and obtain L in (23).
    (b) Compute $\mu_{\mathcal{GP}}$ with Eq. (20), with the period in κ equal to $T = \frac{2\pi}{v \cdot 131}$ for the CPA, where v is the velocity in rad/s.
    (c) Verify stability with Theorem 2.
    (d) Perform an RC experiment at constant velocity v and collect data $\mathcal{D}_N$.

Step 2) GP regression and compensation with $\bar{f}_d$:

    (a) Compute α with (12) and data $\mathcal{D}_N(\phi_N, \hat{d}_N)$ obtained with step 1d, and a set of hyperparameters where T is the spatial period, l a given smoothness and σ_f the variation around the mean.
    (b) Verify stability with Theorem 1.
    (c) Apply (13) online, i.e., at each sample k, to compensate the disturbance for any given $\phi(k)$.

Remark 8. Multiple RC experiments at different velocities can be combined in $\mathcal{D}_N$ to compute α in (12). This allows to average over non-repeating disturbances that are present in $\mathcal{D}_N$, i.e., similar to employing a learning gain in traditional RC [23]. However, taking more data can be time consuming, hence to demonstrate that our proposed method does not require such averaging techniques, only a single velocity is used to experimentally obtain $\mathcal{D}_N$ in the next section.

5. Application to the CPA

In this section, the experimental results on the CPA are presented. First, the results of applying GP-based RC to the CPA with constant velocity are shown. Second, the GP-based spatial compensation function $\bar{f}_d$ that is computed on the basis of the time-domain RC model is shown. Third, the resulting performance for other constant and varying velocities are presented.

5.1. RC results at constant velocity

The GP-based repetitive controller R is designed for 3 different velocities to, (i) experimentally validate that it rejects the disturbance for constant velocities, and (ii) generate training data for GP regression. Each RC is designed with the periodic kernel function (14) where the period T is modified correspondingly as in Step 2 of Procedure 4.5, and the buffer size N is such that one period easily fits in the buffer.

The resulting cumulative amplitude spectra (CAS) are shown in Fig. 17 as a function of the spatial frequency in tooth^{-1}. The dotted lines show the current state-of-the-art performance with the existing commutation and a simple feedback controller; the solid lines show the obtained performance after convergence of the repetitive controllers. This shows that the repetitive controller R removes the repeating parts of the error that correspond with 1 tooth^{-1} and its harmonics for constant velocity. Note that the existing commutation already performs quite well, with errors of only 30 μrad RMS, but the GP-RC still manages to improve this with about a factor 3.
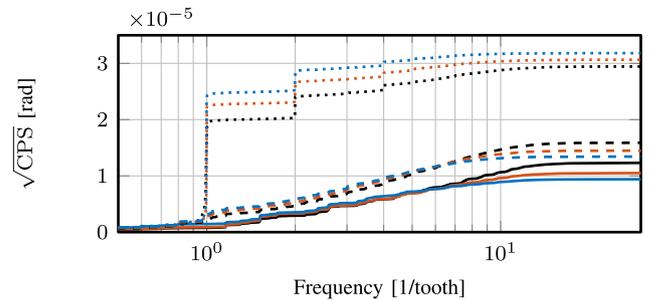
5.2. Learned GP compensation and stability

From the GP-based RC experiment at a single velocity of 0.1 rad/s the data set $\mathcal{D}_N$ is generated, where $\hat{d} = d_t^c$ contains 279 data points. Note that this velocity is considered to be around the upperlimit of what can be expected in the CPA application, and can thus be interpreted as a worst case. A GP-fit is made based on these data points with the hyperparameters: smoothness $l = 0.4$, the period equal to the tooth pitch $T = \frac{2\pi}{131}$, $\sigma_f^2 = 0.001$ the approximate amplitude around the mean, and the estimated noise variance $\sigma_v^2 = 6 \cdot 10^{-4}$ such that the confidence bound corresponds with the variance on the data points. The resulting GP-based compensation function that is periodic with a single tooth pitch β is shown in Fig. 18, including the 99.7% confidence bound.

Stability is assessed with condition (18) in Theorem 1. First, δ in (16) is determined in Fig. 16 yielding that $1/\delta = 0.07$ and hence $\delta \approx 14.3$. Second, Fig. 10 shows that $-\gamma_1 = \gamma_2 = \gamma = 1$ in Assumption 1 is a valid assumption for the unknown non-linear function $\bar{d}(\phi)$. Third, the GP-based compensation function $\bar{f}_d$ shown in Fig. 18 satisfies the sector bound condition in (17) for $\rho = 0.5$. As a result, $\rho + \gamma = 1.5 \ll \delta$ which implies that (18) in Theorem 1 is satisfied and the system is closed-loop stable.

Remark 9. The RC output d_t^c contains noise that is potentially amplified by the learning filter L at high frequencies. Although the GP regression is not sensitive to this noise due to the smoothness parameter l, for convenience and visualization reasons (i.e. to keep the spread on the red dots in Fig. 18 limited) the signal d_t^c is first pre-processed using an anti-causal zero-phase low-pass filter with a cut-off frequency of 1000 Hz (close to the Nyquist frequency) to remove some of the noise before GP regression.

5.3. Results

The learned spatial compensation function is employed at different velocities to compensate for the repeating part of the error. The dashed lines in Fig. 17 show the performance with the GP-based compensation, i.e. with a fixed $\bar{f}_d(\phi)$ in the implementation of Fig. 4. This clearly shows that the periodic component is almost completely suppressed as desired, although the performance is slightly worse compared to RC for a specific rotational velocity. The RMS errors for each velocity
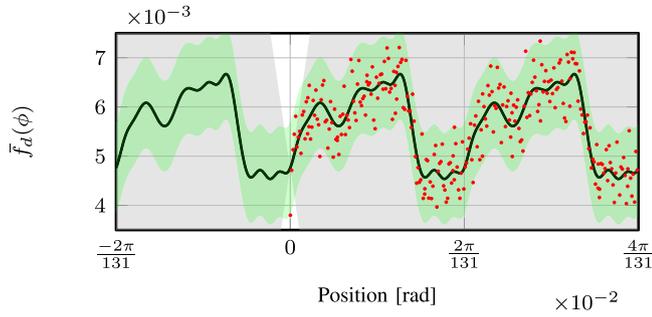
Fig. 18. Learned GP-based compensation function $\bar{f}_d$ with 99.7% confidence bound (■), which is trained with data that is obtained with the RC experiment at 0.1 rad/s. The sector bound condition (17) with $\rho = 0.5$ is shown in (■).
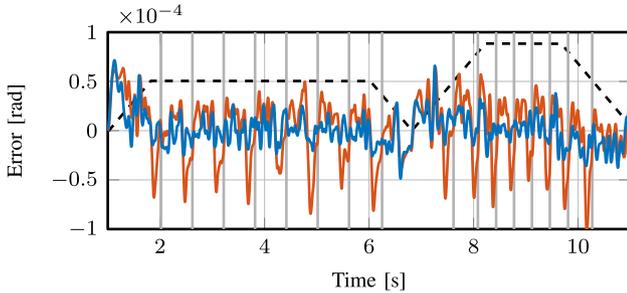
Fig. 19. Experiment with a varying velocity in (– –) (scaled), for the state-of-the-art feedback controller (——) and with the learned GP-based compensation function in (——).

Table 1
Root Mean Squared (RMS) error in micro-radians and the factor performance improvement with respect to feedback control for the three constant velocities.

Velocity [rad/s]	Feedback [μrad]	GP RC [μrad]	GP FB [μrad]
0.08	29.387	12.412 (2.4)	15.864 (1.9)
0.10	30.628	10.568 (2.9)	14.629 (2.1)
0.12	31.852	9.455 (3.4)	13.547 (2.4)
Varying	32.633	n/a	20.021 (1.6)

and control strategy are summarized in Table 1, yielding performance improvement up to a factor 2.4 with $\bar{f}_d$.

In addition, a measurement is carried out where the CPA rotates with a varying velocity for which RC is not applicable. The result is shown in Fig. 19, where (——) shows the error with feedback control and (——) is the error with the learned spatial compensation. This shows that independently of the velocity, the periodic component in the error that repeats with each tooth is significantly attenuated. The overall RMS value of the error is improved by 60%. Moreover, a substantial part of the remaining error can be attributed to reference-induced errors during acceleration and deceleration which is not compensated for by the present approach.

6. Conclusion

Automatic learning of a static compensation function for torque ripple disturbances at arbitrary rotational velocity is successfully derived and demonstrated in this paper. By employing a single repetitive control experiment at constant velocity, a time-domain model of the disturbance is obtained. Consequently, a Gaussian process is used to learn, on the basis of discrete time-domain data, a spatially periodic and continuous function. The learned function can be used for compensation at arbitrary rotational velocities where the torque ripple is not periodic. Moreover, stability checks are provided using the circle criterion, taking into account the non-linearity induced by the obtained

commutation update function. The method is experimentally validated on the coarse pointing assembly (CPA), resulting in a user-friendly automated routine for an average performance improvement of a factor 2 on the CPA at varying velocity, which could prove to be enabling for FSOC applications. This approach can thus successfully be employed during integration of future CPA production series, saving labor-intense testing and calibration steps, and thereby contributing to the cost-effectiveness of the CPA and its applicability in commercial space applications, such as FSOC. Future research focuses on direct learning or updating of the individual commutation functions on the basis of data.

CRediT authorship contribution statement

Noud Mooren: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Visualization. **Max van Meer:** Software, Investigation, Writing – review & editing. **Gert Witvoet:** Conceptualization, Resources, Writing – original draft, Writing – review & editing, Supervision. **Tom Oomen:** Conceptualization, Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix. Proofs

Proof of Theorem 1. The interconnection in Fig. 4 is reformulated in the standard feedback interconnection in Fig. 13 with

$$f(\phi) = \bar{f}_d(\phi) - \bar{d}(\phi)$$

and $H = PS$. Subsequently, Lemma 1 is used to provide stability conditions. First, the process sensitivity PS is stable by design of C_{fb}. Second, condition (16) implies that the Nyquist plot of PS resides in a disk with radius $1/\delta$ around the origin, denoted by $D(-\delta, \delta)$. It remains to show, using the sector bound condition (15) in Lemma 1, that the interconnection is stable if $\rho + \gamma \le \delta$. From the triangular inequality is follows that $|\bar{f}_d(\phi) - \bar{d}(\phi)| \le |\bar{f}_d(\phi)| + |\bar{d}(\phi)|$. Using Lemma 1, the interconnection is globally asymptotically stable if

$$|\bar{f}_d(\phi)| + |\bar{d}(\phi)| \le \delta|\phi|. \tag{A.1}$$

Now use Assumption 1, i.e., $|\bar{d}(\phi)| < \gamma|\phi|$, to rewrite (A.1) into

$$|\bar{f}_d(\phi)| \le (\delta - \gamma)|\phi| = \rho|\phi|, \tag{A.2}$$

or equivalently $-\rho|\phi| \le \bar{f}_d(\phi) \le \rho|\phi|$ for $\rho \le \delta - \gamma$ which corresponds with (18). Finally, as $\bar{f}_d(\phi)$ is periodic with period β, see Definition 1, it is sufficient to check the sector bound condition (A.2) on the interval $\phi \in \left[-\frac{1}{2}\beta, \frac{1}{2}\beta\right)$ which completes the proof.

Proof of Theorem 2. The proof consists of two part, first, (21) is derived in [9, Theorem 2] for stability of the linear part in Fig. 5, i.e., for $\bar{d}(\phi) = 0$, second, it remains to show that (22) implies stability with $\bar{d}(\phi) \neq 0$ under Assumption 1 with Lemma 1. The setting in Fig. 5 is reformulated in the standard feedback interconnection in Fig. 13, yielding that

$$H = P(1 + P(C + R))^{-1} = PSS_{\mathrm{R}},$$

and $f(\phi) = \bar{d}(\phi)$. Using Lemma 1 and the sector bound condition in Assumption 1 the interconnection is stable if the Nyquist plot of H resides inside the disk $D(\gamma, -\gamma)$. This is equivalent to (22) which completes the proof.

References

[1] Kramer L, Peters J, Voorhoeve R, Witvoet G, Kuiper S. Novel motorization axis for a coarse pointing assembly in optical communication systems. IFAC-PapersOnLine 2020;53(2):8426–31.

[2] Miller TJE. Electronic control of switched reluctance machines. Elsevier; 2001.

[3] Vujičić VP. Minimization of torque ripple and copper losses in switched reluctance drive. IEEE Trans Power Electron 2011;27(1):388–99.

[4] Rohrig C, Jochheim A. Identification and compensation of force ripple in linear permanent magnet motors. In: Proceedings of the 2001 american control conference.(Cat. No. 01CH37148), Vol. 3. IEEE; 2001, p. 2161–6.

[5] Ai W, Wu M, Li X, Li X. Active disturbance rejection based iterative learning control for direct torque control of switched reluctance motor drive. In: 2021 IEEE 10th data driven control and learning systems conference (DDCLS). IEEE; 2021, p. 1379–84.

[6] Longman RW. On the theory and design of linear repetitive control systems. Eur J Control 2010;16(5):447–96.

[7] Hara S, Yamamoto Y, Omata T, Nakano M. Repetitive control system: A new type servo system for periodic exogenous signals. IEEE Trans Automat Control 1988;33(7):659–68.

[8] Francis B, Wonham W. The internal model principle of control theory. Automatica 1976;12(5):457–65.

[9] Mooren N, Witvoet G, Oomen T. Gaussian process repetitive control: Beyond periodic internal models through kernels. Automatica 2022;140:110273.

[10] Mooren N, Witvoet G, Oomen T. Gaussian process repetitive control with application to an industrial substrate carrier system with spatial disturbances. IEEE Trans Control Syst Technol 2023;31(1):344–58.

[11] Mooren N, Witvoet G, Oomen T. Gaussian process repetitive control for suppressing spatial disturbances. IFAC-PapersOnLine 2020;53(2):1487–92.

[12] Sahoo SK, Panda SK, Xu JX. Application of spatial iterative learning control for direct torque control of switched reluctance motor drive. In: 2007 IEEE power engineering society general meeting. 2007, p. 1–7.

[13] Yao B, Xu L. Adaptive robust motion control of linear motors for precision manufacturing. Mechatronics 2002;12(4):595–616.

[14] Hu C, Yao B, Chen Z, Wang Q. Adaptive robust repetitive control of an industrial biaxial precision gantry for contouring tasks. IEEE Trans Control Syst Technol 2011;19(6):1559–68.

[15] Steinbuch M, Weiland S, Singh T. Design of noise and period-time robust high-order repetitive control, with application to optical storage. Automatica 2007;43(12):2086–95.

[16] Schmidt RM, Schitter G, Rankers A. The design of high performance mechatronics-: high-tech functionality by multidisciplinary system integration. Ios Press; 2020.

[17] van Zundert J, Oomen T. On inversion-based approaches for feedforward and ILC. Mechatronics 2018;50:282–91.

[18] Tomizuka M. Zero phase error tracking algorithm for digital control. J Dyn Syst Meas Control 1987;109(1):65–8.

[19] Williams C, Rasmussen C. Gaussian processes for machine learning, Vol. 2. MIT press Cambridge, MA; 2006.

[20] Murphy KP. Machine learning: A probabilistic perspective. MIT press Cambridge, MA; 2012.

[21] Khalil HK. Nonlinear systems. 3rd ed.. Prentice Hall; 2002.

[22] Desoer CA, Vidyasagar M. Feedback systems: input-output properties. SIAM; 2009.

[23] Tomizuka M. Dealing with periodic disturbances in controls of mechanical systems. Annu Rev Control 2008;32(2):193–9.

Noud Mooren received the B.Sc. degree from Fontys University of Applied Sciences, and the M.Sc. degree (cum laude) and Ph.D. degree (2022) from the Eindhoven University of Technology, Eindhoven, The Netherlands. He is currently a System Engineer at Sioux Technologies, The Netherlands. He is a recipient of the 2020 AMC Best Paper Award. His research interest is in the field of motion control and learning control techniques for applications in mechatronic systems.

Max van Meer received the M.Sc. degree (cum laude) in Mechanical Engineering from the Eindhoven University of Technology, Eindhoven, The Netherlands in 2021. He is currently a Ph.D. candidate at the Eindhoven University of Technology in the Department of Mechanical Engineering. His research interests include learning control, real-time systems and machine learning for motion control.

Gert Witvoet received the M.Sc. degree (cum laude) and Ph.D. degree from the Eindhoven University of Technology, Eindhoven, The Netherlands. He is a senior dynamics and control specialist at the Netherlands Organisation for Applied Scientific Research (TNO), Delft, The Netherlands, and a part-time associate professor with the Mechanical Engineering department at the Eindhoven University of Technology. He is a recipient of the Unilever Research Prize, and several best Master teacher awards. His research interests are in the application of advanced motion and learning control techniques on high-tech instruments and equipment, with applications in the semiconductor, astronomy, and space markets.

Tom Oomen received the M.Sc. degree (cum laude) and Ph.D. degree from the Eindhoven University of Technology, Eindhoven, The Netherlands. He is currently a professor with the Department of Mechanical Engineering at the Eindhoven University of Technology. He is also a part-time full professor with the Delft University of Technology. He held visiting positions at KTH, Stockholm, Sweden, and at The University of Newcastle, Australia. He is a recipient of the 7th Grand Nagamori Award, the Corus Young Talent Graduation Award, the IFAC 2019 TC 4.2 Mechatronics Young Research Award, the 2015 IEEE Transactions on Control Systems Technology Outstanding Paper Award, the 2017 IFAC Mechatronics Best Paper Award, the 2019 IEEJ Journal of Industry Applications Best Paper Award, and recipient of a Veni and Vidi personal grant. He is Associate Editor of the IEEE Control Systems Letters (L-CSS), IFAC Mechatronics, and IEEE Transactions on Control Systems Technology. He is a member of the Eindhoven Young Academy of Engineering. His research interests are in the field of data-driven modeling, learning, and control, with applications in precision mechatronics.