

Hierarchical Event-triggered Learning for Cyclically Excited Systems with Application to Wireless Sensor Networks*

Jonas Beuchert¹, Friedrich Solowjow², Jörg Raisch¹, Sebastian Trimpe², and Thomas Seel¹

Abstract—Communication load is a limiting factor in many real-time systems. Event-triggered state estimation and event-triggered learning methods reduce network communication by sending information only when it cannot be adequately predicted based on previously transmitted data. This paper proposes an event-triggered learning approach for nonlinear discrete-time systems with cyclic excitation. The method automatically recognizes cyclic patterns in data – even when they change repeatedly – and reduces communication load whenever the current data can be accurately predicted from previous cycles. Nonetheless, a bounded error between original and received signal is guaranteed. The cyclic excitation model, which is used for predictions, is updated hierarchically, i.e., a full model update is only performed if updating a small number of model parameters is not sufficient. A nonparametric statistical test enforces that model updates happen only if the cyclic excitation changed with high probability. The effectiveness of the proposed methods is demonstrated using the application example of wireless real-time pitch angle measurements of a human foot in a feedback-controlled neuroprosthesis. The experimental results show that communication load can be reduced by 70 % while the root-mean-square error between measured and received angle is less than 1°.

Index Terms—sensor networks, statistical learning

I. INTRODUCTION

MANY applications require real-time transmission of signals over communication channels with bandwidth limitations. A typical example is given by wireless sensor networks in feedback-controlled systems. The number of agents (i.e., network nodes) and their communication rate is limited by the amount of information the wireless network can transmit in real-time. It is, therefore, desirable to reduce the communication load without compromising the accuracy of the transmitted signals.

Well known approaches are event-based sampling [1]–[3] and event-triggered state estimation (ETSE [4]–[6], sometimes referred to as model-based event-based sampling [7]): At each sampling instant, the receiving agent independently predicts the state, which is measured by the sender, based on previous




Fig. 1. Event-triggered learning in a two-agent network. If the measured signal (e.g., a human foot angle) can be described by a cyclically excited system model, the model is shared with the receiver and the communication is reduced to those samples that cannot be predicted using the model and previously transmitted data. Inaccuracy of the model (e.g., due to change in walking pattern) is detected and a new model is identified and shared.

estimations and a model. The sender performs the identical prediction and communicates the measured state if and only if the error between prediction and measurement exceeds a predefined threshold. Otherwise, there is no communication and the receiving agent uses the model-based prediction as estimation (cf. Fig. 1).

Since the prediction accuracy heavily depends on the quality of the utilized model, it was recently proposed to learn and update models in an event-triggered fashion as well [8]. Occurrence of communication is treated as a random variable, and the model is updated when empirical data does not fit the probability distribution that would result if the model was the truth.

The present paper builds on the idea of [8] and develops event-triggered learning (ETL) methods for cyclically excited systems. The main contributions are:

- Extension of the concept of ETL to specifically target systems with a locally cyclic excitation. Locally means that cycles close in time are almost identical; however, cycles that are far apart are not necessarily similar. This class of systems is useful for describing both biological and technical processes (e.g., human motion, breath, heartbeat, and production cycles).
- Utilization of a learning trigger tailored to the problem at hand (one-sided Kolmogorov-Smirnov test [9]), which fires with high probability in case of a model change and with low probability otherwise.
- Introduction of the novel idea of a hierarchical model learning strategy, which updates and communicates only a reduced number of model parameters whenever that is sufficient.
- Demonstration of significant communication savings (70 %) using experimental data from cyclic human motion collected with a wearable inertial sensor network.

*The work of F. Solowjow and S. Trimpe was supported in part by the Max Planck Society, the IMPRS-IS, and the Cyber Valley Initiative.

¹Jonas Beuchert, Jörg Raisch, and Thomas Seel are with Control Systems Group, Technische Universität Berlin, 10587 Berlin, Germany, jonas.beuchert@campus.tu-berlin.de, seel@control.tu-berlin.de.

²Friedrich Solowjow and Sebastian Trimpe are with Intelligent Control Systems Group, Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany, solowjow@is.mpg.de, trimpe@is.mpg.de.

This is the first application of ETL on real-world network data.

The paper continues as follows. After defining the considered problem in Sec. II, the event-triggered learning architecture from [8] is briefly explained in Sec. III and then extended to cyclically excited systems. Subsequently, the properties of the method are validated experimentally in Sec. IV. Finally, Sec. V provides conclusions.

II. SYSTEM REPRESENTATION AND PROBLEM FORMULATION

Consider a discrete-time system with sample index $k \in \mathbb{Z}$ and state $x[k] \in \mathbb{R}^n$, which is measured by the sending agent. The system is assumed to be influenced by a cyclic input $u[k] \in \mathbb{R}^m$ with cycle length $N \in \mathbb{N}^+$ such that $u[k+N] = u[k]$ and by zero-mean noise $\varepsilon[k] \in \mathbb{R}^s$, which is distributed according to a time-invariant probability distribution \mathcal{P} . The recursive state update law of the system is characterized by the dynamics $f: \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^s \rightarrow \mathbb{R}^n$, i.e.,

$$x[k] = f(x[k-1], u[k], \varepsilon[k]). \quad (1)$$

While the dynamics f is assumed to be known, the excitation $u[k]$ and its cyclicity N are unknown and may change with time. In the following, whenever the term *model* is used, it refers to an approximation of the cyclic excitation $u[k]$. The distribution \mathcal{P} of the noise $\varepsilon[k]$ is known.

This paper considers architectures with one sending and one receiving agent. However, the methods can be directly applied to multi-agent systems and yield the same advantages therein.

The sending and receiving agents have the following capabilities, which will be made precise in the next section:

- the sender can transmit measured data samples to the receiver, i.e., perform *state updates*;
- sender and receiver can estimate current data samples from previously transmitted data and a model, i.e., perform *predictions*;
- the sender can estimate excitation trajectories from measured data, i.e., perform *model identification*;
- the sender can send model parameters to the receiver, i.e., perform *model updates*.

The main objective is to find a joint strategy for the sender and the receiver such that the amount of communication (state and model updates) is reduced while the error between the actual measurement signal $x[k]$ and the signal estimate $\hat{x}[k] \in \mathbb{R}^n$ on the receiver side remains small in the sense of a suitable metric $\|\cdot, \cdot\|: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}_0^+$.

III. ALGORITHM DESIGN

The proposed event-triggered learning approach for cyclically excited systems is illustrated in Fig. 2. It can be described by the following building blocks:

- Two identical predictors that estimate the measured state $x[k]$ based on an internal state $\hat{x}[k-1]$ and an estimated excitation trajectory $\hat{u} \in \mathbb{R}^{m \times \hat{N}}$ of one cycle with estimated cycle length $\hat{N} \in \mathbb{N}^+$.
- A binary state-update trigger $\gamma_{\text{state}}[k] \in \{0, 1\}$ that determines when to update the internal state $\hat{x}[k]$ of

the predictors with the measurement $x[k]$ to ensure a bounded error $\|x[k], \hat{x}[k]\|$ of the estimation.

- A model learning block that estimates the excitation trajectory \hat{u} used by the predictors or updates the estimated parameters $\vartheta \in \mathbb{R}^v$ (e.g., cycle length or amplitude) of the current trajectory \hat{u} .
- A binary learning trigger $\gamma_{\text{learn}}[k] \in \{0, 1\}$ that determines when to update the internal excitation trajectory model \hat{u} of the predictors. Ideally, learning shall be triggered if and only if the rate of state updates increases due to a false or inaccurate model \hat{u} .

A. Prediction

Let the current model of the cyclic excitation $u[k]$ be described by the aforementioned matrix \hat{u} such that the current excitation is the $(j[k])^{\text{th}}$ column of \hat{u} , i.e., $\hat{u}_{j[k]}$ is the estimated value of $u[k]$, where the index $j[k]$ obeys

$$j[k+1] = \begin{cases} j[k] + 1 & \text{if } \gamma_{\text{learn}}[k] = 0 \wedge j[k] < \hat{N} \\ 1 & \text{if } \gamma_{\text{learn}}[k] = 1 \vee j[k] = \hat{N} \end{cases}. \quad (2)$$

The index $j[k]$ is increased after every prediction step, unless the end of the estimated excitation trajectory is reached ($j[k] = \hat{N}$) or learning is triggered ($\gamma_{\text{learn}}[k] = 1$).

The estimate $\hat{x}[k]$ is determined by

$$\hat{x}[k] = \begin{cases} f(\hat{x}[k-1], \hat{u}_{j[k]}, 0) & \text{if } \gamma_{\text{state}}[k] = 0 \\ x[k] & \text{if } \gamma_{\text{state}}[k] = 1 \end{cases}. \quad (3)$$

B. State-update Trigger

If $d[k] = \|x[k], f(\hat{x}[k-1], \hat{u}_{j[k]}, 0)\|$ reaches or exceeds the predefined threshold δ , then a state update as defined in Sec. III-A is triggered

$$\gamma_{\text{state}}[k] = \begin{cases} 0 & \text{if } d[k] < \delta \\ 1 & \text{if } d[k] \geq \delta \end{cases}. \quad (4)$$

Triggers that compare the actual value to a model-based prediction are common in ETSE (cf. [1], [4]–[7], [10]–[12]). Because a state update ($\gamma_{\text{state}}[k] = 1$) leads to zero error, (3) and (4) together ensure that $\|x[k], \hat{x}[k]\|$ is bounded by δ .

C. Robust Learning Trigger

If the model (the estimated excitation trajectory \hat{u}) is exact, then state updates occur only due to noise $\varepsilon[k]$. If the model no longer yields valid predictions of the current data, the time durations between two consecutive state updates will decrease. The learning trigger aims at detecting this decrease and then triggering a model update.

Let the inter-communication time be the discrete time duration (number of samples) between two consecutive state updates and collect these times $\tau_1, \tau_2, \dots, \tau_i$ in a buffer, which is emptied whenever model learning is triggered ($\gamma_{\text{learn}}[k] = 1$).

Perform Monte Carlo (MC) simulations of (1), (3), and (4) for the case in which the model \hat{u} is perfect and state updates occur only due to noise. This yields a hypothetical

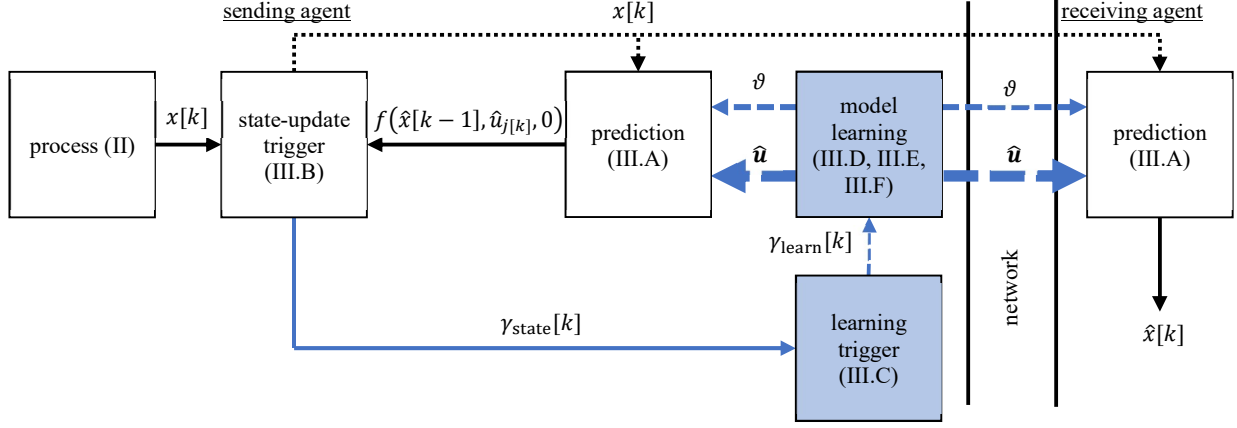


Fig. 2. One sending and one receiving agent with the typical event-triggered state estimation architecture in black and event-triggered learning in blue (extended from [8]). The process provides the measured state $x[k]$ at every sampling instant k . At the same time, the state is estimated by the prediction blocks of the sender and the receiver using the previous estimate $\hat{x}[k-1]$ and a trajectory model of the excitation \hat{u} . If the prediction differs significantly from the measured state, then a state update is triggered, and the internal states of both prediction blocks are set to the measured state. Too frequent state updates indicate poor model quality and, therefore, trigger model learning. This can either lead to only an adjustment of certain parameters ϑ of the current model trajectory or to a completely new excitation trajectory \hat{u} . In any case, the new model information is shared between sender and receiver.

cumulative distribution function $F: \mathbb{N}^+ \rightarrow [0, 1]$ of the inter-communication times. At each sample instant k , the aforementioned buffer provides an empirical cumulative distribution function $\hat{F}_k: \mathbb{N}^+ \rightarrow [0, 1]$ of the previously observed inter-communication times; i.e., for any value $\tau \in \mathbb{N}^+$, $\hat{F}_k(\tau)$ is the proportion of observed inter-communication times less than or equal to τ .

The hypothetical and the empirical distribution function of the inter-communication times are compared by the one-sided two-sample Kolmogorov-Smirnov test (KS-test) [9], [13], [14]. Its null hypothesis is that the empirical inter-communication times come from the hypothetical distribution $F(\tau)$. The alternative hypothesis is that they come from a different one. The result of the test is an estimated probability $p[k] \in [0, 1]$ that the null hypothesis is true. It is compared with a defined significance level $\eta \in [0, 1]$, i.e., the probability that the test rejects the null hypothesis although it is correct (type I error). If the $p[k]$ -value is smaller than η for a predefined minimum holding time duration $t_{\min} \in \mathbb{R}_0^+$, then a model update is triggered

$$\gamma_{\text{learn}}[k] = \begin{cases} 0 & \text{if } \exists l \in [k - t_{\min}, k] : p[l] \geq \eta \\ 1 & \text{if } p[l] < \eta \forall l \in [k - t_{\min}, k] \end{cases}. \quad (5)$$

Note that the null hypothesis is always accepted if no inter-communication times were observed since the last model update. Furthermore, for a large class of systems, the hypothetical distribution $F(\tau)$ will not depend on \hat{u} and, therefore, can be determined beforehand, i.e., no MC simulations must be performed in real time. Finally, for sufficiently simple systems, the distribution might even be determined analytically, and the one-sided one-sample KS-test can be used.

Utilizing the KS-test to design learning triggers was first proposed in [15]. Density-based learning triggers use richer statistical information and have more advantageous properties than learning triggers that are based on the expected value as proposed in [8]. In contrast to [15], the current method uses a one-sided trigger condition because model updates are

not required if the inter-communication times are larger than expected. The robustification with a minimum holding time t_{\min} prevents model learning due to unmodeled short-term effects. Furthermore, if a change of the process behavior occurs, then learning should not be triggered before the change is completed. For the theoretical properties of the statistical tests to hold, $\tau_1, \tau_2, \dots, \tau_i$ are assumed independent and identically distributed (see [8], [15]). While this is not necessarily the case for every system of the form (1), it holds, for example, for the application system considered in Sec. IV.

D. Model Learning – Small Model Update

The model learning block learns a new model trajectory \hat{u} based on the previous measurements $x[k], x[k-1], \dots$. However, transferring the complete trajectory \hat{u} of a cycle from the sending to the receiving agent leads to a significant amount of communication. Therefore, it is assumed that, in some cases, an appropriate new model trajectory $\hat{u}_+ \in \mathbb{R}^{m \times \hat{N}_+}$ with a new cycle length \hat{N}_+ can be derived from the old one by employing a parametric deformation function $g: \mathbb{R}^{m \times \hat{N}} \times \mathbb{R}^v \rightarrow \mathbb{R}^{m \times \hat{N}_+}$ with a small number of parameters $\vartheta \in \mathbb{R}^v$ such that $\hat{u}_+ = g(\hat{u}, \vartheta)$.

Whenever a learning update is triggered, the learning block determines the parameters ϑ that lead to the best approximation of the states measured during the last cycle:

$$e(k, \tilde{\vartheta}) = \sum_{l=k-\hat{N}_++1}^k \|\mathbf{x}[l], \tilde{\mathbf{x}}_{\tilde{\vartheta}}[l]\|^2, \quad (6)$$

$$E[k] = \min_{\tilde{\vartheta}} e(k, \tilde{\vartheta}), \quad \vartheta = \underset{\tilde{\vartheta}}{\operatorname{argmin}} e(k, \tilde{\vartheta}), \quad (7)$$

where $\tilde{\mathbf{x}}_{\tilde{\vartheta}}$ is the state obtained from simulating (1) with excitation $\hat{u}_+ = g(\hat{u}, \tilde{\vartheta})$ and with initial state $x[k - \hat{N}_+]$. This value provides an estimate of how good the model would fit after a small update.

Provided the small update is sufficient (see Sec. III-E), the parameters ϑ are transmitted to the receiver and, subsequently,

the receiver determines the new trajectory using the predefined function g and the previous model trajectory \hat{u} . For a cyclic process, useful parameters could be cycle length, phase shift, or amplitude, which can be estimated using standard signal processing methods in frequency or time domain to lower the computational costs in comparison to solving the optimization problem (6) explicitly (for an example see Sec. IV). Learning of these parameters is always carried out as first step of the block *model learning* in Fig. 2.

E. Learning-type Trigger

The small update described above requires much less communication than an update of the full trajectory \hat{u} , but might not always lead to sufficiently precise predictions. We define a binary trigger $\gamma_{\text{full}}[k] \in \{0, 1\}$ that indicates when a small update is not sufficient to achieve a satisfying performance improvement.

If the error $E[k]$ exceeds a threshold $\alpha \in \mathbb{R}_0^+$, then a full update is triggered

$$\gamma_{\text{full}}[k] = \begin{cases} 0 & \text{if } \gamma_{\text{learn}}[k] = 0 \vee E[k] \leq \alpha \\ 1 & \text{if } \gamma_{\text{learn}}[k] = 1 \wedge E[k] > \alpha \end{cases} \quad (8)$$

and a new full trajectory \hat{u} is identified and transmitted as detailed in Sec. III-F. Otherwise ($\gamma_{\text{full}}[k] = 0$), the sensor sends the parameters ϑ as small model update and both agents deform the previous excitation trajectory to obtain the new model $\hat{u} = \hat{u}_+$. The triggers of both model updates and the state update exhibit a hierarchical dependency: A small model update is only executed if the (communication-wise cheaper) state updates occur too frequently. Likewise, a (communication-wise expensive) full model update is only carried out if a small model update is expected to be insufficient.

F. Model Learning – Full Model Update

In case of a triggered full update, an appropriate new model \hat{u} cannot be obtained from the previous one by applying the deformation g . Therefore, the excitation trajectory for the previously measured cycle is estimated based on the available state measurements. We assume that the dynamics f allow for estimating $u[k]$ from a finite number of state measurements (as is the case in the considered application, see Sec. IV).

This trajectory could be transferred directly to the receiving agent via the network and used as a precise model by both prediction blocks for the following state estimations. However, sending all samples of the trajectory would require the network to transfer many values in a short time interval. Therefore, the sender compresses the trajectory at first with polynomial regression and the receiver performs reconstruction to obtain \hat{u} . Possible alternatives to polynomial regression for compression are, e.g., wavelet transformation with thresholding of the coefficients [16], and (sparse) Gaussian process regression [17], [18].

The entire approach described in Sec. III is summarized in Algorithm 1.

Algorithm 1 ETL for Cyclically Excited Systems (Sender)

```

if  $j < N$  then  $j \leftarrow j + 1$  else  $j \leftarrow 1$  ▷ Update index (2)
 $\hat{x} \leftarrow f(\hat{x}, \hat{u}_j, 0)$  ▷ Predict measurement (3)
if  $\|x, \hat{x}\| < \delta$  then ▷ Prediction error small (4)
     $t \leftarrow t + 1$  ▷ Increase inter-communication time
else ▷ Prediction error large (4)
     $\hat{x} \leftarrow x$  ▷ Update state with measurement (3)
     $\tau \leftarrow [\tau \quad t]$  ▷ Store inter-communication time
     $t \leftarrow 0$  ▷ Reset inter-communication time
if  $\text{KSTest}(\tau)$  then  $l \leftarrow 0$  else  $l \leftarrow l + 1$  ▷ KS-Test (5)
 $x \leftarrow [x \quad x]$  ▷ Store measurement
if  $l \geq t_{\min}$  then ▷ Minimum holding time reached (5)
     $j \leftarrow 0$  ▷ Reset model trajectory index (2)
     $\tau \leftarrow \emptyset$  ▷ Empty inter-communication time buffer
     $\vartheta \leftarrow \text{estimateParam}(x)$  ▷ Small update parameter (6), (7)
     $\hat{u} \leftarrow g(\hat{u}, \vartheta)$  ▷ Small model update (Sec. III-D)
    if  $\text{trajError}(\hat{u}, x) > \alpha$  then ▷ Small update poor (7), (8)
         $\hat{u} \leftarrow \text{identifyU}(x)$  ▷ Full model update (Sec. III-F)

```

IV. APPLICATION: PITCH ANGLE OF HUMAN FOOT

To demonstrate the effectiveness of the proposed algorithms, we consider a network of wearable sensor units comprising at least an inertial sensing chip, a wireless communication module, and a microcontroller, which is attached to a human body during gait. Such measurement systems are used for real-time biofeedback and control of robotic systems and neuroprostheses [19]. The rate at which the network can communicate reliably in real time is limited by the number of sensors. If the communication load between each sensor and the receiver can be reduced, higher base sampling rates or a larger number of sensors can be used. This challenge was recently addressed using heuristic approaches [20]. We now apply the proposed ETL methods to this problem.

We consider the specific example of real-time measurement of the foot pitch angle $x[k] \in \mathbb{R}$ in a feedback-controlled neuroprostheses [19] (cf. Fig. 1). The angle dynamics is modeled by cyclic increments $u[k] \in \mathbb{R}$ and zero-mean additive Gaussian noise $\varepsilon[k] \sim \mathcal{N}(0, \sigma^2)$ with standard deviation $\sigma \in \mathbb{R}_0^+$, i.e., the state update is

$$x[k] = x[k-1] + u[k] + \varepsilon[k]. \quad (9)$$

Equation (9) corresponds to (1) and is used for state predictions. For the trigger (4), the absolute difference $d[k] = |x[k] - \hat{x}[k]|$ is used as a metric.

Validation is based on measurements of approximately 30 minutes of variable human gait including normal walking with frequent speed and ground inclination changes as well as style changes to different kinds of simulated pathological gait (straightened knee, drop foot, and dragged leg). While walking style changes lead to clearly different angle trajectories, small gait velocity changes are well described by time warping of the angle trajectory, i.e., by cycle length changes of the periodic excitation $u[k]$. The small model update is, therefore, designed as follows: Determine the new cycle length $\vartheta_1 = \hat{N}$ and time shift ϑ_2 of \hat{u} that best describe the measurements of the most

recent cycle. Both parameters are estimated in the frequency domain (using the auto- or crosscovariance of the measured states of a small number of previous cycles) and further refined in the time domain with local optimization³.

For a full model update, (9) with $\varepsilon[k]$ being zero-mean implies that an unbiased estimate of $u[k]$ can be obtained by taking the difference of $x[k]$ and $x[k-1]$. We thus calculate $\Delta x[l] := x[l] - x[l-1] \forall l \in [k - \hat{N} + 1, k]$ and estimate the full trajectory as $\hat{u} = [\Delta x[k - \hat{N} + 1] \cdots \Delta x[k]]$, where k is the index of the sampling instant where learning is triggered.

A. Parametrization

The measurement sampling rate is set to 50 Hz. The accuracy of IMU-based foot pitch angle measurements is in the range of 3–4° [21]. Therefore, the maximum admissible angle error δ in (4) is set to 2°. The significance level for the KS-test (5) is set to $\eta = 5\%$. The accompanying minimum holding time is $t_{\min} = 0.35 \text{ s} \cdot 50 \text{ Hz}$, which is about a third of a stride duration. The initial MC simulation to obtain $F(\tau)$ is carried out with 1000 trials and a standard deviation $\sigma = 0.9^\circ$ of the noise. This value is the observed root-mean-square-error (RMSE) during normal walking. The RMSE-threshold of the learning-type trigger (8) is $\alpha = 5^\circ$. Polynomial regression is chosen for compression of \hat{u} . The degree is fixed to 18, which leads to a maximum of 20 communicated parameters including the constant offset and the estimated cycle length \hat{N} . In contrast, the full trajectory of a stride typically contains 40–80 samples if the rate is 50 Hz. More parameters would increase the risk of overfitting.

The MC simulation yields an expected inter-communication time $\hat{\mathbb{E}}[\tau] \approx 8$. This implies that a state update is expected to happen every 8th sample on average if the model is correct. Furthermore, because of $\eta = 5\%$, every 20th state update is expected to trigger an undesired model update, which leads to transmission of either two or 20 model parameters. In total, the parametrization should result in not more than $\frac{V_{\text{ETL}}}{V_{\text{full}}} \leq \frac{n+\eta w}{n \cdot \hat{\mathbb{E}}[\tau]} = \frac{1+0.05 \cdot 20}{8} = 25\%$ of transmitted values in comparison to full communication if the model is correct (i.e., whenever gait style and velocity do not change).

Finally, we initialize the model \hat{u} of the excitation trajectory with a vector containing zeros. The first estimated model is transmitted when the first full model update is triggered.

B. Results and Discussion

The proposed ETL method is compared to alternative strategies by determining the number of transmitted values and the RMSE between the measured angle and the output of the receiver. Results are visualized in Fig. 4. The trivial strategy of sending only every second sample (decim.) yields an RMSE of almost 2° at 50 % communication rate. The same reduction of communication is achieved by pure ETSE (i.e., the scheme in Fig. 2, in black without ETL); however, at an RMSE of less than 1°. ETL achieves a reduction of traffic to even less than




Fig. 4. Number of transferred values V with respect to full communication V_{full} (left) and resulting RMSE (right). We compare full communication (full), decimation with factor 2 (decim), event-triggered state estimation (ETSE), and two parametrizations of event-triggered learning (ETL_h / ETL_o) for about 0.5 h of highly variable gait. The results show that ETL reduces the communication significantly without changing the accuracy of the transferred signal in comparison to ETSE.

30 %, while the RMSE is still below 1°. Finally, optimizing the heuristically chosen parameters (ETL_h) with nested cross-validation (ETL_o) leads to small further improvements (27 %).

Fig. 5 shows how a small model update is triggered due to a change of the cycle length of the measured process, while Fig. 6 demonstrates how a change of the shape of the trajectory causes a full model update. Both figures also provide evidence that the absolute estimation error is bounded by $\delta = 2^\circ$. This represents a remarkable improvement with respect to the results in [20].

V. CONCLUSIONS

A set of methods for event-triggered learning for cyclically excited systems is proposed to minimize communication in sensor networks. The approach automatically recognizes cyclic patterns in data – even when they change repeatedly – and reduces communication load whenever the current data can be accurately predicted from previous cycles. In contrast to previous approaches, the current methods exploit explicitly the periodicity of the dynamics and account for time-variant behavior. Additional major advantages of the methods are that they assure an upper bound on the error of the received signal and that unnecessary model updates occur only with a low probability, which is controlled by a user-defined significance level.

The proposed methods are shown to yield significant resource savings in a wireless body sensor network for orientation tracking of a human foot. The experimental results show that a large reduction of communication load (by 70 %) and a small bounded estimation error (by 2°) can be achieved at the same time. This means that up to three times as many sensors could be used without jeopardizing latency of the real-time communication. Moreover, the examined sampling rate of 50 Hz is relatively low; up to 1 kHz is common in IMU networks (200 Hz in wireless ones), which gives even more potential for resource savings with the proposed methods.

Future research will focus on quaternion-based estimation of full body segment orientations and a learning trigger with less assumptions on inter-communication times.

³Code and detailed description on <http://www.control.tu-berlin.de/EventTriggeredLearning>

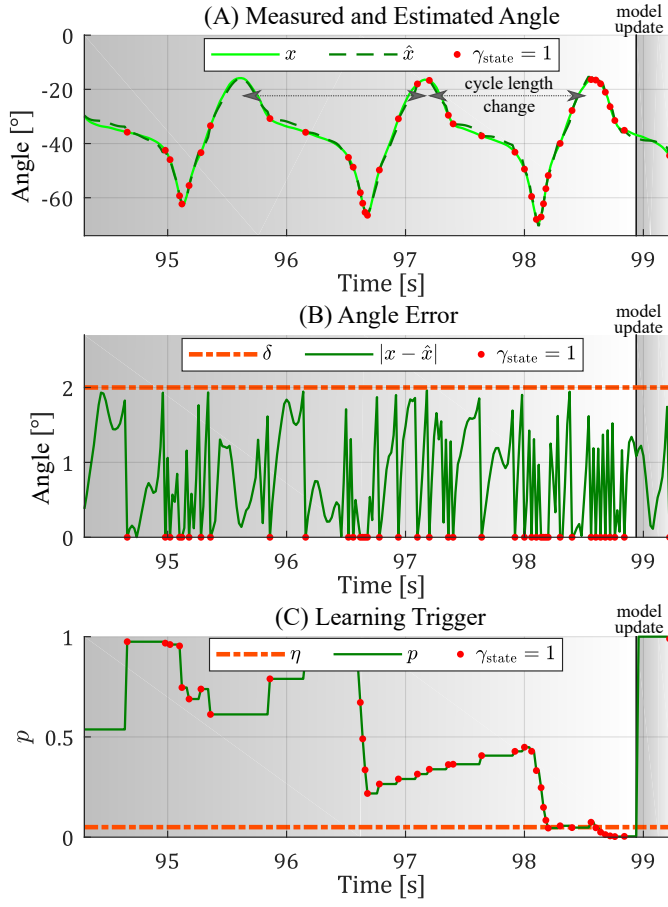


Fig. 5. Small model update ($\gamma_{\text{learn}}[k] = 1 \wedge \gamma_{\text{full}}[k] = 0$) triggered at 98.9 s due to a change of cycle length (gait velocity). (A) measured and estimated state; (B) error $|x[k] - \hat{x}[k]|$ and state-update trigger threshold δ ; (C) probability $p[k]$ of the KS-test (5) and significance level η . Model learning is triggered because too many state updates ($\gamma_{\text{state}}[k] = 1$) occur and $p[k]$ falls below η for the minimum holding time t_{\min} .

REFERENCES

- [1] D. Laidig, S. Trimpe, and T. Seel, "Event-based sampling for reducing communication load in realtime human motion analysis by wireless inertial sensor networks," *Current Directions in Biomedical Engineering*, vol. 2, no. 1, pp. 711–714, 2016.
- [2] M. Miskowicz, *Event-based control and signal processing*. CRC press, 2015.
- [3] M. Lemmon, "Event-triggered feedback in control, estimation, and optimization," in *Networked Control Systems*. Springer, 2010, pp. 293–358.
- [4] D. Shi, L. Shi, and T. Chen, "Event-based state estimation," *Switzerland: Springer*, 2016.
- [5] S. Trimpe and R. D'Andrea, "Event-based state estimation with variance-based triggering," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3266–3281, 2014.
- [6] J. Sijs and M. Lazar, "Event based state estimation with time synchronous updates," *IEEE Transactions on Automatic Control*, vol. 57, no. 10, pp. 2650–2655, 2012.
- [7] Q. Liu, Z. Wang, X. He, and D. Zhou, "A survey of event-based strategies on control and estimation," *Systems Science & Control Engineering: An Open Access Journal*, vol. 2, no. 1, pp. 90–97, 2014.
- [8] F. Solowjow, D. Baumann, J. Garcke, and S. Trimpe, "Event-triggered learning for resource-efficient networked control," in *American Control Conference (ACC)*, 2018, pp. 6506–6512.
- [9] F. J. Massey Jr, "The Kolmogorov-Smirnov test for goodness of fit," *Journal of the American Statistical Association*, vol. 46, no. 253, pp. 68–78, 1951.
- [10] M. Miskowicz, "Send-on-delta concept: An event-based data reporting strategy," *Sensors*, vol. 6, no. 1, pp. 49–63, 2006.

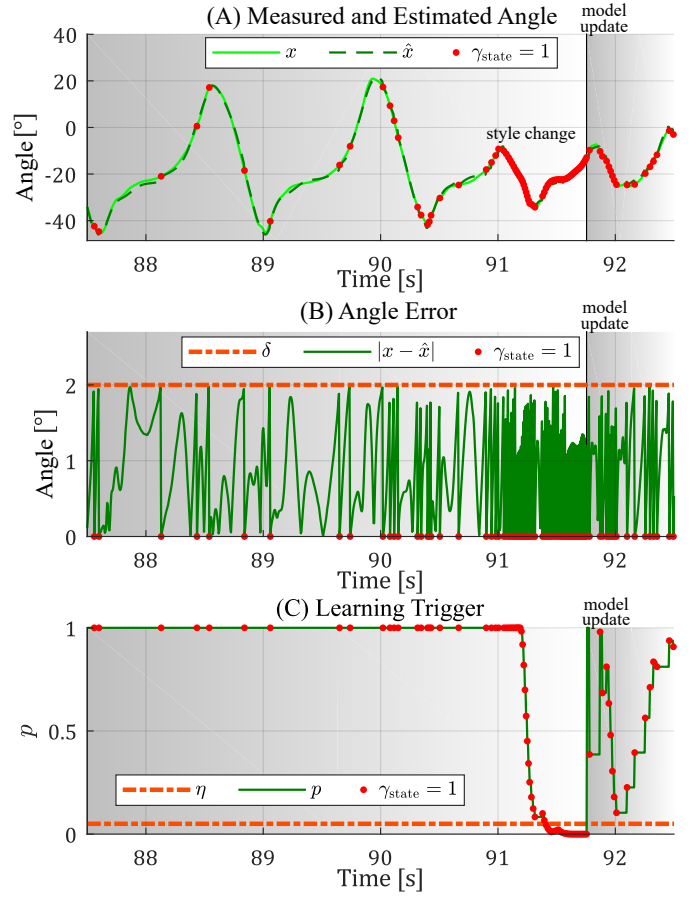


Fig. 6. Full model update ($\gamma_{\text{learn}}[k] = 1 \wedge \gamma_{\text{full}}[k] = 1$) triggered at 91.8 s due to a change of trajectory shape (walking style). Plotted signals and quantities are the same as in Fig. 5.

- [11] V. Nguyen and Y. Suh, "Networked estimation with an area-triggered transmission method," *Sensors*, vol. 8, no. 2, pp. 897–909, 2008.
- [12] S. Trimpe and M. C. Campi, "On the choice of the event trigger in event-based estimation," in *Int. Conference on Event-based Control, Communication, and Signal Processing (EBCCSP)*, 2015, pp. 1–8.
- [13] L. H. Miller, "Table of percentage points of Kolmogorov statistics," *Journal of the American Statistical Association*, vol. 51, no. 273, pp. 111–121, 1956.
- [14] G. Marsaglia, W. W. Tsang, J. Wang *et al.*, "Evaluating Kolmogorovs distribution," *Journal of Stat. Software*, vol. 8, no. 18, pp. 1–4, 2003.
- [15] F. Solowjow and S. Trimpe, "Event-triggered learning," *arXiv:1904.03042*, 2019. [Online]. Available: <https://arxiv.org/abs/1904.03042>
- [16] A. Antoniadis and G. Oppenheim, *Wavelets and statistics*. Springer Science & Business Media, 2012, vol. 103.
- [17] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT Press Cambridge, MA, 2006, vol. 2, no. 3.
- [18] J. Quiñero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate gaussian process regression," *Journal of Machine Learning Research*, vol. 6, no. Dec, pp. 1939–1959, 2005.
- [19] T. Seel, C. Werner, and T. Schauer, "The adaptive drop foot stimulator—multivariable learning control of foot pitch and roll motion in paretic gait," *Med. Eng. & Physics*, vol. 38, no. 11, pp. 1205–1213, 2016.
- [20] T. Zhang, D. Laidig, and T. Seel, "Stop repeating yourself: Exploitation of repetitive signal patterns to reduce communication load in sensor network," in *European Control Conference (ECC)*, Naples, Italy, 2019, accepted.
- [21] T. Seel, D. Graurock, and T. Schauer, "Realtime assessment of foot orientation by accelerometers and gyroscopes," *Current Directions in Biomedical Engineering*, vol. 1, no. 1, pp. 446–469, 2015.