

Event-triggered Learning for Resource-efficient Networked Control

Friedrich Solowjow^{1,2}, Dominik Baumann¹, Jochen Garcke^{2,3}, Sebastian Trimpe¹

Abstract—Common event-triggered state estimation (ETSE) algorithms save communication in networked control systems by predicting agents’ behavior, and transmitting updates only when the predictions deviate significantly. The effectiveness in reducing communication thus heavily depends on the quality of the dynamics models used to predict the agents’ states or measurements. Event-triggered learning is proposed herein as a novel concept to further reduce communication: whenever poor communication performance is detected, an identification experiment is triggered and an improved prediction model learned from data. Effective learning triggers are obtained by comparing the actual communication rate with the one that is expected based on the current model. By analyzing statistical properties of the inter-communication times and leveraging powerful convergence results, the proposed trigger is proven to limit learning experiments to the necessary instants. Numerical and physical experiments demonstrate that event-triggered learning improves robustness toward changing environments and yields lower communication rates than common ETSE.

I. INTRODUCTION

Networked control systems (NCSs) are rapidly gaining in popularity, both in academia and industry. Advancements in control strategies and network technologies enable the systems to closely interact with their environment and share data. Treating communication as a shared resource, as suggested in [1], is an important step to scale NCSs to problems involving many agents.

In this paper, we consider NCSs with multiple spatially distributed agents, whose dynamics are independent, but that are coupled through a joint control objective and communicate via a shared network. Figure 1 depicts two agents representative for one communication link in such an NCS. While communication between agents is beneficial or even necessary for coordination (e.g., formation control [2], or multi-agent balancing [3]), the network constitutes a shared and scarce resource and, hence, its usage shall be limited.

Event-triggered state estimation (ETSE) [4]–[8] has been proposed to reliably exchange sensor or state data between agents, but with limited inter-agent communication. Many ETSE methods utilize dynamics models to predict other agents’ states or measurements (see Fig. 1), in order to anticipate their behavior without the need for continuous data transmissions. Event triggering rules then ensure that

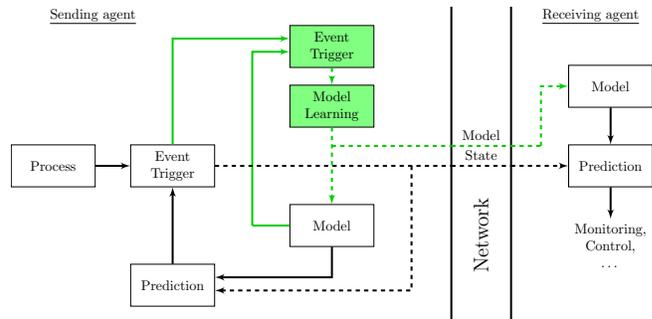


Fig. 1. Proposed event-triggered learning architecture. Based on a typical event-triggered state estimation architecture (in black), we propose model learning (in green) to improve predictions and thus lower communication between Sending and Receiving agents. Learning experiments are triggered by comparing empirical with expected inter-communication times.

an update is sent whenever the prediction is not good enough. Hence, the capability of the dynamics model in making accurate predictions critically determines how much communication can be saved.

In order to improve prediction accuracy, we propose to combine ETSE with model learning. We develop a framework, where an agent uses its input-output data to update the model of its dynamics, and communicates this model to other agents. With improved models, it is possible to achieve superior prediction accuracy and thus lower communication even further compared to standard ETSE, especially when facing changing dynamics. Since learning and communicating models can be resource-intensive operations themselves, we develop triggering rules that quantify the model accuracy and decide when to learn. The result is an *event-triggered learning* scheme, which sits on top of the standard ETSE framework (cf. Fig. 1).

Contributions: In detail, this paper makes the following contributions:

- introducing the novel idea of event-triggered learning;
- derivation of a data-driven learning trigger based on statistical properties of inter-communication times;
- probabilistic guarantees ensuring the effectiveness of the proposed learning triggers;
- concrete implementation of event-triggered learning for linear Gaussian dynamic systems; and
- demonstration of improved communication behavior in numerical simulations and hardware experiments on a cart-pole system.

Related work: Various event-triggering approaches have been proposed for improving resource usage in state estimation; for recent overviews see, e.g., [9]–[12] and ref-

¹Intelligent Control Systems Group, Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany. Email: {fsolowjow, dbaumann, strimpe}@tuebingen.mpg.de

²Inst. for Numerical Simulation, University of Bonn, 53115 Bonn, Germany.

³Fraunhofer SCAI, Schloss Birlinghoven, 53754 St. Augustin, Germany.

This work was supported in part by the German Research Foundation (DFG) Priority Program 1914 (grant TR 1433/1-1), the Cyber Valley Initiative, and the Max Planck Society.

erences therein. Approaches differ, among others, in the type of models that are used for predictions. The send-on-delta protocol [13] assumes a trivial constant model by triggering when the difference between the current and lastly communicated value passes a threshold. This protocol is extended to linear predictions in [14], which are obtained from approximating the signal derivative from data. More elaborate protocols use dynamics models of the observed process, which typically leads to more effective triggering [10], [15]. The vast majority of ETSE methods (e.g., [4]–[12], [15], [16]) employ linear dynamics models, which we also consider herein. Nonlinear prediction models are used in [17], [18], for example. None of these works considers model learning to improve prediction accuracy, as we do herein. While we implement event-triggered learning for linear Gaussian problems, the general idea also applies to other types of estimation problems and dynamics models.

In order to obtain effective learning triggers, we take a probabilistic view on inter-communication times (i.e., the time between two communication events) and trigger learning experiments whenever the expected communication differs from the empirical. A similar interpretation of inter-communication times is considered in [19], where it is used to infer stability results. We extend this idea with statistical convergence results to design the event trigger for learning.

Deciding if learning is necessary is related to the problem of change detection, which seeks to identify times when the distribution governing a random process changes. Change detection has received considerable attention in statistical literature; see [20] for an overview. While one main application of this work is fault detection, we use the proposed trigger to initiate learning experiments to capture the changed distribution. Furthermore, we do not analyze the stochastic process directly, but instead target inter-communication times, which are a natural one-dimensional feature in the ETSE framework and also amenable to theoretical analysis, which is incorporated in the trigger design. Applications of change detection in the context of NCSs focus on fault detection in the presence of delays [21], or network design and performance in general [22]. Iterative learning control has been proposed for multi-agent problems in [23] to improve event-triggered control for repetitive tasks, however, the idea of triggering learning when there is significant change is new to the best of the authors’ knowledge.

II. EVENT-TRIGGERED LEARNING

In this section, we explain the main idea of event-triggered learning using the schematic in Fig. 1. The figure depicts a canonical problem, where one agent (‘Sending agent’ on the left) has information that is relevant for another agent at a different location (‘Receiving agent’). This setting is representative for remote monitoring scenarios or two agents of a multi-agent network, for instance. For resource-efficient communication, a standard ETSE architecture is used (shown in black). The main contribution of this work is to incorporate learning into the ETSE architecture. By designing an event trigger also for model learning (in green),

learning tasks are performed only when necessary. We next explain the core components of the proposed framework.

The sending agent in Fig. 1 monitors the state of a dynamic process (either directly measured or obtained via state estimation) and can transmit this state information to the remote agent via a network link. In order to save network resources, an event-based protocol is used. The receiving agent makes model-based predictions (‘Model’ and ‘Prediction’) to anticipate the state at times of no communication. The sending agent implements a copy of the same prediction and compares it to the current state in the ‘Event Trigger’, which triggers a communication whenever the prediction deviates too much from the actual state. This general scheme is standard in ETSE literature (see [9]–[12] and references therein). The effectiveness of this scheme in reducing communication will generally depend on the accuracy of the prediction, and thus the quality of the model.

The key idea of this work is to trigger learning experiments and learn improved models from data when prediction performance of the current model is poor. The updated model is then shared with the remote agent to improve its predictions. Because performing a learning task is costly itself (e.g., involving computation and communication resources, as well as possibly causing deviation from the control objective), we propose event-triggering rules also for model learning.

While the idea of using event triggering to save communication in estimation or control is quite common by now, this work proposes event triggering also on a higher level. Triggering of learning tasks yields improved prediction models, which are the basis for ETSE at the lower level.

The general idea of event-triggered learning applies to diverse scenarios. In the following, we make the idea concrete for ETSE of linear Gaussian systems (introduced in Sec. III). For this case, model learning can be solved with standard least-squares estimation (Sec. IV). We propose a trigger design that is based on probabilistic analysis of inter-communication times (Sec. V and VI). By means of numerical examples and physical experiments on a cart-pole system (Sec. VII and VIII), we show that communication can effectively be reduced through event-triggered learning.

III. LINEAR GAUSSIAN PROBLEM

In this section, we make the problem of event-triggered learning precise for linear Gaussian dynamic systems. For this, we introduce all standard elements of the ETSE architecture shown in Fig. 1 (black).

A. Process

We assume the linear dynamics (‘Process’ in Fig. 1)

$$x(k+1) = Ax(k) + Bu(k) + \epsilon(k), \quad (1)$$

with discrete-time index k , state $x(k) \in \mathbb{R}^n$, control input $u(k) \in \mathbb{R}^q$, system matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times q}$, and independent identically distributed (i.i.d.) Gaussian random variables $\epsilon(k) \sim \mathcal{N}(0, \Sigma)$. For simplicity, we assume the state $x(k)$ can be measured, but the framework readily extends to the case where the state is locally estimated.

The system (A, B) is assumed stabilizable. Hence, stable closed loop dynamics can be achieved through state feedback

$$u(k) = Fx(k) + r(k), \quad (2)$$

where $r(k)$ is a known reference. The reference can be used for tracking problems or to excite the system, which is particularly important during learning experiments and will be further discussed later. The closed loop dynamics then reads

$$x(k+1) = A_{\text{CL}}x(k) + Br(k) + \epsilon(k), \quad (3)$$

with stable matrix $A_{\text{CL}} := (A + BF)$.

B. Predictions

The ‘Prediction’ block, which is implemented on the sending and the receiving agent (cf. Fig. 1), utilizes a model (\hat{A}, \hat{B}) to predict the true process. After initializing $\hat{x}(0) = x(0)$, both agents iterate

$$\hat{x}(k+1) = \hat{A}_{\text{CL}}\hat{x}(k) + \hat{B}r(k). \quad (4)$$

The prediction (4) is deterministic and deviates from the true value (3) over time. We bound this error through an event trigger: whenever a certain error threshold is reached, the actual state $x(k+1)$ is transmitted and the prediction reset

$$\hat{x}(k+1) = \begin{cases} \hat{A}_{\text{CL}}\hat{x}(k) + \hat{B}r(k) & \text{if } \gamma_{\text{state}} = 0 \\ x(k+1) & \text{if } \gamma_{\text{state}} = 1 \end{cases}, \quad (5)$$

where the binary variable γ_{state} denotes positive ($\gamma_{\text{state}} = 1$) or negative triggering decisions.

C. Event Trigger

The ‘Event Trigger’ on the sending agent has access to both, prediction and true state. It can thus track the error

$$\|z(k)\|_2 := \|x(k) - \hat{x}(k)\|_2 \quad (6)$$

and trigger a communication when it becomes too large:

$$\gamma_{\text{state}} = 1 \iff \|z(k)\|_2 \geq \delta. \quad (7)$$

This way, communication is reduced to the necessary instants given by a significant prediction error.

D. Problem Formulation

Precise models $(\hat{A}, \hat{B}, \hat{\Sigma})$ are key to reduce communication rates. In this work, we address mismatches between model $(\hat{A}, \hat{B}, \hat{\Sigma})$ and true process (A, B, Σ) , which may stem from imprecise initial models or changing dynamics, for example. The development of a model learning framework to improve the effectiveness of ETSE in reducing communication is the main objective of this paper. This includes, in particular, a decision rule (the *learning trigger* γ_{learn}) for deciding when a new model $(\hat{A}, \hat{B}, \hat{\Sigma})$ ought to be learned.

IV. MODEL LEARNING

This section addresses learning of the dynamics model (4) from input-output data of the system (3) (‘Model Learning’ in Fig. 1). In the numerical and physical experiments of Sec. VII and VIII, the data is obtained from dedicated learning experiments. In other settings, data could be recorded during normal operation. We here present standard least squares to estimate the system matrices, but any other technique for linear system identification [24] could be used.

Rewriting (3) as

$$x(k+1) = \begin{bmatrix} A_{\text{CL}} & B \end{bmatrix} \begin{bmatrix} x(k) \\ r(k) \end{bmatrix} + \epsilon(k) \quad (8)$$

and stacking M samples yields

$$Y = \begin{bmatrix} A_{\text{CL}} & B \end{bmatrix} X + \begin{bmatrix} \epsilon(0) & \dots & \epsilon(M-1) \end{bmatrix}, \quad (9)$$

with

$$X := \begin{bmatrix} x(0) & \dots & x(M-1) \\ r(0) & \dots & r(M-1) \end{bmatrix}, \quad Y := \begin{bmatrix} x(1) & \dots & x(M) \end{bmatrix} \quad (10)$$

Thus, the ordinary least squares (OLS) estimator for the model matrices yields $[\hat{A}_{\text{CL}} \ \hat{B}] = YX^T(XX^T)^{-1}$.

Due to the auto-regressive structure of the system (3), we can ensure identifiability for certain types of input signals r . Sinusoidal input signals with increasing frequencies (chirp) are known to yield invertible matrices XX^T (see condition of persistent excitation, e.g., [24]). We thus use chirp signals as reference r to excite the system when a learning experiment shall be performed. Since the estimator is designed to achieve minimal variance it is straightforward to obtain $\hat{\Sigma}$ by averaging over the squared residuals.

The specific learning method only affects the ‘Model Learning’ block in Fig. 1 and has no influence on the decision whether learning is necessary (‘Event Trigger’). Hence, the general event-triggered learning approach is agnostic to what learning or identification method is used.

V. FOUNDATIONS IN STOCHASTIC ANALYSIS

In this section, we briefly discuss theoretical results from stochastic analysis as background for deriving the event trigger for model learning in the next section. More detailed treatments of stochastic processes and stochastic differential equations (SDEs) can be found in [25] and [26], for example.

Inter-communication times (i.e., the time between two triggering events (7)) will play a key role in the proposed triggering scheme for model learning. In the following, we express inter-communication times as random variables (stopping times) and deduce statistical properties directly from the model (in Sec. V-B). To be amenable to stopping time analysis, we shall first describe the event-triggered estimation scheme in terms of continuous-time SDEs (Sec. V-A).

The analysis in this section is done under the assumption that the prediction model (4) is perfect (i.e., $\hat{A}_{\text{CL}} = A_{\text{CL}}$, $\hat{B} = B$, and $\hat{\Sigma} = \Sigma$), which will be motivated in the next section.

A. Ornstein-Uhlenbeck processes

Consider (3) with $r = 0$, i.e.,

$$x(k+1) = A_{\text{CL}}x(k) + \epsilon(k). \quad (11)$$

Because of the perfect model assumption, the reference signal $r(k)$ cancels out in the later analysis step (15) and is hence irrelevant. Transforming system (11) to continuous time yields an Ornstein-Uhlenbeck (OU) process

$$dX(t) = \mathcal{A}X(t)dt + QdW(t), \quad (12)$$

with $t \in \mathbb{R}$, $X(t) \in \mathbb{R}^n$, and $W(t) \in \mathbb{R}^n$ a multidimensional Wiener process. The matrices \mathcal{A} , $Q \in \mathbb{R}^{n \times n}$ can be computed from A_{CL} , the sampling time of the discrete-time system (11), and the covariance Σ (refer to [26], [27] for details). Continuous- and discrete-time processes coincide in the sense that they have the same distribution for any given sampling time.

The OU process (12) can be written explicitly in integrated form as

$$X(t) = e^{\mathcal{A}t}X(0) + \int_0^t e^{\mathcal{A}(t-s)}Q dW(s). \quad (13)$$

With the perfect model assumption, the discrete-time predictions (4) transform to

$$\hat{X}(t) = e^{\mathcal{A}t}X(0) \quad (14)$$

in continuous time. Using (13) and (14), we thus obtain for the continuous-time prediction error

$$Z(t) = X(t) - \hat{X}(t) = \int_0^t e^{\mathcal{A}(t-s)}Q dW(s). \quad (15)$$

By comparing to (13), it follows that $Z(t)$ is an OU process starting in zero. We will further analyze the behavior of $\|Z(t)\|_2$ and make a connection between stopping times of this process and communication behavior.

B. Stopping Times

Due to the existence and uniqueness result for SDEs [25], continuity of sample paths is assured for OU processes. We will leverage this property to pinpoint the exact moment the process crosses a given threshold. Exactly like in the discrete-time setting, state information is communicated whenever $\|Z(t)\|_2 \geq \delta$, which resets $Z(t)$ to 0. Accordingly, the stopping time τ is defined as the first time when the stochastic process $Z(t)$ exits an n -dimensional sphere with radius δ :

$$\tau := \inf\{t : \|Z(t)\|_2 \geq \delta\}. \quad (16)$$

The random variable τ hence corresponds to the inter-communication time, the time between communication events. Because they coincide in the setting herein, ‘inter-communication time’ and ‘stopping time’ will be used synonymously hereafter. In Sec. VI, learning triggers will be proposed that are based on a comparison of empirical and expected inter-communication times. The computation of expected inter-communication times $\mathbb{E}[\tau|Z(0) = 0]$ is discussed next.

For certain classes of SDEs, it is possible to compute $\mathbb{E}[\tau|Z(0) = x]$ as the solution $v(x)$ to specific partial differential equations (PDEs). The following lemma states this result for the OU process.

Lemma 1: Assume the boundary value problem

$$\begin{aligned} \frac{\partial v(x)}{\partial x} \mathcal{A}x + \frac{1}{2} \text{trace} \left[Q^\top \frac{\partial^2 v(x)}{\partial x^2} Q \right] &= -1 \quad \text{in } \Omega, \\ v(x) &= 0 \quad \text{on } \partial\Omega, \end{aligned} \quad (17)$$

with gradient $\frac{\partial v(x)}{\partial x}$, Hessian $\frac{\partial^2 v(x)}{\partial x^2}$, and Ω an n -dimensional sphere with radius δ , has a unique bounded solution $v(x)$. Then, $\mathbb{E}[\tau|Z(0) = 0] = v(0)$.

Proof: The result is obtained by using the more general Andronov-Vitt-Pontryagin formula from [26] and adapting it to the OU process (as was also done in [19]). ■

The one-dimensional case can be solved analytically (see Example 4.2 on page 111 in [26]) and gives interesting insights. From the solution, dependencies between $\mathbb{E}[\tau|Z(0) = x]$ and the parameters \mathcal{A} and Q can be seen. The magnitude of stochastic effects Q is pushing the process out of the domain, while the stable part \mathcal{A} drives it back to zero. Hence, more stable \mathcal{A} leads to longer stopping times, while larger Q leads to shorter.

For general dimension n , one could attempt solving the PDE (17), which is, however, challenging because it is non-linear and possibly high dimensional. Typically, one has to resort to numerical solutions, which usually yield the whole function $v(x)$ and are computationally intensive. Since we actually only require $v(0)$, an alternative is to approximate $\mathbb{E}[\tau|Z(0) = 0]$ through Monte Carlo simulations, which we use in the experiments herein. Because the error $Z(t)$ is always reset to 0 at triggering instants for the scenario herein, we shall omit the dependence on the initial value and write $\mathbb{E}[\tau]$ instead of $\mathbb{E}[\tau|Z(0) = 0]$ in the following.

VI. EVENT TRIGGER FOR MODEL LEARNING

In this section, we design the event trigger for model learning (green ‘Event Trigger’ block in Fig. 1). We first discuss the general idea of how to trigger learning experiments, and then state two concrete instances.

A. General idea

The learning trigger must reliably detect when the prediction accuracy of the current model is poor, and thus learning a new model from data promises improved predictions. We cannot directly compare the model $(\hat{A}_{\text{CL}}, \hat{B}, \hat{\Sigma})$ to the true process $(A_{\text{CL}}, B, \Sigma)$ because it is unknown. Owing to the stochasticity of the process, it is also not straightforward to quantify model accuracy from data. Since we ultimately care about the communication behavior that results from the models, we propose to utilize inter-communication times to trigger learning in the following way.

If an accurate model $(\hat{A}_{\text{CL}}, \hat{B}, \hat{\Sigma})$ is used, average inter-communication times are expected to be equal to $\mathbb{E}[\tau]$ (as computed according to Sec. V-B under the assumption $(\hat{A}_{\text{CL}}, \hat{B}, \hat{\Sigma}) = (A_{\text{CL}}, B, \Sigma)$). If, however, observed inter-communication times deviate on average from $\mathbb{E}[\tau]$, we

conclude an inaccurate model and trigger a learning experiment ($\gamma_{\text{learn}} = 1$). Hence, the proposed learning trigger compares empirically observed inter-communication times with the expected inter-communication time computed from the current model, and triggers learning when the two are significantly apart.

B. Exact learning trigger

Based on the foregoing discussion, we propose the following learning trigger:

$$\gamma_{\text{learn}} = 1 \iff \left| \frac{1}{N} \sum_{i=1}^N \tau_i - \mathbb{E}[\tau] \right| \geq \kappa_{\text{exact}}, \quad (18)$$

where $\gamma_{\text{learn}} = 1$ indicates that a new model shall be learned; $\mathbb{E}[\tau]$ is computed according to Sec. V-B assuming a perfect model; and $\tau_1, \tau_2, \dots, \tau_N$ are the last N empirically observed inter-communication times (τ_i the duration between two triggers (7)). The moving horizon N is chosen to yield robust triggers. The threshold parameter κ_{exact} quantifies the error we are willing to tolerate. We denote (18) as the *exact learning trigger* because it involves the exact expected value $\mathbb{E}[\tau]$, as opposed to the trigger derived in the next subsection.

Even though the trigger (18) is meant to detect inaccurate models, there is always a chance that the trigger fires not due to an inaccurate model, but instead due to the randomness of the process (and thus randomness of inter-communication times τ_i). Even for a perfect model, (18) may trigger at some point. This is inevitable due to the stochastic nature of the problem. Therefore, we propose to chose κ_{exact} to yield effective triggering with a user-defined confidence level. For this, we make use of Hoeffding's inequality:

Lemma 2 (Hoeffding's inequality [28]): Let τ_1, \dots, τ_N be i.i.d. bounded random variables with $\tau_i \in [0, \tau_{\text{max}}]$ for all i . Then

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \mathbb{E}[\tau] \right| \geq \kappa \right] \leq 2e^{-\frac{2N\kappa^2}{\tau_{\text{max}}^2}}. \quad (19)$$

The application of Hoeffding's inequality requires inter-communication times τ_i to be upper bounded by some τ_{max} . This is easily enforced in practice by triggering at the latest when τ_{max} is reached. Moreover, we specify a maximal probability η that we are willing to tolerate for the difference being caused through randomness. We then have the following result for the trigger (18):

Theorem 1 (Exact learning trigger): Let the parameters η , N , and τ_{max} be given, and assume a perfect model $(\hat{A}_{\text{CL}}, \hat{B}, \hat{\Sigma}) = (A_{\text{CL}}, B, \Sigma)$. For

$$\kappa_{\text{exact}} = \tau_{\text{max}} \sqrt{-\frac{1}{2N} \ln \frac{\eta}{2}}, \quad (20)$$

we have

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \mathbb{E}[\tau] \right| \geq \kappa_{\text{exact}} \right] \leq \eta. \quad (21)$$

Proof: Substituting (20) for κ_{exact} into the right hand side of Hoeffding's inequality yields the desired result. ■

The theorem quantifies the expected convergence rate of the empirical mean to the expected value for a perfect model. This result can be used as follows: the user specifies the desired confidence level η , the number N of inter-communication times considered in the empirical average, and the maximum inter-communication time τ_{max} . By choosing κ_{exact} as in (20), the exact learning trigger (18) is guaranteed to make incorrect triggering decisions (false positives) with a probability of less than η .

C. Approximated learning trigger

As discussed in Sec. V-B, obtaining $\mathbb{E}[\tau]$ can be difficult and computationally expensive. Instead of solving the PDE (17), we propose to approximate $\mathbb{E}[\tau]$ by sampling τ . For this, we simulate the OU process $Z(t)$ (15) until it reaches a sphere with radius δ for M times, and average the obtained stopping times $\tau_1^{\text{sim}}, \dots, \tau_M^{\text{sim}}$. Hence,

$$\mathbb{E}[\tau] \approx \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}}. \quad (22)$$

This yields the *approximated learning trigger*

$$\gamma_{\text{learn}} = 1 \iff \left| \frac{1}{N} \sum_{i=1}^N \tau_i - \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}} \right| \geq \kappa_{\text{approx}}. \quad (23)$$

This approximation leads to a different choice of κ_{approx} .

Theorem 2 (Approximated learning trigger): Let the parameters η , N , $M > N$, and τ_{max} be given, and assume $(\hat{A}_{\text{CL}}, \hat{B}, \hat{\Sigma}) = (A_{\text{CL}}, B, \Sigma)$. For

$$\kappa_{\text{approx}} = \tau_{\text{max}} \sqrt{-\frac{2}{N} \ln \frac{\eta}{4}}, \quad (24)$$

we have

$$\mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}} \right| \geq \kappa_{\text{approx}} \right] < \eta. \quad (25)$$

Proof: We insert $\mathbb{E}[\tau]$, use the triangle inequality, and additivity of probability measures:

$$\begin{aligned} & \mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}} \right| \geq \kappa \right] \\ &= \mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \mathbb{E}[\tau] + \mathbb{E}[\tau] - \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}} \right| \geq \kappa \right] \\ &\leq \mathbb{P} \left[\left| \frac{1}{N} \sum_{i=1}^N \tau_i - \mathbb{E}[\tau] \right| \geq \frac{\kappa}{2} \right] + \\ & \quad \mathbb{P} \left[\left| \frac{1}{M} \sum_{i=1}^M \tau_i^{\text{sim}} - \mathbb{E}[\tau] \right| \geq \frac{\kappa}{2} \right] \\ &\leq \frac{\eta}{2} + \frac{\eta}{2}, \end{aligned}$$

where the last step follows with $M > N$ and Hoeffding's inequality for $\frac{\kappa}{2}$ and $\frac{\eta}{2}$. ■

We avoid solving the PDE (17), but the obtained trigger (23) is less efficient than (18), i.e., the required sample size N for equal accuracy η is higher.

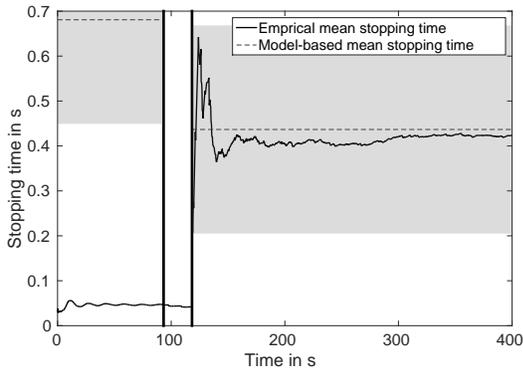


Fig. 2. Communication and learning behavior for the simulation example. The solid line shows the empirical inter-communication time computed as moving average $\frac{1}{N} \sum_{i=1}^N \tau_i$. The dashed line represents the model-based stopping time (22) with the triggering interval $\pm \kappa_{\text{approx}}$ in gray. The two vertical lines represent the start and the end of a learning experiment. After learning, the empirical and model-based stopping times match well.

Remark 1 (practical implementation): In the experiments reported in the next sections, we use trigger (23), which is easier to compute than (18). In order to improve the trigger's robustness, we trigger learning experiments only when (23) holds for a certain duration rather than instantaneous. This way, the trigger is less prone to unmodeled short-term effects, which also decreases the probability of false positives.

VII. SIMULATION

This section illustrates the proposed event-triggered learning scheme with a numerical example. For this, we consider two agents as in Fig. 1 and a one-dimensional process

$$x(k+1) = 0.9x(k) + 0.01r(k) + \epsilon(k) \quad (26)$$

with sample time $T_s = 0.01$ s, reference $r(k) = \cos(0.2kT_s)$, and $\epsilon(k) \sim \mathcal{N}(0, 2.5 \cdot 10^{-5})$. To simulate imperfect model information, we assume that only a slightly perturbed model is available for making predictions,

$$\hat{x}(k+1) = 0.85\hat{x}(k) + 0.005r(k). \quad (27)$$

We implement all main components of the proposed event-triggered learning scheme as shown in Fig. 1; that is, the ETSE components as described in Sec. III with $\delta = 0.02$, model learning as in Sec. IV, and the approximated learning trigger (23). As trigger parameters, we chose a confidence level $\eta = 0.05$, as well as $N = 2000$, $M = 10000$, and $\tau_{\text{max}} = 3.5$ s. Theorem 2 then yields $\kappa_{\text{approx}} \approx 0.23$. In order to obtain a robust learning trigger, N in the empirical inter-communication time average, $\frac{1}{N} \sum_{i=1}^N \tau_i$, is reset after a learning experiment, and (23) is only checked after having observed at least $N = 2000$ stopping times τ_i .

Figure 2 illustrates the behavior of the event-triggered estimation and learning system for this numerical example. The actual communication is captured by the empirical inter-communication time, which is the inverse of the communication rate and computed as the moving average $\frac{1}{N} \sum_{i=1}^N \tau_i$, where N is equal to 2000 or the number of stopping times observed since the start or last learning experiment. Since the initial model (27) is inaccurate, we observe a

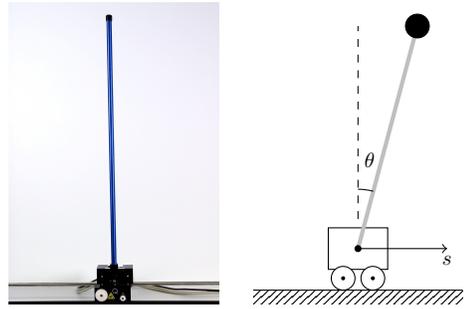


Fig. 3. The event-triggered learning approach is demonstrated in experiments of the depicted cart-pole system.

significant deviation between empirical inter-communication time and model-based one (22). At around 90 s, 2000 inter-communication times τ_i have been observed, and the trigger (23) is checked. Because the triggering condition (right-hand side of (23)) is clearly true, a learning experiment is performed (until about 120 s). Directly after learning, the moving average of the empirical inter-communication time is reset, hence its fluctuations. With the updated model, empirical and model-based stopping time coincide well; hence, no further learning experiment is triggered thereafter.

VIII. PHYSICAL EXPERIMENT

We demonstrate the proposed event-triggered learning approach in experiments on a cart-pole system (see Fig. 3), which is a common benchmark in control [29]. We consider balancing about the upright equilibrium, whose dynamics are approximately linear (1) with position s , velocity \dot{s} , angle θ , angular velocity $\dot{\theta}$ as its states, and motor voltage as input u . The dynamics are stabilized with a standard controller (2), and we consider remote state estimation of the stabilized process (3) as per the architecture in Fig. 1.

For making predictions in ETSE, we initially use a linear model supplied by the manufacturer [30], and we take $\delta = 0.075$ as triggering threshold. As the learning trigger, we use (23), where the empirical mean stopping time is computed over a fixed duration $T = 60$ s, which is an alternative to including a fixed number of samples in the average. While the physical system does not exactly match the assumption in Sec. V-A (e.g., no strictly linear Gaussian dynamics), the theoretical analysis of Sec. VI is still helpful in guiding our choice of a triggering threshold κ . We choose $\kappa = 2.5$ in the trigger (23), which is slightly lower than what we would obtain through Theorem 2. Since we additionally enforce the triggering condition to be true for 120 s, we reduce the risk of false positives.

Figure 4 shows the results of the experiment. While the model used for predictions is not perfect, there is still a good match between the empirical and model-based stopping time: the empirical stopping time mostly remains within the gray area and thus no learning is triggered initially. At 240s, we add weights on top of the pole (cf. Fig. 4), thus changing the system dynamics. The empirical inter-communication time drops (i.e., indicating more communication) and clearly captures the change in dynamics. After the

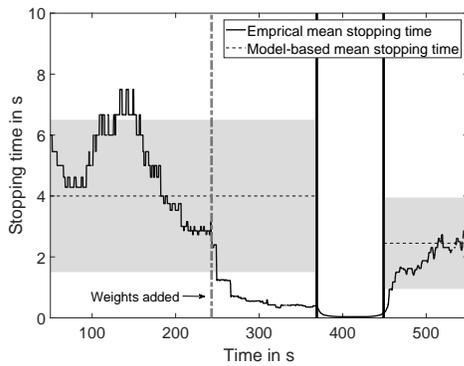


Fig. 4. Communication and learning results for the cart-pole experiments. Color coding for the graphs is the same as in Fig. 2. The vertical dashed line indicates a change of the system dynamics (increasing the pole mass).

triggering condition being true (the black graph being outside the gray area) for 120 s, a learning experiment is triggered. A new model is then learned according to Sec. IV (an open-loop model is identified instead of a closed-loop one, which makes no difference here). After updating the model, the empirical and model-based stopping times coincide very well. The event-triggered learning thus successfully detected and compensated for the changed dynamics, and reduced communication thereafter.

IX. CONCLUSION

Event-triggered learning is proposed in this paper as a novel concept to trigger model learning when needed. The concept is applied to event-triggered state estimation (ETSE) and shown to lead to reduced communication in simulation and physical experiments. For this setting, we obtained (provably) effective learning triggers by means of statistical stopping time analysis.

This paper is the first to develop the concept of event-triggered learning. Extending the method to nonlinear dynamic systems is an obvious next step we plan for future work. While event-triggered learning has been motivated as an extension to ETSE, the concept generally addresses the fundamental question of *when to learn*, and thus potentially has much wider relevance.

ACKNOWLEDGMENT

The authors thank Felix Grimminger for his support with the cart-pole system, and Alonso Marco for his help with speeding up the learning experiments.

REFERENCES

- [1] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proc. IEEE Special Issue on Technology of Networked Control Systems*, vol. 95, no. 1, pp. 138–162, Jan. 2007.
- [2] Y. Cao, W. Yu, W. Ren, and G. Chen, "An overview of recent progress in the study of distributed multi-agent coordination," *IEEE Trans. on Industrial Informatics*, vol. 9, no. 1, pp. 427–438, 2013.
- [3] S. Trimpe and R. D'Andrea, "The Balancing Cube: A dynamic sculpture as test bed for distributed estimation and control," *IEEE Control Systems Magazine*, vol. 32, no. 6, pp. 48–75, Dec. 2012.
- [4] J. K. Yook, D. M. Tilbury, and N. R. Soparkar, "Trading computation for bandwidth: reducing communication in distributed control systems using state estimators," *IEEE Transactions on Control Systems Technology*, vol. 10, no. 4, pp. 503–518, Jul. 2002.
- [5] S. Trimpe and R. D'Andrea, "An experimental demonstration of a distributed and event-based state estimation algorithm," in *18th IFAC World Congress*, 2011, pp. 8811–8818.
- [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] J. Wu, Q.-S. Jia, K. H. Johansson, and L. Shi, "Event-based sensor data scheduling: Trade-off between communication rate and estimation quality," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 1041–1046, 2013.
- [8] 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.
- [9] D. Shi, L. Shi, and T. Chen, *Event-Based State Estimation: A Stochastic Perspective*. Springer, 2015.
- [10] S. Trimpe and M. C. Campi, "On the choice of the event trigger in event-based estimation," in *Int. Conf. on Event-based Control, Communication, and Signal Processing*, 2015, pp. 1–8.
- [11] J. Sijs, B. Noack, M. Lazar, and U. D. Hanebeck, "Time-periodic state estimation with event-based measurement updates," in *Event-Based Control and Signal Processing*. CRC Press, 2016.
- [12] S. Trimpe, "Event-based state estimation: An emulation-based approach," *IET Control Theory & Applications*, vol. 11, no. 11, pp. 1684–1693, Jul. 2017.
- [13] M. Miskowicz, "Send-on-delta concept: an event-based data reporting strategy," *Sensors*, vol. 6, no. 1, pp. 49–63, 2006.
- [14] Y. S. Suh, "Send-on-delta sensor data transmission with a linear predictor," *Sensors*, vol. 7, no. 4, pp. 537–547, 2007.
- [15] J. Sijs, L. Kester, and B. Noack, "A study on event triggering criteria for estimation," in *17th International Conference on Information Fusion*, July 2014, pp. 1–8.
- [16] M. Muehlebach and S. Trimpe, "Distributed event-based state estimation for networked systems: An lmi approach," *IEEE Transactions on Automatic Control*, vol. 63, no. 1, pp. 269–276, Jan. 2018.
- [17] S. Trimpe and J. Buchli, "Event-based estimation and control for remote robot operation with reduced communication," in *IEEE Int. Conf. on Robotics and Automation*, 2015, pp. 5018–5025.
- [18] M. Martínez-Rey, F. Espinosa, A. Gardel, and C. Santos, "On-board event-based state estimation for trajectory approaching and tracking of a vehicle," *Sensors*, vol. 15, no. 6, pp. 14 569–14 590, 2015.
- [19] Y. Xu and J. P. Hespanha, "Communication logics for networked control systems," in *American Control Conference*, 2004, pp. 572–577.
- [20] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 613–658, 1995.
- [21] Y. Wang, S. X. Ding, H. Ye, and G. Wang, "A new fault detection scheme for networked control systems subject to uncertain time-varying delay," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 5258–5268, 2008.
- [22] M. Haghghi and C. J. Musselle, "Dynamic collaborative change point detection in wireless sensor networks," in *Int. Conf. on Cyber-Enabled Distributed Computing and Knowledge Discovery*, 2013, pp. 332–339.
- [23] T. Zhang and J. Li, "Event-triggered iterative learning control for multi-agent systems with quantization," *Asian Journal of Control*, 2016.
- [24] L. Ljung, *System identification: Theory for the User*. Prentice Hall, New Jersey, 1999.
- [25] B. Øksendal, "Stochastic differential equations," in *Stochastic differential equations*. Springer, 2003, pp. 65–84.
- [26] Z. Schuss, *Theory and Applications of Stochastic Processes: An Analytical Approach*, ser. Applied Mathematical Sciences. Springer New York, 2009.
- [27] K. J. Åström, *Introduction to stochastic control theory*. Dover Publications, 2006.
- [28] U. von Luxburg and B. Schölkopf, *Statistical Learning Theory: Models, Concepts, and Results*. Amsterdam, Netherlands: Elsevier North Holland, May 2011, vol. 10, pp. 651–706.
- [29] O. Boubaker, "The inverted pendulum benchmark in nonlinear control theory: A survey," *International Journal of Advanced Robotic Systems*, vol. 10, no. 5, p. 233, 2013.
- [30] Quanser Inc., "IP02 - self-erecting single inverted pendulum (SESIP) - linear experiment #6: PV and LQR control - instructor manual."