

Finite Sample Guarantees for Quantile Estimation: An Application to Detector Threshold Tuning

David Umsonst¹, Justin Ruths², and Henrik Sandberg¹, *Senior Member, IEEE*

Abstract—In threshold-based anomaly detection, we want to tune the threshold of a detector to achieve an acceptable false alarm rate. However, tuning the threshold is often a non-trivial task due to unknown detector output distributions. A detector threshold that provides an acceptable false alarm rate is equivalent to a specific quantile of the detector output distribution. Therefore, we use quantile estimators based on order statistics to estimate the detector threshold. The estimation of quantiles from sample data has a more than a century-long tradition and we provide three different distribution-free finite sample guarantees for a class of quantile estimators. The first is based on the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality, the second utilizes the Vysochanskij–Petunin inequality, and the third is based on exact confidence intervals for a beta distribution. These guarantees are then compared and used in the detector threshold tuning problem. We use both simulated data as well as data obtained from an experimental setup with the Temperature Control Lab to validate the guarantees provided.

Index Terms—Detector threshold tuning, fault detection, finite sample guarantees, quantile estimation.

I. INTRODUCTION

IN A HIGHLY automated society the automatic detection of anomalies is of utmost importance. The failure of detecting anomalies can have dire consequences, especially when the anomaly occurs in infrastructures critical to our daily life such as power grids and water distribution networks. Two notable incidences of undetected anomalies in critical infrastructures are the Northeast Blackout in 2003 [1], where the software did not notify the operators about an anomaly, which then led to a cascading failure of the power grid, and the attack on the Ukrainian power grid [2], where attackers managed to take over a distribution power grid.

A detector needs to not only be able to detect anomalies, but to not trigger on nominal behavior. Alarms during nominal behavior are called *false alarms*. False alarms increase the cost of detectors and make detectors unreliable. For example, in the survey [3], out of 460 anesthesiologists who stated that they deliberately turned off an alarm device, 68.2% named too many false alarms as a reason for turning off the alarm device.

Manuscript received 17 April 2022; accepted 19 July 2022. Date of publication 19 September 2022; date of current version 23 February 2023. This work was supported in part by the Swedish Research Council under Grant 2016-00861 and in part by the Swedish Civil Contingencies Agency under Grant MSB 2020-09672. Recommended by Associate Editor P. Mhaskar. (Corresponding author: David Umsonst.)

David Umsonst and Henrik Sandberg are with the Division of Decision and Control Systems, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, 10044 Stockholm, Sweden (e-mail: umsonst@kth.se; hsan@kth.se).

Justin Ruths is with the Department of Mechanical Engineering, The University of Texas at Dallas, Richardson, TX 75080 USA (e-mail: jruths@utdallas.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TCST.2022.3199668>.

Digital Object Identifier 10.1109/TCST.2022.3199668

Therefore, both the detection rate of anomalies as well as the false alarm rate need to be taken into account when tuning the anomaly detector. One tool to evaluate the performance of a detector is the receiver-operator-characteristic (ROC) curve [4]. The ROC curve plots the detection rate over the false alarm rate for different detector tuning and the higher the detection rate is for a smaller false alarm rate the better the detector performs. With the emerging threat of cyber-attacks on cyber-physical systems in recent years, Urbina *et al.* [5] argue that the impact of an attacker should also be taken into account when tuning the detector threshold.

Often we do not have exact knowledge about the statistics of the nominal and the anomalous behavior, but we have access to data which can be used to both design and tune the detector. Three different approaches are presented in [6] for this task. The first approach is supervised learning, where data from both nominal and anomalous behavior are used to tune the detector. The second approach is unsupervised learning, where data is available but the algorithm has to determine, what is anomalous and what is normal behavior. The third approach provides a middle ground since it is a semi-supervised learning approach, which uses only nominal data to tune the detector. A method to evaluate the performance of semi-supervised tuning of detectors is proposed in [7].

In this work, we utilize a semi-supervised learning approach and use independent and identically distributed (i.i.d.) samples of the detector output under nominal behavior to estimate a detector threshold that leads to a false alarm rate, which is close to the acceptable false alarm rate with a high probability. Since the threshold that guarantees a pre-defined acceptable false alarm rate is equivalent to a certain quantile of the detector output distribution, we utilize a sample-based quantile estimator to estimate a detector threshold. These quantile estimators typically use one- or two-order statistics of the sample data to determine the quantile and are simple to implement. Therefore, these quantile estimators are commonly used in statistical software packages [8].

It is not clear though how many samples of the detector output are needed to be close to the acceptable false alarm rate when using a threshold estimate based on sample data. Therefore, the contribution of this work is twofold. First, we provide three different finite guarantees to determine the sample size needed to be close to the acceptable false alarm rate with high probability. The first finite guarantee, which we proposed in [9], is based on the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality [10], the second finite guarantee is based on the Vysochanskij–Petunin inequality [11], and the third finite guarantee is based on exact confidence intervals of beta random variables [12]. All three finite guarantees are based on samples from the detector output only and the results

are distribution-free and independent of how the anomaly detector determines its output. Since we use a quantile estimator to estimate the threshold these distribution-free finite guarantees are also finite guarantees for the estimation of quantiles. Second, we perform a thorough validation of the finite guarantees with both simulated, and real data obtained from an experimental setup.

In the literature, it is quite common to use nominal data or make assumptions on the nominal behavior when determining the threshold that guarantees an acceptable false alarm rate. For example, under the assumption of a Gaussian distribution for the detector input, Murguia and Ruths [13] give a closed-form solution to tune a χ^2 detector and approximations of solutions to a cumulative sum (CUSUM) detector with resetting for an acceptable false alarm rate. Since the true nominal detector distribution is often not known, in [14] a distributionally robust approach is proposed, which makes assumptions on the finiteness of the moments of the input to a χ^2 detector.

More recently attention has turned to sample-based methods that can detect anomalies without requiring a formal model of behavior. Although having more samples is intuitively better, it is important to establish the minimum number of samples necessary for detector tuning so that the detector threshold might be adjusted adaptively over time. Sample guarantees also provide characterizations that show how detection confidence can be improved if the detector has access to more than the minimum number of samples required. Li and Martínez [15] propose a new detector based on the Wasserstein distance, which uses a sample-based tuning approach to achieve an acceptable false alarm rate. The tuning method uses the detector inputs under nominal behavior and assumes a light-tailed distribution for the detector inputs. Another approach to learn detectors from nominal behavior based on M -estimation is provided in [16]. Our approach has the advantage that no certain detector structure needs to be assumed and no knowledge about distributions is necessary such as in [13], [14], and [15] because the approach is purely based on samples.

Notation: Let \mathbb{R} and \mathbb{Q} denote the set of real and rational numbers, respectively. We call $\gamma = n_1/n_2$ the irreducible fraction of $\gamma \in \mathbb{Q}$ if and only if n_1 and n_2 are coprime integers. Let $x \in \mathbb{R}$, then $|x|$, $\lceil x \rceil$, and $\lfloor x \rfloor$ denote the absolute value of x , the smallest integer larger than or equal to x , and the largest integer smaller than or equal to x , respectively. Given a set $\{x_i\}_{i=1}^N$, the i th order statistic, $x_{(i)}$ is the i th largest element in $\{x_i\}_{i=1}^N$, such that the set of order statistics $\{x_{(i)}\}_{i=1}^N$ is $\min_i x_i = x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(N)} = \max_i x_i$. A random variable X that follows a beta distribution with parameters m and n is denoted as $X \sim \text{Beta}(m, n)$. Given an event E , its probability, expected value, variance, and the indicator function of E are given by $\text{prob}\{E\}$, $\mathbb{E}\{E\}$, $\text{Var}\{E\}$, and $\mathbf{1}_E$, respectively.

II. PROBLEM FORMULATION

In this section, we present the detector tuning problem, how it relates to the quantile of a random variable, and formulate

the problem of determining a sample-based threshold, which with high probability guarantees only a small deviation from the acceptable false alarm.

A. Tuning Detector Thresholds

The problem of anomaly detection occurs in many different fields. For example, in healthcare when devices monitor a patient or in governmental agencies to detect tax fraud. In this work, we look at anomaly detection in the context of a control system, where the feedback system is equipped with an anomaly detector on the controller side (see Fig. 1). The input to the anomaly detector can depend both on the measurements received as well as the actuator signals determined by the controller, which are not necessarily scalar variables. For example, if a Kalman filter is used the input to the anomaly detector are the actuator signals and the sensor measurements, which are used to determine the difference, r , between the received and predicted measurements. This difference can be further processed to determine the output y_D , e.g., $y_D = r^\top r$.

In the control system example, a small detector output $y_D \in \mathbb{R}$ typically indicates that the system works as predicted, while large outputs indicate an unpredicted behavior. However, in other applications, the detector output can also be a similarity measure, where a small value for y_D indicates anomalous behavior (see, for example, [16]). In this work, we use the interpretation that a large output indicates anomalous behavior. Therefore, an alarm is triggered when y_D exceeds a threshold $J_D \in \mathbb{R}$, i.e., $y_D > J_D$, and no alarm is triggered when $y_D \leq J_D$. Due to random processes, such as measurement noise, the detector output is also of a stochastic nature.

Assumption 1: In the nominal case, the detector output, y_D , is a random variable with a continuous cumulative distribution function $F(y_D)$.

For the sake of simplicity, we assume a continuous CDF for the detector output. In an industrial process control context, the process is often controlled around a desired steady-state value, which can be seen as stationary. Hence, assuming that the detector output is a random variable with a fixed distribution is a reasonable choice. Furthermore, if the plant has linear dynamics and the noise processes are Gaussian, a Kalman filter converges to a stationary distribution for its residual signals, which are often used as the input of the anomaly detector. If the anomaly detector has no internal dynamics, such as a neural network or norm-based detector, then the output of the detector is also a random variable with a fixed distribution.

The stochasticity of y_D leads to alarms in the nominal case, so-called false alarms, where the rate of false alarms depends on the threshold. Since false alarms incur unnecessary costs and too many false alarms make a detector unreliable, we can choose a large threshold to avoid too many false alarms. However, a too large threshold leads usually to a smaller detection rate of anomalies. Hence, there is a trade-off between the detection rate of anomalies and the false alarm rate in the nominal case when tuning the threshold. Furthermore, Urbina *et al.* [5] point out that there is also a trade-off between the impact an attacker can have that wants to keep y_D below




Fig. 1. Plant is controlled with a controller and the controller side is equipped with an anomaly detector that utilizes both the measurements and the actuator signals to determine its output y_D .

the threshold and the false alarm rate when choosing the detector threshold.

In addition to that, the nature of anomalies is in most cases unknown. Take, for example, a complex large-scale system such as the power grid, where many different types of anomalies, such as sensor and generator failures or attacks, can occur. Therefore, we often do not know which anomalies will occur and what detection rate we will obtain for a certain threshold J_D . Since it is simpler to focus only on the nominal case instead of trying to consider all possible anomalous behaviors, we will focus on finding a threshold that guarantees an acceptable false alarm rate. More specifically, we want to find the smallest threshold J_D such that

$$\text{prob}\{y_D \leq J_D\} \geq \gamma \quad (1)$$

holds, where $\gamma \in (0, 1)$. The threshold J_D will result in a false alarm probability of at most $1 - \gamma$. We want to find the smallest threshold because a trivial solution to guaranteeing an acceptable false alarm rate is to choose an arbitrarily large threshold, which in turn will also reduce the detection rate of anomalies.

In special cases, a closed form solution for the threshold exists, see, for example, [13], or the detector threshold can be approximated in a distributionally robust fashion, see [14]. If the output of the detector depends on the threshold as well, such as for the CUSUM detector with resetting, it is more difficult to determine a threshold that guarantees a desired false alarm rate. Detectors whose output depends on the threshold are one avenue of future work and will not be considered in this work.

Assumption 2: The detector output y_D does not depend on the threshold J_D .

B. Problem of Guaranteeing a False Alarm Rate

Before we formulate the problem we consider, we want to define the notion of a γ -quantile.

Definition 1: The γ -quantile J_D of a probability distribution is defined as

$$J_D = \inf\{y_D : F(y_D) \geq \gamma\}$$

where $\gamma \in (0, 1)$.

Note that J_D in Definition 1 is unique and finite because of the infimum operator and the fact that $\gamma \in (0, 1)$. Since $F(J_D) = \text{prob}\{y_D \leq J_D\}$, we can immediately see that

the threshold we are looking for in (1) is equivalent to the γ -quantile of the detector output distribution. Further, γ -quantiles can be obtained as optimizers of convex optimization problems (see [17]), which theoretically makes them efficient to calculate. However, it is often not possible to find an expression for the probability distribution of the detector output. A reason for that is that either the plant dynamics, the controller dynamics, the detector dynamics or all of them are highly nonlinear, such that the distribution of y_D does not have a closed-form solution. Samples from the monitored process under nominal conditions are usually available, such that we can obtain samples from the detector output under nominal conditions. Due to the equivalence with γ -quantiles, we use sample-based quantile estimators for a sample-based estimate \tilde{J}_D of the threshold J_D . In this work, we use N independent and identically distributed (i.i.d.) samples of y_D , $\{y_{D,i}\}_{i=1}^N$, to estimate the detector threshold as

$$\tilde{J}_D = \beta y_{D,(\lfloor N\gamma \rfloor)} + (1 - \beta) y_{D,(\lfloor N\gamma \rfloor + 1)} \quad (2)$$

where $\beta \in (0, 1)$. Note that (2) has the form of common quantile estimators used in software packages [8]. In our previous work, we showed how the quantile estimators can be derived from a sample approximation of the conditional value-at-risk [9].

Although the true threshold J_D can be approximated via (2), an open problem is to determine how many samples we need to get a good approximation. A good approximation is often characterized by assuming that $|J_D - \tilde{J}_D|$ is small with high probability. Distribution-free bounds on the bias of quantile estimates can be found in [18], where these bounds are always fulfilled and not only with high probability. However, depending on the shape of the probability distribution even small deviations from the threshold J_D can lead to large deviations in the false alarm rate. Therefore, we are more interested in how close the false alarm rate $1 - F(\tilde{J}_D)$ is to the acceptable false alarm rate $1 - F(J_D) = 1 - \gamma$, since the acceptable false alarm rate is an important variable for the system operator.

Problem 1: Determine the number, N , of i.i.d. detector output samples needed such that

$$\text{prob}\{|F(\tilde{J}_D) - \gamma| \leq \epsilon\} \geq 1 - \rho$$

holds, where $\gamma \in (0, 1)$, $\epsilon \in (0, 1)$, $\rho \in (0, 1)$, and \tilde{J}_D is an estimator of the form given in (2).

Enforcing that the false alarm rate is close to the acceptable false alarm rate with a high probability, which means that ϵ and ρ are close to zero, will intuitively lead to threshold estimates that are not just trivially large to guarantee the acceptable false alarm rate.

III. FINITE SAMPLE GUARANTEES

In this section, we use three different approaches to obtain distribution-free finite sample guarantees that solve Problem 1 and compare the finite guarantees with each other.

A. Finite Guarantees From the DKW Inequality

The first sample guarantee is based on the DKW inequality [10], which gives us probabilistic bounds on how close the

empirical distribution function $F_N(y_D) = 1/N \sum_{i=1}^N \mathbf{1}_{\{y_{D,i} \leq y_D\}}$ is to the true cumulative distribution function $F(y_D)$

$$\text{prob}\{|F(y_D) - F_N(y_D)| \leq \epsilon\} \geq 1 - 2e^{-2N\epsilon^2}. \quad (3)$$

While the DKW inequality gives us bounds for the whole CDF, we will evaluate it only at the point of interest, which is $y_D = \tilde{J}_D$, to obtain our finite guarantees. Note that the result was presented in our previous paper [9], but is restated for the sake of completeness.

Proposition 1 (Proposition 3 in [9]): Assume we have N i.i.d. samples, $\{y_{D,i}\}_{i=1}^N$, of the detector output y_D and $\gamma \in \mathbb{Q}$ such that $\gamma = n_1/n_2$ is an irreducible fraction. A solution to Problem 1 is given by $N = \lceil \ln(2\rho^{-1}) / (2\epsilon^2 n_2) \rceil n_2$ if $\beta \in [0, 1)$ in (2).

Proof: First, we choose $N = kn_2$, where $k \in \mathbb{N}$, such that $N\gamma = kn_1$ is an integer as well. Further, with $\beta \in [0, 1)$ we obtain that $F_N(\tilde{J}_D) = \gamma$. Evaluating the DKW inequality (3) at $y_D = \tilde{J}_D$ leads then to

$$\text{prob}\{|F(\tilde{J}_D) - \gamma| \leq \epsilon\} \geq 1 - 2e^{-2kn_2\epsilon^2}.$$

Finally, we set $\rho = 2e^{-2kn_2\epsilon^2}$ and solving for k leads to $k = \lceil \ln(2\rho^{-1}) / (2\epsilon^2 n_2) \rceil$, which concludes the proof. ■

Since the DKW inequality takes the whole CDF into account, this finite guarantee can be conservative. Here, we are only interested in the point $y_D = \tilde{J}_D$ and not the complete probability distribution. Therefore, we propose two more finite guarantees that evaluate the CDF at $y_D = \tilde{J}_D$ in the following.

B. Finite Guarantees From Vysochanskij–Petunin Inequality

Instead of focusing on the whole probability distribution as in the previous section, we now use the statistics of order statistics to determine a finite sample guarantee which utilizes the Vysochanskij–Petunin inequality [11]. The Vysochanskij–Petunin inequality is given by

$$\text{prob}\{|X - \mathbb{E}\{X\}| \geq \epsilon\} \leq \frac{4\text{Var}\{X\}}{9\epsilon^2} \quad (4)$$

if $3\epsilon^2 \geq 8\text{Var}\{X\}$, where X is a unimodal random variable with a finite mean and variance. Since for the Vysochanskij–Petunin inequality the expected value and variance of a random variable are needed, we introduce the expected value and variance of the CDF of a random variable evaluated at the m th order statistic.

Lemma 1: Let $y_{D,(m)}$ be the m th order statistics of N i.i.d. samples with CDF $F(\cdot)$. Then $F(y_{D,(m)})$ has a unimodal beta distribution with parameters m and $N + 1 - m$, i.e., $F(y_{D,(m)}) \sim \text{Beta}(m, N + 1 - m)$, and the expected value and the variance of $F(y_{D,(m)})$ are given by

$$\mathbb{E}\{F(y_{D,(m)})\} = \frac{m}{N + 1} \quad (5)$$

and

$$\text{Var}\{F(y_{D,(m)})\} = \frac{m(N + 1 - m)}{(N + 1)^2(N + 2)} \quad (6)$$

respectively.

Proof: From Chapter 2 in [19], we know that $F(y_{D,(m)}) \sim \text{Beta}(m, N + 1 - m)$. Hence, (5) and (6) are the expected value

and variance of the beta distribution with parameters m and $N + 1 - m$, respectively. Further, since $m \in \{1, \dots, N\}$, both parameters of the beta distribution are larger than or equal to one, which indicates that the beta distribution is unimodal (see Chapter 2 in [20]). ■

Interestingly, neither the expected value nor the variance of the CDF at the m th order statistic depend on the distribution itself. This is used in the following to determine a distribution-free finite sample guarantee.

Theorem 1: Assume we have N i.i.d. samples $\{y_{D,i}\}_{i=1}^N$ of the detector output y_D and let $\gamma \in \mathbb{Q}$ such that $\gamma = n_1/n_2$ is its irreducible fraction. A solution to Problem 1 is given by $N = kn_2 - 1$, where

$$k = \left\lceil \frac{1}{n_2} \left(\frac{4\gamma(1-\gamma)}{9\rho\epsilon^2} - 1 \right) \right\rceil \quad (7)$$

if $4\gamma(1-\gamma) > 9\rho\epsilon^2$, $6\rho \leq 1$, and $\tilde{J}_D = y_{D,(\lfloor N\gamma \rfloor + 1)}$.

Proof: Let $\tilde{J}_D = y_{D,(m)}$. Then the false alarm rate of this threshold is given by $1 - F(y_{D,(m)})$, which is a random variable that depends on the samples obtained. Hence, we use the Vysochanskij–Petunin inequality (4) and Lemma 1 to obtain

$$\text{prob}\left\{\left|F(y_{D,(m)}) - \frac{m}{N + 1}\right| \geq \epsilon\right\} \leq \frac{4m(N + 1 - m)}{9\epsilon^2(N + 1)^2(N + 2)}.$$

Next, with $k \in \mathbb{N}$ we set $N = kn_2 - 1$ and $m = \lfloor N\gamma \rfloor + 1 = kn_1$ such that $m/(N + 1) = \gamma$, which leads to

$$\text{prob}\{|F(y_{D,(\lfloor N\gamma \rfloor + 1)}) - \gamma| \geq \epsilon\} \leq \frac{4\gamma(1-\gamma)}{9\epsilon^2(kn_2 + 1)}.$$

By introducing $\rho = (4\gamma(1-\gamma))/(9\epsilon^2(kn_2 + 1))$ and solving for k we obtain (7) by making sure that k is an integer. Next, to guarantee that $k \geq 1$, we need to introduce the condition $4\gamma(1-\gamma) > 9\rho\epsilon^2$. Finally, for the Vysochanskij–Petunin inequality (4) to hold we need $3\epsilon^2 \geq 8\text{Var}\{y_{D,(\lfloor N\gamma \rfloor + 1)}\}$ for the determined sample size N . Let us look at an upper bound for the variance first

$$\begin{aligned} \text{Var}\{y_{D,(\lfloor N\gamma \rfloor + 1)}\} &= \frac{\gamma(1-\gamma)}{\left[\frac{1}{n_2} \left(\frac{4\gamma(1-\gamma)}{9\rho\epsilon^2} - 1 \right) \right] n_2 + 1} \\ &\leq \frac{\gamma(1-\gamma)}{\frac{1}{n_2} \left(\frac{4\gamma(1-\gamma)}{9\rho\epsilon^2} - 1 \right) n_2 + 1} = \frac{9}{4}\epsilon^2\rho. \end{aligned}$$

From this upper bound we obtain that if $6\rho \leq 1$ then $3\epsilon^2 \geq 8\text{Var}\{y_{D,(\lfloor N\gamma \rfloor + 1)}\}$ holds. ■

We would like to point out that the condition $\rho \leq 1/6$ is not too restrictive, since ρ is typically chosen to be small to achieve guarantees with high probability. Instead of using the Vysochanskij–Petunin inequality one could also use Chebyshev's inequality [21], which leads to more relaxed conditions on ρ and ϵ at the expense of having more conservative sample guarantees.

C. Finite Guarantees From Confidence Intervals of the Beta Distribution

In the previous section, we determined finite guarantees based on the Vysochanskij–Petunin inequality, which uses the expected value and variance of $F(y_{D,(m)})$. In this section,

we will use that $F(y_{D,(m)})$ has a beta distribution and determine sample guarantees based on confidence intervals of the beta distribution. We begin by introducing a result on the confidence interval of a beta-distributed random variable.

Lemma 2: Let X be distributed according to a beta distribution with parameters m and $N + 1 - m$, i.e., $X \sim \text{Beta}(m, N + 1 - m)$, and let $\hat{\gamma} = m/N \in (0, 1)$. Then

$$\text{prob}\{\hat{\gamma} - \epsilon_l \leq X \leq \hat{\gamma} + \epsilon_u\} \geq 1 - \rho$$

where $\rho \in (0, 1)$

$$\begin{aligned} \epsilon_u &= \frac{\sqrt{\hat{\gamma} - \hat{\gamma}^2}}{\sqrt{N}} z_{\frac{\rho}{2}} + \frac{2(0.5 - \hat{\gamma})z_{\frac{\rho}{2}}^2 - 1 - \hat{\gamma}}{3N} \\ \epsilon_l &= \frac{\sqrt{\hat{\gamma} - \hat{\gamma}^2}}{\sqrt{N}} z_{\frac{\rho}{2}} - \frac{2(0.5 - \hat{\gamma})z_{\frac{\rho}{2}}^2 - 1 - \hat{\gamma}}{3N} \end{aligned}$$

up to order $\mathcal{O}(N^{-3/2})$, and $z_{\rho/2}$ is the upper $\rho/2$ -quantile of the standard Gaussian distribution.

Proof: From Lemma 1 and Theorem 1 in [12], we obtain that $\text{prob}\{X \leq \gamma_q\} = \rho_q$ holds for

$$\gamma_q = \hat{\gamma} - \frac{\sqrt{\hat{\gamma} - \hat{\gamma}^2}}{\sqrt{N}} z_{\rho_q} + \frac{2(0.5 - \hat{\gamma})z_{\rho_q}^2 - 1 - \hat{\gamma}}{3N}$$

up to order $\mathcal{O}(N^{-3/2})$. Using $\rho_q = \rho/2$ and $\rho_q = 1 - \rho/2$ to obtain γ_L and γ_U , respectively, we determine that

$$\text{prob}\{\gamma_L \leq X \leq \gamma_U\} = 1 - \rho.$$

Finally, since $z_{1-\rho/2} = -z_{\rho/2}$, we obtain ϵ_l and ϵ_u by arithmetic operations, which concludes the proof. ■

Note that the bounds ϵ_l and ϵ_u in Lemma 2 are functions of $\hat{\gamma}$, ρ , and N . For the sake of readability, we omit the parameters of these bounds. Further, note that Lemma 2 provides a potentially asymmetric confidence interval, i.e., $\epsilon_l \neq \epsilon_u$, depending on the value of $\hat{\gamma}$ and ρ .

Theorem 2: Assume $\gamma \in [0.5, 1)$ is a rational number, such that $\gamma = n_1/n_2$ is the irreducible fraction of γ , and that we have N i.i.d. samples of y_D , $\{y_{D,i}\}_{i=1}^N$. Under these assumptions, a solution to Problem 1 is given by $N = kn_2$, and (8), as shown at the bottom of the page, if $\tilde{J}_D = y_{D,(N\gamma)}$, where $z_{\rho/2}$ is the upper $\rho/2$ -quantile of the standard Gaussian distribution.

Proof: First, recall that $F(y_{D,(m)}) \sim \text{Beta}(m, N + 1 - m)$ from Lemma 1 and note that by setting $m = N\gamma$ we have $\hat{\gamma} = \gamma$, since $N = kn_2$. Next, by assuming that $\gamma \in [0.5, 1)$, we determine that $\epsilon_l \geq \epsilon_u$. This shows us that $[\gamma - \epsilon_l, \gamma + \epsilon_u] \subseteq [\gamma - \epsilon_l, \gamma + \epsilon_l]$ for $\gamma \in [0.5, 1)$. Therefore, with $\tilde{J}_D = y_{D,(N\gamma)}$, the true probability of triggering no alarm is given by $F(\tilde{J}_D)$ such that we obtain the confidence interval $\text{prob}(|F(\tilde{J}_D) - \gamma| \leq \epsilon_l) \geq 1 - \rho$ from Lemma 2. By setting

$N = kn_2$ and $\epsilon_l = \epsilon$ in Lemma 2 we solve for \sqrt{k} and obtain

$$\begin{aligned} \sqrt{k} &= \frac{z_{\frac{\rho}{2}} \sqrt{\gamma - \gamma^2}}{2\epsilon \sqrt{n_2}} \\ &\pm \sqrt{\left(\frac{z_{\frac{\rho}{2}} \sqrt{\gamma - \gamma^2}}{2\epsilon \sqrt{n_2}} \right)^2 + \frac{2(\gamma - 0.5)z_{\frac{\rho}{2}}^2}{3n_2\epsilon} + \frac{1 + \gamma}{3n_2\epsilon}}. \end{aligned} \quad (9)$$

Since $\sqrt{k} \geq 0$, we discard the solution with a negative sign. After squaring and rounding up to the next larger integer, we obtain (8). This concludes the proof. ■

Note that it is not restrictive to only consider $\gamma \in [0.5, 1)$ in Theorem 2. This is because

$$\begin{aligned} \text{prob}\{|F(\tilde{J}_D) - \gamma| \leq \epsilon\} &= \text{prob}\{|1 - F(\tilde{J}_D) - (1 - \gamma)| \leq \epsilon\} \\ &= \text{prob}\{|F(\hat{J}_D) - (1 - \gamma)| \leq \epsilon\} \end{aligned}$$

where \hat{J}_D is the threshold that approximates a false alarm rate of γ , which exists due to the continuity of the CDF. Thus, the integer k obtained from (8) for a certain $\gamma \in [0.5, 1)$ is the same as for $1 - \gamma$.

D. Discussion and Comparison of the Bounds

In this section we will discuss the three finite sample guarantees obtained previously, investigate their scaling in the parameters ϵ and ρ , and compare the sample sizes. For a given γ , ϵ , and ρ , let N_{DKW} , N_{VP} , and N_{beta} be the sample sizes obtained from Proposition 1, Theorem 1, and Theorem 2, respectively. First, we want to discuss how to choose ϵ . A reasonable bound for ϵ is $\epsilon \leq \min(\gamma, 1 - \gamma)$, because this choice results in $[\gamma - \epsilon, \gamma + \epsilon] \subseteq [0, 1]$. This means that the $\pm\epsilon$ -band around the acceptable false alarm rate contains only reasonable false alarm rates, i.e., false alarm rates in $[0, 1]$. Further, with that choice we have $(4/9)\gamma(1 - \gamma) \geq (4/9)\min(\gamma, 1 - \gamma)^2 \geq (4/9)\epsilon^2 > (1/6)\epsilon^2 \geq \rho\epsilon^2$ if $\rho \leq 1/6$. Hence, if $\rho \leq 1/6$ the condition on both ρ and ϵ in Theorem 1 is fulfilled for this bound of ϵ .

Next, we want to investigate the scaling of each guarantee in the parameters ϵ and ρ . All finite guarantees scale with ϵ^{-2} , while they differ in their scaling in ρ . The guarantee scales with $\ln(\rho^{-1})$, ρ^{-1} and $z_{\rho/2}$ when it is obtained from Proposition 1, Theorem 1, and Theorem 2, respectively. Typically, we desire high probability guarantees such that ρ is close to zero. It follows that Theorem 2 has the best scaling for $\rho \in (0, 0.5]$ since then $z_{\rho/2} < \ln(\rho^{-1}) < \rho^{-1}$.

Now, we compare the sample sizes from the three different finite guarantees when $\rho = 0.05$ and $\epsilon = 0.01$ for $\gamma \in [0.01, 0.99]$. Fig. 2 shows the results of the comparison. We see that the required sample size obtained from the DKW inequality is almost constant for different γ . This is expected because the DKW inequality applies for the whole cumulative distribution function. Further, we observe that Theorem 2

$$k = \left\lceil \left(\frac{z_{\frac{\rho}{2}} \sqrt{\gamma - \gamma^2}}{2\epsilon \sqrt{n_2}} + \sqrt{\left(\frac{z_{\frac{\rho}{2}} \sqrt{\gamma - \gamma^2}}{2\epsilon \sqrt{n_2}} \right)^2 + \frac{2(\gamma - 0.5)z_{\frac{\rho}{2}}^2}{3n_2\epsilon} + \frac{1 + \gamma}{3n_2\epsilon}} \right)^2 \right\rceil \quad (8)$$




Fig. 2. We compare the sample sizes obtained from the DKW inequality (dash-dotted line), from the Vysochanskij–Petunin inequality (dashed line) and from the beta confidence intervals (solid line), where $\gamma \in [0.01, 0.99]$, $\rho = 0.05$, and $\epsilon = 0.01$.

produces the smallest sample sizes out of the three finite guarantees for all investigated γ . We also observe that $N_{VP} \leq N_{DKW}$ if $\gamma \notin (0.294, 0.706)$. So only for γ close to 0 and 1 the finite guarantees from Theorem 1 perform better than the finite guarantees from Proposition 1.

Both Fig. 2 and the scaling in ρ and ϵ show us that the sample sizes we obtain from Theorem 2 are smaller than the sample sizes from Proposition 1 and Theorem 1 for all investigated γ . Furthermore, according to [12], the confidence bounds ϵ_u and ϵ_l are (nearly) exact if $N \geq 40$, since the influence of the higher order terms disappears.

IV. NUMERICAL EXAMPLES

In the first part of section, we evaluate the sampling guarantees numerically for three different detector output distributions. In the second part, we use detector output data from an experimental setup to tune the threshold of a CUSUM detector without resetting. The code to reproduce these results can be found at <https://github.com/DavidUmsonst/FiniteSampleGuaranteesForQuantileEstimation>.

A. Evaluation of Sample Guarantees

We begin by evaluating the finite sample guarantees. The idea is to approximate the threshold of the anomaly detector with (2) based on N samples, $\{y_{D,i}\}_{i=1}^N$, of the detector output and we do this approximation $N_T = 1000$ times to obtain N_T different approximations \tilde{J}_D of J_D . Then we draw 10^6 new samples of the detector output and calculate the empirical false alarm rate of each threshold.

Here, we investigate three different cases. First, we assume that the detector output has a $\chi^2(4)$ distribution, where we have four degrees of freedom. The samples are then taken i.i.d. from a $\chi^2(4)$ distribution. Second, we assume that the detector output has a Lévy distribution and obtain i.i.d. samples for the output from a Lévy distribution. Third, we assume the samples are taken from the trajectory of a non-parametric CUSUM detector [13] without resetting, given by

$$y_D(k+1) = \max(0, y_D(k) + \|r(k)\|_2^2 - \delta) \quad (10)$$

where $y_D(0) = 0$, $\delta = 6$, and $r(k)$ is the input of the detector and is drawn i.i.d. from a four-dimensional, zero-mean, multivariate Gaussian distribution at each time step k .

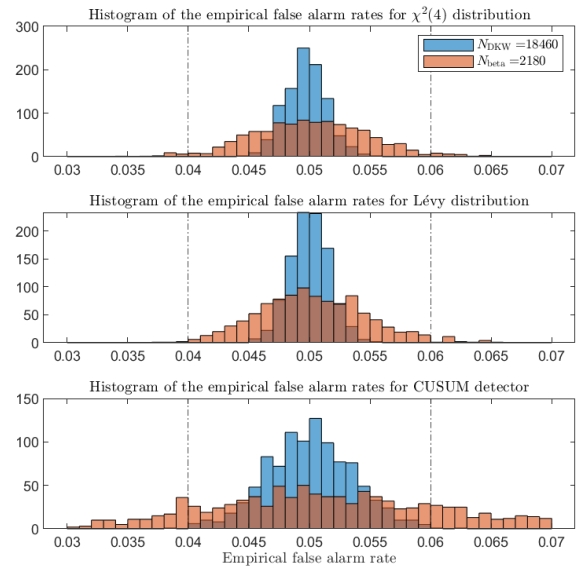


Fig. 3. Empirical false alarm rate from $N_T = 1000$ thresholds is evaluated over a dataset of 10^6 samples for a $\chi^2(4)$ distribution (upper plot), a Lévy distribution (center plot), and samples obtained from a CUSUM detector (lower plot).

Furthermore, we use $y_{D,i} = y_D(i)$ to obtain the samples, that is, the samples are the trajectory of the CUSUM detector. Therefore, the i.i.d. assumption on the samples is not fulfilled in this case.

For our simulation, we choose $\gamma = 0.95$, $\epsilon = 0.01$, and $\rho = 0.05$, which means that the empirical false alarm rate should be in the interval $[0.04, 0.06]$ with a probability of 95%. Moreover, Section III-D showed us that for these values of γ , ϵ , and ρ , $N_{\text{beta}} < N_{VP} < N_{DKW}$. Therefore, we investigate only the smallest and largest sample sizes, N_{beta} and N_{DKW} , respectively, for the sake of clarity. In addition to that, we set $\beta = 0$ when approximating the threshold with (2).

Fig. 3 shows the histogram of the empirical false alarm rate for the $\chi^2(4)$ distribution (upper plot), the Lévy distribution (center plot), and samples obtain from the CUSUM detector (lower plot) for each of the two sample sizes investigated. The two upper plots confirm our theoretical results, where the empirical false alarm rate lies inside the desired interval with a probability larger than 95%. More specifically, when $N = N_{\text{beta}}$ and y_D has a χ^2 distribution only 4.2%, that is, 42 out of 1000, of the empirical false alarm rates are outside of the desired interval. Similarly, when $N = N_{\text{beta}}$ and y_D has a Lévy distribution only 2.4% of the empirical false alarm rates are outside of the desired interval. If $N = N_{DKW}$ none of the empirical false alarm rates are outside of the desired interval. Therefore, we see that the sample size provided by Theorem 2 are very close to the desired guarantees of at most 5% of false alarm rates being outside the desired interval, while the sample size provided by Proposition 1 has much better probabilistic guarantees but has also a conservative amount of samples. Furthermore, the two upper plots in Fig. 3 verify that the guarantees hold for both light-tailed and heavy-tailed distributions.

However, in the lower plot, where the samples are not i.i.d., we observe that 34.1% of the empirical false alarm

rates lie outside of the desired interval when $N = N_{\text{beta}}$ is used. For $N = N_{\text{DKW}}$ only 0.3% of the empirical false alarm rates lie outside the desired interval. Hence, we see that the probabilistic guarantees are not fulfilled, when $N = N_{\text{beta}}$. Although in this case, the choice of $N = N_{\text{DKW}}$ provides enough conservatism to outweigh the effect of the non-i.i.d. sampling, in principle, there is no guarantee in either case. Therefore, the sample size obtained in Theorem 2 is the smallest sample size that fulfills the probabilistic guarantees given N_{beta} i.i.d. samples, but it is sensitive to violations of the i.i.d. assumption.

One reasonable approach to deal with correlated data is to collect more data than the minimum amount for i.i.d. data and randomly sample an approximately i.i.d. subset of the correlated data. Proposition 1, Theorem 1, and Theorem 2 can help with determining the size of the approximately i.i.d. subset of correlated data, but not the minimum size of the correlated data needed. This approach of dealing with correlated data is used in the next section, when we obtain samples from our experimental setup. Determining the minimum size of the correlated dataset for threshold estimation needs further investigation and is an avenue for future work.

B. Tuning With Real Data

In this last section, we evaluate the sample guarantees provided by Theorem 2 with data obtained from an experimental setup, where a real process is controlled. The process used is the Temperature Control Lab (TCLab), which consists of two heaters and one temperature sensor for each heater (for more details see [22]). In the experimental set up, we control both heaters to have a temperature of 40°C. We implement an LQG controller, which is based on a linearized and discretized model around the steady-state temperature 40°C for both heaters. Here we use a sampling time of 1 s to obtain measurements. From the TCLab's sensors we obtain two measurements, $y(k) \in \mathbb{R}^2$, and with the predicted sensor measurements $\hat{y}(k)$ from the steady-state Kalman filter inside the LQG controller the residual signal is $\bar{r}(k) = y(k) - \hat{y}(k)$. Here, $r(k) = \hat{\Sigma}^{-1/2}(\bar{r}(k) - \hat{\mu})$ serves as the input to (10), where $\hat{\mu}$ and $\hat{\Sigma}$ are approximations of the mean and the covariance matrix of the residual signal, respectively, used to normalize $\bar{r}(k)$, $\delta = 3$ in (10), and the initial state of the CUSUM detector is again set to zero. Note that, to the best of our knowledge, there exists no closed-form solution for the threshold of a CUSUM detector that guarantees an acceptable false alarm rate $1 - \gamma$. Therefore, we will use N_{beta} samples of the detector to approximate the threshold according to (2) with $\beta = 0$ for an acceptable false alarm rate of $1 - \gamma$.

We let the experiment run for 15 379 s to gather data. Since the detector output is assumed to be a random variable with a fixed distribution (see Assumption 1), we only work with samples from the steady state of the system, which is approximately reached for $k \geq 780$. We obtain $\hat{\mu}$ and $\hat{\Sigma}$ from the first 1000 samples of the residual signal in steady state, i.e., the dataset $\mathcal{D}_{\mu\Sigma} = \{r(k)\}_{k=780}^{1779}$. With $\hat{\mu}$ and $\hat{\Sigma}$ available we determine the detector output $y_D(k+1)$ such that we use the dataset $\mathcal{D} = \{y_D(k)\}_{k=1779}^{15379}$ of detector outputs to estimate




Fig. 4. Box plot of the empirical false alarm rate over the acceptable false alarm rate $1 - \gamma$ when $N_T = 10\,000$ threshold estimates with corresponding their empirical false alarm rates are obtained from the dataset \mathcal{D} via repeated training/test splits.

the threshold, where $y_D(1779) = 0$, and validate the empirical false alarm rate of the threshold estimate.

In the following, we apply the approach outlined at the end of the previous section. This means that, for a given γ , we randomly choose N_{beta} samples from \mathcal{D} without replacement to obtain an approximately i.i.d. dataset for determining \tilde{J}_D . Since we only have a finite amount of samples in \mathcal{D} , the remaining samples of \mathcal{D} are compared to the threshold to produce the empirical false alarm rate induced by \tilde{J}_D . Repeatedly applying this approach is known as repeated training/test splits [23], which lets us simultaneously estimate and evaluate \tilde{J}_D with a finite dataset \mathcal{D} . This is done repeatedly to obtain $N_T = 10\,000$ empirical false alarm rates for a given γ . For tuning the detector, we choose $\epsilon = 0.01$ and $\rho = 0.05$ and investigate nine different values for γ , that is, $\{\gamma_i\}_{i=1}^9$, where $\gamma_i = 0.95 + (i - 1) \cdot 0.005$. For these values of ϵ and ρ , the amount of samples for testing is at least 11 421 for each of the investigated values of γ .

Fig. 4 shows a box plot of the empirical false alarm rates over the investigated acceptable false alarm rates $1 - \gamma$. We observe that the median value of the empirical false alarm rate is almost exactly located at the acceptable false alarm rate for all investigated values of γ . Furthermore, the box plots are concentrated around the acceptable false alarm rate. Only 5% of the empirical false alarm rates should lie outside of the $\pm\epsilon$ -band (shaded area in Fig. 4) around the acceptable false alarm rate, since $\rho = 0.05$. Here, the largest percentage of empirical false alarm rates outside the $\pm\epsilon$ -band is 4.77% for an acceptable false alarm rate of 0.045, i.e., $\gamma = 0.955$. Hence, although the data in \mathcal{D} is strongly correlated due the CUSUM detector dynamics (10), we achieve the desired guarantees with the number of samples determined from Theorem 2. The reason for that is that due to the random selection of N_{beta} samples from \mathcal{D} to determine \tilde{J}_D , it is unlikely that several adjacent samples, e.g., $y_D(2000)$, $y_D(2001)$, and $y_D(2002)$, are chosen such that the samples in the training set are not highly correlated anymore.




Fig. 5. Empirical false alarm rate from $N_T = 10\,000$ thresholds is evaluated over a dataset \mathcal{D} when repeated training/test splits are used to determine \hat{J}_D , where the area between the vertical dash-dotted lines represents the desired $\pm\epsilon$ -band.

Note that the samples in the test set are still strongly correlated. This shows that the guarantees for the threshold determined from the uncorrelated samples holds for the correlated samples in the test set as well.

Finally, we look at the histogram of the empirical false alarm rates for $\gamma = 0.95$ (see Fig. 5). We observe that in contrast to the lower plot in Fig. 3 the histogram for $N = N_{\text{beta}}$ is now concentrated around $1 - \gamma = 0.05$ and only 4.64% of the empirical false alarm rates are located outside of the desired interval marked by the vertical dash-dotted lines. This demonstrates that by selecting training data using random sampling (as opposed to sequential sampling) we are able to employ the sample guarantees in Theorem 2 to find a threshold for the highly correlated output of the CUSUM detector.

V. CONCLUSION

In this work, we considered the tuning of detector thresholds and pointed out the equivalence of the detector threshold and a specific quantile of the detector output distribution. We derived three different finite guarantees for the estimation of a quantile. The first is based on the DKW inequality, which takes the whole cumulative density function into account. The second is based on the Vysochanskij–Petunin inequality and uses the expected value and variance of the CDF evaluated at a specific order statistic to determine the sample guarantees. The third is based on the confidence interval of a beta distribution and utilizes a closed-form solution of the confidence interval bounds.

When comparing the three guarantees, we saw that the third guarantee has the best scaling in the confidence parameter and leads to the smallest sample size. Simulations showed that the i.i.d. assumption is important and can lead to violations of the guarantees. However, we showed in our experimental setup that using random instead of sequential samples to tune a threshold can be an effective way to mitigate the adverse effects of a non-i.i.d. detector output without changing the sequential implementation of the detector operation.

Avenues for future work involve relaxing the assumptions we have made in this work. We aim to develop extended results to directly take the non-i.i.d. nature of the detector output into

account rather than to use the indirect random sampling approach we propose here. We would also like to be able to address non-stationary detector output distributions, including results that provide guidance on real-time threshold selection.

REFERENCES

- [1] J. Minkel. (Aug. 2008). *The 2003 Northeast Blackout—Five Years Later*. Accessed: Aug. 12, 2020. [Online]. Available: <https://www.scientificamerican.com/article/2003-blackout-five-years-later/>
- [2] R. M. Lee, M. J. Assante, and T. Conway, “Analysis of the cyber attack on the Ukrainian power grid. Defense use case,” in *Proc. E-ISAC*, 2016, pp. 1–29.
- [3] J. W. R. McIntyre, “Ergonomics: Anaesthetists’ use of auditory alarms in the operating room,” *Int. J. Clin. Monitor. Comput.*, vol. 2, no. 1, pp. 47–55, Mar. 1985.
- [4] C. D. Brown and H. T. Davis, “Receiver operating characteristics curves and related decision measures: A tutorial,” *Chemometric Intell. Lab. Syst.*, vol. 80, no. 1, pp. 24–38, Jan. 2006.
- [5] D. I. Urbina *et al.*, “Limiting the impact of stealthy attacks on industrial control systems,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Oct. 2016, pp. 1092–1105.
- [6] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Comput. Surv.*, vol. 41, no. 3, pp. 1–58, Jul. 2009.
- [7] N. Goix, “How to evaluate the quality of unsupervised anomaly detection algorithms?” 2016, *arXiv:1607.01152*.
- [8] R. J. Hyndman and Y. Fan, “Sample quantiles in statistical packages,” *Amer. Statist.*, vol. 50, no. 4, pp. 361–365, Nov. 1996.
- [9] D. Umsonst, J. Ruths, and H. Sandberg, “Sample-based anomaly detector tuning with finite sample guarantees,” in *Proc. Amer. Control Conf. (ACC)*, May 2021, pp. 3248–3253.
- [10] P. Massart, “The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality,” *Ann. Probab.*, vol. 18, no. 3, pp. 1269–1283, 1990.
- [11] F. Pukelsheim, “The three sigma rule,” *Amer. Statist.*, vol. 48, no. 2, pp. 88–91, 1994.
- [12] M. Thulin, “The cost of using exact confidence intervals for a binomial proportion,” *Electron. J. Statist.*, vol. 8, no. 1, pp. 817–840, 2014.
- [13] C. Murguia and J. Ruths, “CUSUM and chi-squared attack detection of compromised sensors,” in *Proc. IEEE Conf. Control Appl. (CCA)*, Sep. 2016, pp. 474–480.
- [14] V. Renganathan, N. Hashemi, J. Ruths, and T. H. Summers, “Higher-order moment-based anomaly detection,” *IEEE Control Syst. Lett.*, vol. 6, pp. 211–216, 2022.
- [15] D. Li and S. Martinez, “High-confidence attack detection via Wasserstein-metric computations,” *IEEE Control Syst. Lett.*, vol. 5, no. 2, pp. 379–384, Apr. 2021.
- [16] S. Cléménçon and J. Jakubowicz, “Scoring anomalies: A M-estimation formulation,” in *Proc. 16th Int. Conf. Artif. Intell. Statist.*, vol. 31, C. M. Carvalho and P. Ravikumar, Eds. Scottsdale, AZ, USA: PMLR, May 2013, pp. 659–667.
- [17] R. Rockafellar and S. Uryasev, “Conditional value-at-risk for general loss distributions,” *J. Banking Finance*, vol. 26, no. 7, pp. 1443–1471, Jul. 2002.
- [18] A. Okolewski and T. Rychlik, “Sharp distribution-free bounds on the bias in estimating quantiles via order statistics,” *Statist. Probab. Lett.*, vol. 52, no. 2, pp. 207–213, Apr. 2001.
- [19] H. A. David and H. N. Nagaraja, *Order Statistics*, 3rd ed. Hoboken, NJ, USA: Wiley, 2003.
- [20] *Handbook of Beta Distribution and its Applications*, Statistics: Textbooks and Monographs, Marcel Dekker, New York, NY, USA, 2004, vol. 174.
- [21] A. Papoulis, *Probability, Random Variables, and Stochastic Processes* (McGraw-Hill Series in Electrical Engineering Communications and Signal Processing), 3rd ed. New York, NY, USA: McGraw-Hill, 1991.
- [22] J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren, “Benchmark temperature microcontroller for process dynamics and control,” *Comput. Chem. Eng.*, vol. 135, Apr. 2020, Art. no. 106736.
- [23] M. Kuhn and K. Johnson, *Over-Fitting Model Tuning*. New York, NY, USA: Springer, 2013, pp. 61–92.