

Generalized chi-squared detector for LTI systems with non-Gaussian noise

Navid Hashemi¹ and Justin Ruths¹

Abstract—Previously, we derived exact relationships between the properties of a linear time-invariant control system and properties of an anomaly detector that quantified the impact an attacker can have on the system if that attacker aims to remain stealthy to the detector. A necessary first step in this process is to be able to precisely tune the detector to a desired level of performance (false alarm rate) under normal operation, typically through the selection of a threshold parameter. To-date efforts have only considered Gaussian noises. Here we generalize the approach to tune a chi-squared anomaly detector for noises with non-Gaussian distributions. Our method leverages a Gaussian Mixture Model to represent the arbitrary noise distributions, which preserves analytic tractability and provides an informative interpretation in terms of a collection of chi-squared detectors and multiple Gaussian disturbances.

I. INTRODUCTION

Model-based anomaly detection uses a predictor to forecast the evolution of a dynamical system. This prediction is compared with the actual measured value to test for an appreciable discrepancy, which may indicate the presence of anomaly in the system. If the predictor is perfect, then this task is easy - any discrepancy is enough to cause concern. The task is made more challenging because the predictor has some uncertainty in its forecast, which implies that relatively small discrepancies may be due to this uncertainty rather than anomalous behavior. One of the most fundamental sources of uncertainty in conventional control systems is captured in the terms that quantify system and measurement noise. The overwhelming majority of work on stochastic model-based detection assumes system and measurement noises that are Gaussian distributed, see e.g., [1], [2], [3]. In the context of linear time-invariant systems especially, normal distributions preserve the capability of finding analytic solutions. In this work, we do not assume any structure of the noise probability density functions; however, we manage to retain some analytic tractability by employing a Gaussian Mixture Model (GMM) representation of these arbitrary distributions.

In a stochastic context, detectors necessarily trade off sensitivity (rate of true positives - true alarms) for fall-out (rate of false positives - false alarms). The machinery we present to tune detectors is a necessary step to (a) understand the shape of the Receiver Operating Characteristic (ROC) curve, (b) quantify how that shape depends on system and design parameters, and (c) select a point on the ROC that has a desired level of performance - as measured by a

tolerable true/false positive balance. In past work we have provided methods to tune several popular detectors when the uncertainty is driven by Gaussian system and measurement noises [4], [5]. When noise distributions deviate significantly from this assumption, then new tools are required to produce an effective means of tuning detectors. Here we provide a generalized chi-squared detector that can be tuned to any desired rate of false alarms through the appropriate selection of the detector sensitivity threshold.

The process of tuning model-based detectors is to propagate the uncertainty, here the noise distributions, through the system and into the detector. By quantifying the distributions the detector expects to see under normal operation (with no anomalies), we can quantify the trade-off between the true and false positive rates based on the sensitivity threshold we select. We begin by reviewing the tools for the Gaussian case and conclude by demonstrating the backwards compatibility of the results we develop with the Gaussian case and by showing an example.

II. MODEL-BASED ANOMALY DETECTION

Consider a general discrete-time linear time-invariant (LTI) system,

$$\begin{aligned} x_{k+1} &= Fx_k + Gu_k + v_k \\ y_k &= Cx_k + \eta_k \end{aligned} \quad (1)$$

with time index $k \in \mathbb{N}$, state $x_k \in \mathbb{R}^n$, output $y_k \in \mathbb{R}^p$, input $u_k \in \mathbb{R}^m$, matrices F , G , and C of appropriate dimensions, and iid multivariate noises $v_k \in \mathbb{R}^n$ and $\eta_k \in \mathbb{R}^p$ with covariance matrices $R_1 \in \mathbb{R}^{n \times n}$, $R_1 \geq 0$ and $R_2 \in \mathbb{R}^{p \times p}$, $R_2 \geq 0$ respectively. The random processes v_k and η_k are mutually independent. We assume that (F, G) is stabilizable and (F, C) is detectable such that there are no unstable unobservable or uncontrollable modes.

The main idea behind fault detection theory is to use an estimator to forecast the evolution of the system in the absence of fault [6]. In this analysis we consider a model-based estimator of Luenberger form,

$$\hat{x}_{k+1} = F\hat{x}_k + Gu_k + L(y_k - C\hat{x}_k), \quad (2)$$

where the estimator has perfect model knowledge, L is the observer gain matrix, and \hat{x}_{k+1} is the predicted state. This prediction is compared to the actual measurement received from the sensors. If the difference between the measurement and the predicted output is large, there may be an anomaly in the system. This difference is typically called the residual,

$$r_k = y_k - C\hat{x}_k, \quad (3)$$

¹These authors are with the Departments of Mechanical and Systems Engineering at the University of Texas at Dallas, Richardson, Texas, USA navid.hashemi@utdallas.edu, jruths@utdallas.edu

where y_k is the actual measurement and $C\hat{x}_k$ is the estimated output. Different model-based detectors use the residual in different ways to quantify deviation away from the estimate. One of the simplest and most widely used approach constructs the distance measure z_k from a quadratic form of the residual,

$$z_k = (r_k - \mu)^T \Sigma^{-1} (r_k - \mu), \quad z_k > \alpha \longrightarrow \text{alarm} \quad (4)$$

where $\mu \in \mathbb{R}^p$ and $\Sigma \in \mathbb{R}^{p \times p}$ are the mean and covariance of the residual random vector, r_k , under normal operation (no faults or attacks). This detector raises alarms if z_k exceeds an assigned threshold, i.e., $z_k > \alpha$, $\alpha \in \mathbb{R}_{>0}$.

When the system and sensor noises are zero-mean and Gaussian distributed, then the residual is also a zero-mean Gaussian random variable, i.e., $r_k \sim \mathcal{N}(0, \Sigma)$, with covariance $\Sigma = CPC^T + R_2$, where P represents the asymptotic covariance of the estimation error, $\lim_{k \rightarrow \infty} P_k = \lim_{k \rightarrow \infty} E[e_k e_k^T] = P$, $e_k = x_k - \hat{x}_k$, and is the solution of the Lyapunov equation [7],

$$(F - LC)P(F - LC)^T - P + R_1 + LR_2L^T = 0. \quad (5)$$

The distance measure z_k is then a sum of squared Gaussian variables making it chi-squared distributed. For this reason, the detector (4) is conventionally called the *Chi-Squared Detector*. This name underscores the prevailing assumption that system and sensor noises are Gaussian distributed - the simplifying assumption we lift in this work.

In past work, we characterized the relationship to select the threshold α in (4) to yield a desired false alarm rate. Effectively capturing the receiver operator curve (ROC) of the detector, this characterization is critical for comparisons across detectors because all detectors must be tuned for comparable performance. The following Lemma states the relationship for a chi-squared detector with zero-mean Gaussian system and measurement noises.

Lemma 1 ([5]): Assume that there are no anomalies present in an LTI system (1) driven by zero-mean Gaussian system and measurement noises such that the residual $r_k \sim \mathcal{N}(0, \Sigma)$ and consider the chi-squared detector (4) with threshold $\alpha \in \mathbb{R}_{>0}$. To achieve a desired false alarm rate \mathcal{A}^* set $\alpha = \alpha^* := 2P^{-1}(1 - \mathcal{A}^*, \frac{p}{2})$, where $P^{-1}(\cdot, \cdot)$ denotes the inverse regularized lower incomplete gamma function.

III. NON-GAUSSIAN NOISES

Although assuming noises fall according to Gaussian distributions is standard for the theoretical treatment of control systems, practical implementations for tuning chi-squared detectors on real systems can fall short when using Lemma 1 directly. System noise is typically used to aggregate model uncertainty. For many control concerns, increasing the covariance of Gaussian noise is an effective, conservative approach for capturing model imperfections. However, when we aim to use this noise to predict the expected behavior of the system it is important that the noise distribution is an accurate representation of the model discrepancy. In addition, injecting conservatism into the system noise directly reduces the sensitivity of the detector. For sensor noise, many sensors

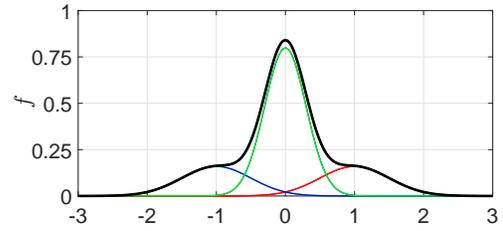


Fig. 1: Given the probability density function (black) a Gaussian Mixture Model (GMM) expression can be constructed. Here the GMM is composed of three Gaussian modes.

exhibit nonlinear behavior in the distribution of their reported values or evidence strong quantization effects; both of which can have a dramatic effect on the accuracy of Lemma 1.

We revisit the chi-squared detector defined in (4) with system and sensor noises that are arbitrary distributed. These distributions can be expressed as convex mixtures of Gaussian distributions [8]. We construct the so-called *Gaussian mixture model* (GMM) expansion of the distributions for system noise (v_k) and measurement noise (η_k), respectively,

$$f_\eta = \sum_{j=1}^{m_1} p_j^\eta \mathcal{N}(x | \mu_j^\eta, K_j^\eta), \quad (6)$$

$$f_v = \sum_{j=1}^{m_2} p_j^v \mathcal{N}(x | \mu_j^v, K_j^v), \quad (7)$$

where m_1 is the number of Gaussian modes for the measurement noise and m_2 is the number of Gaussian modes for the system noise, $\mu_j^\eta \in \mathbb{R}^p$ and $\mu_j^v \in \mathbb{R}^n$ are the mean value of each Gaussian mode, $K_j^\eta \in \mathbb{R}^{p \times p}$ and $K_j^v \in \mathbb{R}^{n \times n}$ are their corresponding covariances. It is important that the scalar values p_j^η and p_j^v satisfy

$$\sum_{j=1}^{m_1} p_j^\eta = 1, \quad \text{and} \quad \sum_{j=1}^{m_2} p_j^v = 1, \quad (8)$$

which follows from interpreting (6) (resp., (7)) as a total law of probability between Gaussian conditional probabilities with probability p_j^η (resp., p_j^v). The EM algorithm supplies a reliable mechanism for computing GMMs from data with guarantees on the convergence of these models [9]. See Fig. 1 for an example of a GMM.

The challenge, now, for anomaly detection is to accurately characterize the distribution of the residual and distance measure, ideally maintaining an analytic relationship so that the role of different design variables (e.g., observer gain) can be understood. This paper employs a generalized version of the chi-squared detector and in using the GMM representation of the noise distributions allows us to develop the GMM representation of the residual distribution and then compute statistics about the distance measure distribution.

A. Residual Distribution

We first characterize the residual distribution.

Lemma 2: Given the LTI system (1) driven by system and sensor noises whose probability density functions can be expressed as the Gaussian mixture models in (6) and (7), the probability density function of the residual r_k at time $k \in \mathbb{N}$ can be written as a Gaussian mixture model of $m_k = m_1^k m_2^{k-1}$ Gaussian modes

$$f_{r_k}(x) = \sum_{j=1}^{m_k} \tau_j \mathcal{N}(x | \beta_j, \Theta_j), \quad (9)$$

where τ_j represents the mixture probabilities of the Gaussian modes, β_j are the means, and Θ_j are the covariances,

$$\beta_j = \sum_{\kappa=1}^k A_\kappa^T \mu_{n_\kappa}^\eta + \sum_{\kappa=1}^{k-1} B_\kappa^T \mu_{n_{\kappa+k}}^v$$

$$\Theta_j = \sum_{\kappa=1}^k A_\kappa K_{n_\kappa}^\eta A_\kappa^T + \sum_{\kappa=1}^{k-1} B_\kappa K_{n_{\kappa+k}}^v B_\kappa^T$$

$$\tau_j = \prod_{\kappa=1}^k p_{n_\kappa}^\eta \times \prod_{\kappa=1}^{k-1} p_{n_{\kappa+k}}^v$$

$$A_\kappa = \begin{cases} I & \kappa = 1 \\ -C(F - LC)^{\kappa-2}L & 2 \leq \kappa \leq k \end{cases}$$

$$B_\kappa = C(F - LC)^{\kappa-1} \quad 1 \leq \kappa \leq k-1$$

where n_κ^j captures all the possible permutations of combining terms from the system and measurement noise GMM modes, following the rules: leftmargin=1.3cm

Init: $n_\kappa^1 = 1$ for $\kappa = 1, \dots, 2k-1$

Add: $n_1^{j+1} = n_1^j + 1$

Wrap: if $n_\kappa^{j+1} > \begin{cases} m_1 & \text{for } \kappa = 1, 2, \dots, k \\ m_2 & \text{for } \kappa = k+1, \dots, 2k-1 \end{cases}$

then $n_{\kappa+1}^{j+1} = n_{\kappa+1}^j + 1$ and $n_\kappa^{j+1} = 1$.

Remark 1: Since any probability density function can be expressed as a Gaussian mixture model, the key insight from Lemma 2 is not that the residual can be expressed as a GMM, but rather that the residual GMM can be found analytically from the GMMs of the system and measurement noises. The notation above is unavoidably complicated due to a large number of terms. Their definitions are straightforward by looking at the proof below - namely the expansion of the product in (16). In Section V we will discuss an effective means of reducing the number of GMM terms needed to represent the density function.

Proof: In the absence of anomalies, from the definition of the residual (3) and the system equations (1), we can recursively solve for the expression of r_k as a function of only the system and sensor noises (alternatively use the z-transform),

$$r_k = \eta_k - \sum_{\kappa=1}^{k-1} C(F - LC)^{\kappa-1} (L\eta_{k-\kappa} - v_{k-\kappa}). \quad (10)$$

Thus the residual is a linear combination of sequential iid samples of noise such that its distribution can be expressed

as a convolution integral in terms of the constituent noise distributions [10], i.e.,

$$f_{r_k} = f_{A_k \eta_1} * \dots * f_{A_1 \eta_k} * f_{B_{k-1} v_1} * \dots * f_{B_1 v_{k-1}} \quad (11)$$

where $A_\kappa = -C(F - LC)^{\kappa-2}L$, $2 \leq \kappa \leq k$, $A_1 = I$ and $B_\kappa = C(F - LC)^{\kappa-1}$, $1 \leq \kappa \leq k-1$. Since the noise samples are iid, we can drop the time dependence to yield,

$$f_{r_k} = f_{A_1 \eta} * \dots * f_{A_k \eta} * f_{B_1 v} * \dots * f_{B_{k-1} v}. \quad (12)$$

The characteristic function of a random variable X is defined $\varphi_X(\omega) = \mathbb{E}[e^{i\omega^T X}]$, where $\mathbb{E}[\cdot]$ denotes expectation. Using the characteristic function to represent the distributions involved in (12), changes the convolution from an integral to a product. The characteristic functions of the residual and noises are,

$$\varphi_{r_k}(\omega) = \int_{-\infty}^{\infty} f_{r_k}(x) e^{i\omega^T x} dx, \quad (13)$$

$$\varphi_\eta(\omega) = \int_{-\infty}^{\infty} f_\eta(x) e^{i\omega^T x} dx, \quad (14)$$

$$\varphi_v(\omega) = \int_{-\infty}^{\infty} f_v(x) e^{i\omega^T x} dx. \quad (15)$$

Note that the characteristic function of an affine transformation of a random variable $Y = QX + R$ is

$$\varphi_Y(\omega) = e^{i\omega^T R} \varphi_X(Q^T \omega),$$

thus $\varphi_{A_\kappa \eta}(\omega) = \varphi_\eta(A_\kappa^T \omega)$ and $\varphi_{B_\kappa v}(\omega) = \varphi_v(B_\kappa^T \omega)$. Therefore, $\varphi_{r_k}(\omega)$ is a product of characteristic functions for transformed system and measurement noises,

$$\varphi_{r_k}(\omega) = \prod_{\kappa=1}^k \varphi_\eta(A_\kappa^T \omega) \times \prod_{\kappa=1}^{k-1} \varphi_v(B_\kappa^T \omega). \quad (16)$$

Using the Gaussian mixture model with the characteristic function representation makes it possible to expand the system noise and measurement noise, e.g., for the measurement noise

$$\varphi_\eta(\omega) = \sum_{j=1}^{m_1} p_j^\eta \int_{-\infty}^{\infty} \mathcal{N}(x | \mu_j^\eta, K_j^\eta) e^{i\omega^T x} dx, \quad (17)$$

$$= \sum_{j=1}^{m_1} p_j^\eta e^{(i\omega^T \mu_j^\eta - \frac{1}{2} \omega^T K_j^\eta \omega)}, \quad (18)$$

where the simplification from (17) to (18) comes from identifying that the integral in (17) is the characteristic function of a Gaussian distribution which has a closed form expression [10]. If we replace the variable ω with its transformed $A_\kappa^T \omega$ and perform the same process for the system noise,

$$\varphi_\eta(A_\kappa^T \omega) = \sum_{j=1}^{m_1} p_j^\eta e^{(i\omega^T A_\kappa \mu_j^\eta - \frac{1}{2} \omega^T A_\kappa K_j^\eta A_\kappa^T \omega)}, \quad (19)$$

$$\varphi_v(B_\kappa^T \omega) = \sum_{j=1}^{m_2} p_j^v e^{(i\omega^T B_\kappa \mu_j^v - \frac{1}{2} \omega^T B_\kappa K_j^v B_\kappa^T \omega)}. \quad (20)$$

Substituting these expressions into (16) reveals that $\varphi_{r_k}(\omega)$ is a linear combination of $m_k = m_1^k m_2^{k-1}$ exponential terms

$$\varphi_{r_k}(\omega) = \sum_{j=1}^{m_k} \tau_j e^{(i\omega^T \beta_j - \frac{1}{2}\omega^T \Theta_j \omega)} \quad (21)$$

where the expressions for β_j , Θ_j , and τ_j in the statement of Lemma 2 can be derived by substituting (19)-(20) and expanding the product in (16). Conveniently, (21) is in the form of the characteristic function of a linear combination of Gaussian functions, thus (21) demonstrates that the residual distribution at time step k can be expressed as

$$f_{r_k}(x) = \sum_{j=1}^{m_k} \tau_j \mathcal{N}(x | \beta_j, \Theta_j), \quad (22)$$

where τ_j represents the mixture probabilities of the m_k Gaussian modes, β_j are the means of the modes, and Θ_j are the covariances of the modes. ■

In the absence of anomalies, the time dependence of the residual is governed by the convergence of the estimator, since the system is time-invariant and the noises are iid. In most practical situations the estimator is designed to converge relatively quickly and so it is reasonable to assume for the rest of this work that sufficient convergence of the estimation has already been achieved. Thus we seek the steady state distribution of the residual, which permits some simplification by removing the dependence on time.

Lemma 3: If the system is stable, then the distribution of the residual converges to f_{r_∞} as $k \rightarrow \infty$, i.e., given some error tolerance $\epsilon \in \mathbb{R}_{>0}$ there exists a $k^* \in \mathbb{N}$ such that

$$\|f_{r_\infty} - f_{r_{k^*}}\| < \epsilon, \quad (23)$$

implying that $f_r := f_{r_{k^*}}$ provides an arbitrarily close approximation of the steady state distribution.

Proof: To not belabor an intuitive result, we provide a sketch of the proof without explicitly writing all details regarding the convergence. Because the system is stable $\rho(F - LC) < 1$, which means

$$\lim_{\kappa \rightarrow \infty} A_\kappa = O_{p \times p}, \quad \text{and} \quad \lim_{\kappa \rightarrow \infty} B_\kappa = O_{n \times p},$$

where O is the zero matrix of the designated size. Therefore, in steady state the characteristic functions characterizing the system and measurement noise become

$$\begin{aligned} \lim_{\kappa \rightarrow \infty} \varphi_\eta(A_\kappa^T \omega) &= \lim_{\kappa \rightarrow \infty} \sum_{j=1}^{m_1} p_j^\eta e^{(i\omega^T A_\kappa \mu_j^\eta - \frac{1}{2}\omega^T A_\kappa K_j^\eta A_\kappa^T \omega)} \\ &= \sum_{j=1}^{m_1} p_j^\eta = 1, \end{aligned} \quad (24)$$

$$\begin{aligned} \lim_{\kappa \rightarrow \infty} \varphi_v(B_\kappa^T \omega) &= \lim_{\kappa \rightarrow \infty} \sum_{j=1}^{m_2} p_j^v e^{(i\omega^T B_\kappa \mu_j^v - \frac{1}{2}\omega^T B_\kappa K_j^v B_\kappa^T \omega)} \\ &= \sum_{j=1}^{m_2} p_j^v = 1. \end{aligned} \quad (25)$$

The convergence of these characteristic functions imply the convergence of the characteristic function of the residual to $\varphi_{r_\infty}(\omega)$ such that for a given $\tilde{\epsilon} \in \mathbb{R}_{>0}$ there exists a k^* such that

$$|\varphi_{r_\infty}(\omega) - \varphi_{r_{k^*}}(\omega)| < \tilde{\epsilon}, \quad (26)$$

where φ_{r_k} is defined as in (16). This provides an approximation for the steady state that is made arbitrarily accurate by selecting an arbitrarily small $\tilde{\epsilon}$ (and hence large k^*),

$$\begin{aligned} \varphi_{r_\infty}(\omega) &= \prod_{\kappa=1}^{\infty} \varphi_\eta(A_\kappa^T \omega) \times \prod_{\kappa=1}^{\infty} \varphi_v(B_\kappa^T \omega) \\ &\approx \prod_{\kappa=1}^{k^*} \varphi_\eta(A_\kappa^T \omega) \times \prod_{\kappa=1}^{k^*-1} \varphi_v(B_\kappa^T \omega). \end{aligned} \quad (27)$$

As mentioned earlier, k^* can be interpreted as the settling time of the control system and estimator. As before, the characteristic function of the residual corresponds to a probability density function that is composed of the sum of multiple Gaussian modes - a GMM. This distribution can also be made arbitrarily accurate by selecting a smaller $\tilde{\epsilon}$ which corresponds to a larger k^* and a smaller ϵ in (23),

$$f_{r_\infty}(x) \approx f_r(x) := f_{r_{k^*}}(x) = \sum_{j=1}^m \pi_j \mathcal{N}(x | \mu_j, K_j), \quad (29)$$

where $m = m_1^{k^*} m_2^{k^*-1}$ is the number of Gaussian modes used to represent the steady state residual distribution. As before, the values of π_j , μ_j , and K_j are the mixture probabilities, means, and covariances, respectively, of the m Gaussian modes and are computed by substituting (19)-(20) and expanding the product in (28). ■

B. Distance Measure & False Alarm Rate Tuning

Most detectors construct a scalar-valued distance measure from the residual to quantify how different the measurement is from what is expected. In this paper we use the generalized chi-squared detector in (4) where μ is the overall mean value and Σ is the overall covariance of the steady state residual. From (29),

$$\mu = \sum_{j=1}^m \pi_j \mu_j \quad \text{and} \quad \Sigma = \sum_{j=1}^m \pi_j K_j + \gamma_j \gamma_j^T \quad (30)$$

where $\gamma_j = \mu - \mu_j$ is the difference between the individual GMM (mode) means and overall mean. Our aim in this section, and ultimately of this paper, is to characterize the expected rate of false alarms given a chosen threshold value, α , of the detector. Recall that alarms are generated if $z_k > \alpha$ for any $k \in \mathbb{N}$. Thus the probability of drawing a distance measure value from its distribution that leads to an alarm is $P(z_k > \alpha)$. Figure 2 plots the multivariate, scalar-valued distance measure function z_k over the domain of the residual r_k , in the two dimensional ($p = 2$) case. Due to the quadratic form of the distance measure in (4), the surface is a paraboloid. The region \mathcal{D} is the area contained by the projection of the level set $z_k = \alpha$ onto the r_k plane. Theorem 1 is the generalized version of Lemma 1 for tuning the

$$M_j = \frac{\det \mathcal{C}}{\sqrt{(2\pi)^p \det K_j}} \int_0^{2\pi} \int_0^\pi \cdots \int_0^\pi \int_0^{\sqrt{\alpha}} e^{(\mathcal{C}^T \rho_k + \gamma_j)^T K_j^{-1} (\mathcal{C}^T \rho_k + \gamma_j)} d\rho \quad (\star)$$

$$d\rho = |\rho_k|^{p-1} \sin(\phi_1)^{p-2} \sin(\phi_2)^{p-3} \cdots \sin(\phi_{p-2}) d|\rho_k| d\phi_1 d\phi_2 \cdots d\phi_{p-1}$$

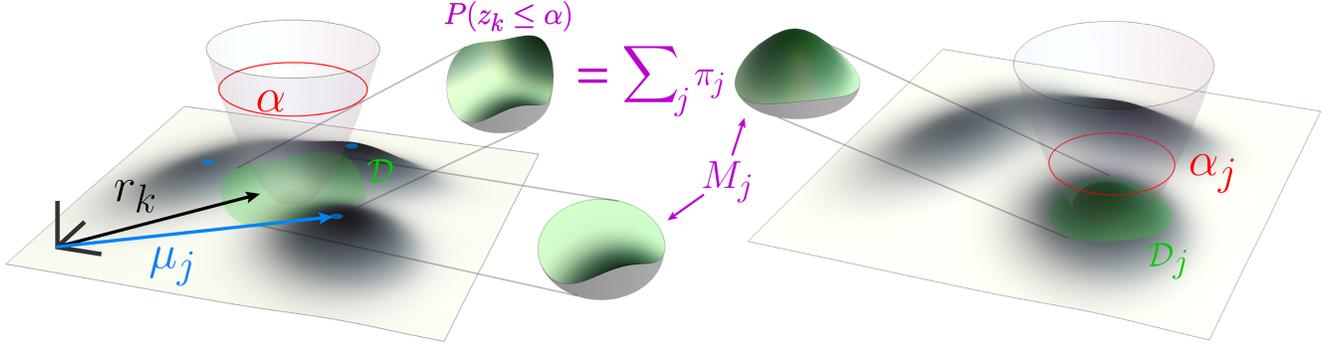


Fig. 2: The 3D surface of the residual GMM probability density function (shaded to reveal height mapping) constructed from $m = 3$ Gaussian modes with means μ_j denoted by blue dots. The quadratic form of the distance measure z_k forms a paraboloid over the domain of the residual r_k centered at μ (the mean of the residual distribution). The cumulative probability $P(z_k \leq \alpha)$ corresponds to the amount of probability in the residual distribution within the region \mathcal{D} (33) (green area), the area corresponding to the domain formed by projecting the $z_k = \alpha$ level set of the paraboloid (red line). The probability $P(z_k \leq \alpha) = \sum_j \pi_j M_j$, where M_j is the volume of probability contributed by the Gaussian mode j . In Section IV, we show that we can interpret the generalized detector as the probabilistic combination of m chi-squared detectors, where the j -th detector is tuned such that the false alarm rate is $\mathcal{A}_j = 1 - M_j$, which means the threshold α_j is selected such that the volume of probability contained within the region \mathcal{D}_j (38) (the projection of the α_j -level set of the paraboloid centered at μ_j) is equal to M_j .

detector for a desired level of performance (desired false alarm rate) in the case that the system and measurement noises are no longer Gaussian.

Theorem 1: Assume that there are no anomalies present in an LTI system (1) driven by arbitrary system and measurement noises such that the residual $r_k \sim f_r$ (29), i.e., an m -th order Gaussian mixture model, and consider the generalized chi-squared detector (4) with threshold $\alpha \in \mathbb{R}_{>0}$. The expected false alarm rate is

$$\mathcal{A} = 1 - P(z_k \leq \alpha) = 1 - \sum_{j=1}^m \pi_j M_j, \quad (31)$$

where M_j is given by (\star) , in the box above.

Proof: The false alarm rate is the probability that z_k exceeds the threshold α , $\mathcal{A} = P(z_k > \alpha) = 1 - P(z_k \leq \alpha)$. The cumulative probability of the nonnegative distance measure is given by

$$P(z_k \leq \alpha) = \int_0^\alpha f_{z_k}(z) dz = \iint_{\mathcal{D}} f_r(\mathbf{r}) d\mathbf{r}, \quad (32)$$

where $d\mathbf{r}$ is the differential area element over the region

$$\mathcal{D} = \{r_k \mid z_k = (r_k - \mu)^T \Sigma^{-1} (r_k - \mu) \leq \alpha\}, \quad (33)$$

which is in general a p -dimensional ellipsoid and whose boundary is defined by the projection of the level set $z_k = \alpha$ onto the r_k plane, and threshold α is the assigned threshold

of the generalized chi-squared detector. Effectively, (32) expresses that the probability of having $z_k \leq \alpha$ is equal to the volume under the f_r distribution, restricted to the r_k values that generate a z_k value less than or equal to α . Since f_r is composed of a mixture of Gaussian modes, Fig. 2 depicts how this formulation sums the volume of probability contained under the various Gaussian modes and within the region \mathcal{D} . Replacing f_r in equation (32) explicitly expressing the m GMM modes leads to

$$P(z_k \leq \alpha) = \frac{1}{\sqrt{(2\pi)^p}} \sum_{j=1}^m \frac{\pi_j}{\sqrt{\det K_j}} \iint_{\mathcal{D}} e^{(r_k - \mu_j)^T K_j^{-1} (r_k - \mu_j)} d\mathbf{r}.$$

In order to write this equation in terms of the assigned threshold on the detector, we write it in normalized spherical form, changing the volume element $d\mathbf{r}$ to the volume element $d\rho$ which is a volume element over an p -sphere characterized with radius $\sqrt{\alpha}$; hence,

$$r_k - \mu = \mathcal{C}^T \rho_k \rightarrow r_k - \mu_j = \mathcal{C}^T \rho_k + \gamma_j,$$

$$\text{and } d\rho = \frac{d\mathbf{r}}{\det \mathcal{C}},$$

where ρ_k is a vector varying inside the p -sphere and is an affine transformation of the steady state residual $\rho_k = \mathcal{C}^{-T} (r_k - \mu)$ and the matrix \mathcal{C} is the Cholesky decomposition of covariance Σ . This transformation simplifies the new

region of integration to be a p -sphere. Making this substitution and expressing the limits of integration in spherical form yield the final form in (31) and (\star), where ρ_k in spherical form is a function of its norm $|\rho_k|$ and angles ϕ_i , $i = 1 \dots p-1$, where p is the dimension of the measurement. ■

Remark 2: Now that we have the relationship (31) and (\star) between assigned threshold of the detector α and the corresponding CDF for the distance measure in the no fault/attack case $P(z_k < \alpha)$, we can use this result to find the threshold α that provides a desired false alarm rate. There is a mapping between the threshold α and the false alarm rate \mathcal{A} . Armed with Theorem 1, it is possible to use, for example, a bisection method to find the threshold value to yield a desired false alarm rate.

IV. INTERPRETATION

Using an approach that leverages the Gaussian mixture model representation of arbitrary noise distributions not only recovers analytic tractability, it also provides an intuitive and instructive backwards compatibility with the standard chi-square detector driven by Gaussian noise. Suppose that in using the tools presented in this paper, we find that the GMM of the residual distribution is the combination of three distinct Gaussian modes (e.g., see Fig. 1),

$$f_r(x) = \sum_{j=1}^3 \pi_j \mathcal{N}(x | \mu_j, K_j), \quad (34)$$

where μ_j are the mean values of the GMM modes, K_j are the corresponding covariances, and π_j are the mixing probabilities.

A way to interpret the GMM residual probability distribution (34) is that at each time k the residual r_k is drawn from the first GMM mode with probability π_1 , drawn from the second GMM mode with probability π_2 , and drawn from the third GMM mode with probability π_3 . For each of these modes of the residual GMM, we can reverse engineer a hypothetical Gaussian measurement noise that if applied to the system in isolation would generate a Gaussian residual equal to that mode of the residual GMM. This hypothetical Gaussian noise would have mean a_j and covariance C_j ,

$$a_j = E^{-T} \mu_j \quad \text{and} \quad C_j = E^{-1} K_j E^{-T}, \quad (35)$$

for $j \in \{1, 2, 3\}$ and the matrix E is defined as

$$E = \sum_{\kappa=1}^{\infty} A_{\kappa}.$$

If the measurement noise was Gaussian with mean a_j and covariance C_j , $j \in \{1, 2, 3\}$, then we could tune a conventional chi-squared detector using Lemma 1 to determine a threshold α_j to yield a false alarm rate \mathcal{A}_j . If further, we selected $\mathcal{A}_j = 1 - M_j$, where M_j is defined in (\star), then (31) becomes

$$\mathcal{A} = 1 - \pi_1(1 - \mathcal{A}_1) - \pi_2(1 - \mathcal{A}_2) - \pi_3(1 - \mathcal{A}_3), \quad (36)$$

$$= \pi_1 \mathcal{A}_1 + \pi_2 \mathcal{A}_2 + \pi_3 \mathcal{A}_3, \quad (37)$$

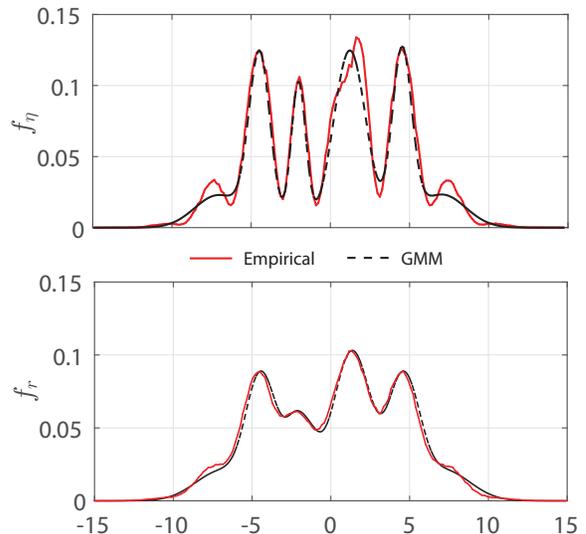


Fig. 3: The measurement noise distribution (top) is approximated by a Gaussian mixture model with six modes and leads to a corresponding complex residual distribution (bottom). The GMM approximations agree well with the empirical distributions from Monte-Carlo simulations.

since $\pi_1 + \pi_2 + \pi_3 = 1$. In this context, then M_j can also be interpreted the probability under the j -th Gaussian mode distribution (characterized by μ_j and K_j) over the integration region

$$\mathcal{D}_j := \{r_k | z_k = (r_k - \mu_j)^T K_j^{-1} (r_k - \mu_j) \leq \alpha_j\}, \quad (38)$$

which is defined by the level set $z_k = \alpha_j$ of the paraboloid characterized by μ_j and K_j . Thus M_j and α_j are related according to Lemma 1

$$\alpha_j = 2P^{-1} \left(\frac{M_j}{2}, \frac{p}{2} \right). \quad (39)$$

This interpretation is depicted in Fig. 2.

V. EXAMPLE

Consider a single output system and estimator characterized by the following matrices and driven by measurement noise (no system noise) distributed according to the probability density function shown in Fig. 3,

$$F = \begin{bmatrix} 0.8 & 0.2 \\ -0.25 & 0.1 \end{bmatrix}, \quad C = \begin{bmatrix} 0.5 & 0.5 \end{bmatrix}, \quad L = \begin{bmatrix} 0.3 \\ -0.3 \end{bmatrix}.$$

We select a threshold for the detector $\alpha = 0.75$ and use Theorem 1 to calculate the expected false alarm rate for this threshold.

Figure 3 shows the fit of the GMM of the measurement noise compared to the empirical distribution attained from a Monte-Carlo simulation with 5×10^6 samples. Here we select a mixture of six ($m_1 = 6$) Gaussian modes and Table I presents the means and covariances of each mode.

One of the challenges of our method is the possibility of having a large number of terms in the residual distribution

	1	2	3	4	5	6
p^η	0.0847	0.2012	0.1184	0.3200	0.1889	0.0869
μ^η	-7.0877	-4.4709	-2.0082	1.2318	4.5240	7.0504
K^η	2.1997	0.4471	0.2062	1.0392	0.3858	2.2329

TABLE I: The coefficients, means, and covariances of the GMM modes ($m_1 = 6$) of the measurement noise in Fig. 3.

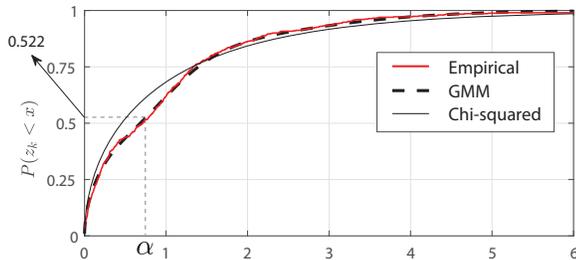


Fig. 4: The cumulative distribution function of the distance measure corresponding to the measurement noise in Fig. 3 evidencing noticeable deviations away from a comparable chi-squared distribution (which would correspond to Gaussian measurement noise). The GMM distribution agrees very well with the empirical distribution found through Monte-Carlo simulation. For a threshold $\alpha = 0.75$, the GMM distribution, using Theorem 1, predicts a false alarm rate of $\mathcal{A} = 1 - 0.522 = 0.478$. The empirical distribution has a false alarm rate of $\mathcal{A} = 1 - 0.516 = 0.484$.

GMM: selecting the settling time $k^* = 10$, $m = m_1^{k^*} = 6^{10}$ since we have no system noise. In practice, however, many of these terms have mean and covariance values that are extremely similar. We can then greatly simplify the GMM expression by aggregating terms that are roughly the same. To do this we define a threshold on the normed difference between mean values and covariance values, d_μ and d_K . If the normed difference between mean values and covariance values for any pair of modes is less than these thresholds,

$$\|\mu_i - \mu_j\| \leq d_\mu \quad \text{and} \quad \|K_i - K_j\| \leq d_K,$$

for $i \neq j \in \{1, \dots, m_1\}$, we consider those terms the same and their coefficients are added together. This procedure provides an arbitrarily accurate approximation with significantly fewer terms. Here we select $d_\mu = 0.0747$ and $d_K = 0.0917$ using a heuristic based on the spread of the distribution and these choices lead to $\tilde{m} = 282$ terms that remain, significantly fewer than the original $m = 6^{10}$. Using this reduced number of modes, the GMM expression of the distribution of residual is shown in Fig. 3 and compares favorably with the empirical residual distribution computed (through Monte-Carlo simulation) from the true noise distribution. By using Theorem 1, we calculate the expected false alarm rate

$$\mathcal{A} = 1 - P(z_k \leq 0.75) = 0.478$$

This expected false-alarm rate is shown in Fig. 4 and the value of \mathcal{A} compares well with the value attained through Monte-Carlo simulation (0.484).

VI. CONCLUSION

In this paper, for discrete-time LTI systems subject to arbitrary sensor and measurement noise, we provide tools to tune model-based detectors and characterize the trade-off between true and false positives. We have generalized one of the most widely used fault detector, the chi-squared detector, for use with general noise distributions. Our approach uses a Gaussian mixture model expression of the disturbances which preserves some of the appealing analytic tractability of working with Gaussian noises on an LTI system.

REFERENCES

- [1] Z. Guo, D. Shi, K. H. Johansson, and L. Shi, "Optimal Linear Cyber-Attack on Remote State Estimation," *IEEE Transactions on Control of Network Systems*, vol. PP, no. 99, pp. 1–10, 2016.
- [2] Y. Mo and B. Sinopoli, "On the performance degradation of cyber-physical systems under stealthy integrity attacks," *IEEE Transactions on Automatic Control*, vol. 61, pp. 2618–2624, 2016.
- [3] C. Z. Bai, F. Pasqualetti, and V. Gupta, "Security in stochastic control systems: Fundamental limitations and performance bounds," in *American Control Conference (ACC)*, 2015, pp. 195–200.
- [4] C. Murguia and J. Ruths, "Characterization of a cusum model-based sensor attack detector," in *proceedings of the 55th IEEE Conference on Decision and Control (CDC)*, 2016.
- [5] —, "Cusum and chi-squared attack detection of compromised sensors," in *proceedings of the IEEE Multi-Conference on Systems and Control (MSC)*, 2016.
- [6] J. Chen and R. J. Patton, *Robust Model-based Fault Diagnosis for Dynamic Systems*. Norwell, MA, USA: Kluwer Academic Publishers, 1999.
- [7] K. Astrom and B. Wittenmark, *Computer-controlled Systems (3rd Ed.)*. Prentice-Hall, Inc., 1997.
- [8] A. G. Bacharoglou, "Approximation of probability distributions by complex mixtures of gaussian measures," *Proceedings of the American Mathematical Society*, vol. 138, no. 7, pp. 2619–2628, jul 2010.
- [9] L. Xu and M. I. Jordan, "On convergence properties of the em algorithm for gaussian mixtures," *Neural Computation*, vol. 8, no. 1, pp. 129–151, 1996.
- [10] H. Stark and J. W. Woods, *Probability, Statistics, and Random Processes for Engineers, 4th Edition*. Pearson Education, Inc., 2014.