Monitoring Model Drift in Feedback-Controlled Systems via Efficacy and Efficiency Metrics

Leonardo Baldo[1], Ferhat Tamssaouet[2], Jorge F. Silva[3], and Marcos E. Orchard[4]

[1] *Machine Health Intelligence Group, Zurich University of Applied Sciences, 8400 Winterthur, Switzerland*
leonardo.baldo@zhaw.ch

[2] *LAAS-CNRS, Université de Toulouse (UT), Toulouse, France*
ferhat.tamssaouet@laas.fr

[3,4] *Department of Electrical Engineering and Data and Artificial Intelligence Initiative, Faculty of Physical and Mathematical Sciences, University of Chile, Santiago, Chile*
josilva@ing.uchile.cl
morchard@ing.uchile.cl

ABSTRACT

This work explores the possibility of monitoring model drift inside a control loop by leveraging efficacy and efficiency metrics. The methodology is based on a fundamental energy–error relationship that links tracking performance to the control effort required to sustain it. The proposed strategy is explicitly designed to deal with feedback-controlled systems and relies only on basic loop signals and simple indicators on time windows: a control-energy metric and an error-performance metric. Based on the similarity between model drift and fault detection problems, the method enables early detection and tracking of progressive faults or drifts that would otherwise remain hidden. The monitoring relies on residuals obtained by comparing the joint evolution of the two metrics against a nominal-condition baseline model. The approach is demonstrated through a MATLAB/ Simulink simulations across various degradation scenarios, showing consistent sensitivity to drifts and enabling their timely detection well before loss of performance becomes apparent at the output level. These results support a lightweight, explainable pathway for model drift monitoring in feedback-controlled systems without requiring additional sensors, the development of complex machine learning/high-fidelity physics models, or structural modifications to the control architecture.

1. INTRODUCTION

Closed-loop control systems have been the foundation of control strategy since their introduction and are extensively utilized in applications where a system or process is required to follow specified reference signals. In closed-loop settings, control data naturally captures the system's input–output behaviour over time, enabling monitoring strategies that leverage deviations from nominal models. In fact, the increasing availability of operational data and the ongoing push towards Industry 4.0 have accelerated the adoption of data-driven modeling and monitoring solutions, often framed under the definition of Prognostics and Health Management (PHM) (Zio, 2022). When applied to controlled systems, degradation may not manifest only as parameter changes, but also as progressive alterations in the input–output relationship captured by nominal models (Kordestani, Saif, Orchard, Razavi-Far, & Khorasani, 2021). In this context, model drift can be defined as a change over time in the statistical relationship between a system's inputs and outputs, resulting in a loss of validity of a model that previously accurately represented the system's nominal behavior (Lu et al., 2019). Given this definition, model drift monitoring is therefore closely related to Fault Detection and Isolation (FDI), since prognosable evolving faults often manifest exactly as alterations in the underlying input–output mechanism. In this perspective, faults naturally correspond to actual drifts, as they alter the system's internal dynamics and therefore modify the input–output relationship, making drift detection a possible strategy for FD (Ramírez, Silva, Tamssaouet, Rojas, & Orchard, 2025).

Although model drift monitoring may appear straightforward

in closed-loop systems, the feedback regulation mechanism significantly complicates the problem. In fact, it is the intrinsic capability of the control loop to guide and force the system to always follow the reference command that leads to a masking of the drift of the model given that the controller always tries to reduce the error no matter what. This challenge closely mirrors a well-known issue in FDI for closed-loop systems, where regulation attenuates fault signatures and sustains apparent nominal tracking until compensation limits are reached. Moreover, closed-loop regulation attenuates the effect of disturbances and compensates for emerging anomalies. To all intents and purposes, feedback regulation diminishes fault signatures and typically manages to obtain near-nominal output tracking until the controller is saturated (Marzat, Piet-Lahanier, Damongeot, & Walter, 2012a).

Motivated by these established insights, this paper treats model drift monitoring in closed-loop as a natural extension of performance monitoring in feedback systems. Given the currently limited body of literature explicitly addressing drift monitoring under closed-loop regulation, we build on the more developed foundations of closed-loop FDI to derive and validate monitoring indicators.

Traditional fault identification methods are largely conceived for open-loop configurations and, as a result, they do not transfer directly to closed-loop controlled systems (Marzat, Piet-Lahanier, Damongeot, & Walter, 2012b). This situation represents a common challenge in many PHM projects. The existing literature offers a perspective on addressing these issues through both physics-based and Machine Learning (ML) techniques. Model-based approaches remain prevalent during the early development stages, employing analytical redundancy and residual generation. However, the absence of detailed physical models for complex processes often impedes scalability, particularly for older and more complex systems where developing physics-based models is both difficult and impractical. As always, the alternative approach in PHM focuses on data-driven strategies, which aim to develop system models based on extensive operational data. Nevertheless, for legacy systems that are already in operation, limited data availability can restrict the effectiveness of these approaches, as they often require a sufficient number of sensors and a high sampling rate to effectively train and validate ML models. In the following paragraphs, the main literature approaches for fault detection of closed-loop systems are analysed.

In the context of model-based strategies, active fault diagnosis has emerged as a promising direction for closed-loop systems. In particular, an auxiliary excitation is intentionally injected to improve fault observability and discrimination. In other words, an auxiliary input is introduced into the system to assist in distinguishing between faulty and normal behaviors. In this sense, Matei et al. (Matei, Zhenirovsky, de Kleer, & Goebel, 2023) introduce a joint control–diagnosis disambiguation scheme for closed-loop systems, where small control perturbations are computed through an optimization similar to Model Predictive Control (MPC) to deliberately steer the plant into state regions that make intermittent actuator faults observable, demonstrating the approach on a quadrotor. Other UAV studies implement the active fault diagnosis concept by superimposing a small signal component on the command of suspected actuators. This is the approach used by Bateman et al. (Bateman, Noura, & Ouladsine, 2008) using a sinusoidal signal as an excitement signal. Another example is provided by Stoustrup and Niemann (Stoustrup & Niemann, 2006; Niemann, 2006) which propose an intrusive active fault diagnosis strategy for parametric/multiplicative faults in closed-loop systems by temporarily modifying the controller so that the faulty plant becomes unstable, making the fault rapidly observable, with a minimal effect on the complete system. Using injected excitation to improve fault detection is promising but risky during normal operation, as perturbations may compromise safety and system stability. In another model-based methodology, Patton and Chen (Patton & Chen, 1991) survey parity-space fault detection and isolation, where faults are revealed by constructing residuals from finite-window analytical redundancy in input–output data and, in closed-loop, via equivalent observer-based residual generators.

For data driven strategies, an example is provided by Sun et al. (Sun, Wang, He, Zhou, & Gu, 2019), who proposed an improved deep neural network framework for fault identification in closed-loop control systems, combining a sliding-window strategy with mode classification. They validate their framework on three scenarios: a controlled numerical simulation, the Tennessee Eastman Process benchmark and a satellite attitude control system with sensor faults. Fuzzy modeling techniques have been approached in Liu et al. (Liu, Yang, Li, & Han, 2022) and reinforcement learning in Zaccaria et al. (Zaccaria, Sartor, Del Favero, & Susto, 2024). In a more control theory perspective, Baïkeche et al (Baïkeche, Marx, Maquin, & Ragot, 2006) address FDI in linear closed-loop systems by deriving transfer-function–based residual generators.

A complementary approach is the one proposed by Marzat et al. (Marzat et al., 2012a), who presents a technique from a perspective that aligns with performance monitoring. Their methodology is based on the idea that the control loop has been designed and tuned to satisfy the requirement of the equipment in nominal condition. Their methodology aims to detect system faults by analyzing the adequacy of the system response. In other words, the discrepancy between the expected and actual control objectives is directly related to the health of the component. In particular, they apply their methodology for a surface-to-air missile model and to monitor the control performance they employ flight mechanics related specific metrics: velocity-pursuit guidance, and propor-

tional navigation guidance law to finally expand the description to observer-based fault detection in linear systems.

Against this backdrop, the present paper proposes a solution which places itself in the middle between MPC and the main idea proposed by Marzat et al. This paper, in particular, integrates the approach proposed by Marzat et al. for fault detection through monitoring control loop performance, while extending the methodology by incorporating the Energy-Error Relationship.

The remainder of this paper is organized as follows: Section 2 introduces the methodology rationale and formalizes the Energy-Error Relationship as a fundamental trade-off for monitoring. Section 3 details the model drift framework, the residual generation process, and the statistical decision logic. Section 4 describes the modeling framework, including the simulation environment and the drift injection mechanism. Section 5 evaluates the selection of the metrics of the system efficacy and efficiency. Finally, Section 6 presents the results and a discussion on the control loop masking effect, followed by concluding remarks and future work in Section 7.

2. METHODOLOGY RATIONALE & ENERGY-ERROR RELATIONSHIP

Let us consider a controlled dynamical system operating in closed-loop:

$$\dot{x}(t) = f(x(t), u(t), d(t), \theta), \qquad y(t) = h(x(t), \theta), \quad (1)$$

where $u(t)$ is the control input (control effort), $d(t)$ denotes external disturbances, θ represents system parameters (possibly health-related), and $y(t)$ is the measured output, By considering $r(t)$ as the reference signal, we define the tracking error as :

$$e(t) = r(t) - y(t)$$

Under nominal closed-loop operation (correct controller design, healthy plant, normal disturbances, i.e., within the envelope for which the controller was designed to compensate) a well designed and tuned system should comply with the dynamic requirements and its mission objectives should be satisfied. This supports the hypothesis that there exists a trade-off between control effort and tracking performance and that this trade-off could be formalized into a mathematical equation. In particular, if the time is partitioned into time windows of duration T, indexed by k, over each window, we can define two scalar indicators summarizing the closed-loop behaviour: namely the control energy indicator reported in Equation 2 and the Error-performance indicator reported in Equation 3.

$$E_u(k) = \int_{kT}^{(k+1)T} \phi_u\big(u(t)\big)\, dt, \quad (2)$$

typically with $\phi_u(u) = \|u\|^2$, giving an input "energy" over the window.

$$J_e(k) = \int_{kT}^{(k+1)T} \phi_e\big(e(t)\big)\, dt, \quad (3)$$

with choices such as $\phi_e(e) = \|e\|^2$, $|e|$, or classical indices such as Integral Time-averaged Absolute / Squared Errors (ITAE/ITSE).

Under the mentioned assumptions, a simple functional relationship between the two indicators can be defined as:

$$J_e(k) \approx \alpha\, E_u(k) + \beta, \quad (4)$$

or, more generally,

$$J_e(k) \approx f_\star\big(E_u(k)\big), \quad (5)$$

for some smooth function $f_\star$, typically monotonic or convex, reflecting the classical trade-off between control effort and tracking performance. The rationale behind this is rooted in the fact that, for a fixed controller and operating regime, tracking performance is obtained at the price of actuation effort. In nominal conditions, the closed-loop system regulates the plant by continuously shaping the input $u(t)$ so as to reject disturbances and minimize the tracking error $e(t)$. As a result, when the loop is healthy and correctly tuned, the pair $(E_u(k), J_e(k))$ tends to evolve along a relationship that reflects the intrinsic effort–performance balance imposed by the controller and by the plant dynamics. This perspective is particularly relevant in the presence of gradual drift. Under closed-loop operation, drifts are often compensated by the controller and therefore remain weakly observable at the output level. Nevertheless, compensation is not free: for the same reference and disturbance environment, a drifted model typically requires larger or more aggressive actuation to maintain comparable tracking performance. This produces a systematic change in the joint statistics of $(E_u(k), J_e(k))$, which can be captured as a deviation from the nominal effort–performance relationship $f_\star\big(E_u(k)\big)$. This deviation can arise well before any obvious loss of tracking is visible in $y(t)$, precisely because the loop preserves performance by increasing effort until control authority becomes saturated.

3. MODEL DRIFT

A model drift framework can then be developed based on the residuals methodology as often performed in recent literature. The strategy adopted in this work consists of first learning the system/model behaviour in nominal conditions and then monitor how actual measurements deviate from this reference. A monitoring model is trained on data collected during normal operation to capture the expected relationships among signals. During monitoring, the same model generates the corresponding "healthy" reconstructions, which are compared against the measured signals. The resulting mismatch, commonly referred to as a residual, provides a compact in-

dicator of abnormal behaviour. The analysis of the trend and values of the residuals can trigger further diagnosis decisions.

From historical healthy data, it is possible to estimate the model of the system in nominal condition or nominal mapping $f_\star$, e.g. via regression:

$$J_e(k) = f_\star\big(E_u(k)\big) + \varepsilon(k), \tag{6}$$

with $\varepsilon(k)$ a modelling residual.

During online operation then, we compute:

$$r(k) = J_e(k) - \hat{f}_\star\big(E_u(k)\big), \tag{7}$$

where $\hat{f}_\star$ is the learned nominal model.

3.1. Decision logic

- If $r(k)$ remains small (within a nominal tolerance band), then the control performance is consistent with the expected nominal behaviour.
- If $r(k)$ exhibits a significant deviation, then:
  - transient deviation followed by recovery suggests a *disturbance*,
  - persistent deviation suggests a *drift*.

A set of standard statistical change–detection tests was then applied to the residual sequence $r(k)$ in order to identify departures from nominal behaviour. In particular:

- **CUSUM:** a cumulative-sum detector designed to reveal persistent small shifts in the mean of $r(k)$ by integrating its incremental deviations from the nominal level.
- **Shewhart distance:** an instantaneous test based on monitoring the normalized magnitude of $r(k)$ with respect to a nominal dispersion estimate (i.e., standard deviation).
- **SPRT:** the Sequential Probability Ratio Test, formulated as a log-likelihood ratio accumulated over time to discriminate between a null hypothesis (nominal residual statistics) and an alternative hypothesis (shifted mean and/or variance).

The overall methodology is reported in Figure 1.

3.2. Disturbance versus fault

A disturbance temporarily perturbs the relation $J_e \leftrightarrow E_u$, but the points $(E_u(k), J_e(k))$ eventually return to the nominal curve characterised by $f_\star$.

A fault, however, changes the structure of the system, giving rise to a new mapping

$$J_e(k) \approx f_{\text{fault}}\big(E_u(k)\big), \tag{8}$$

or, in a linearised form,

$$J_e(k) \approx \alpha_{\text{fault}}\, E_u(k) + \beta_{\text{fault}}, \tag{9}$$

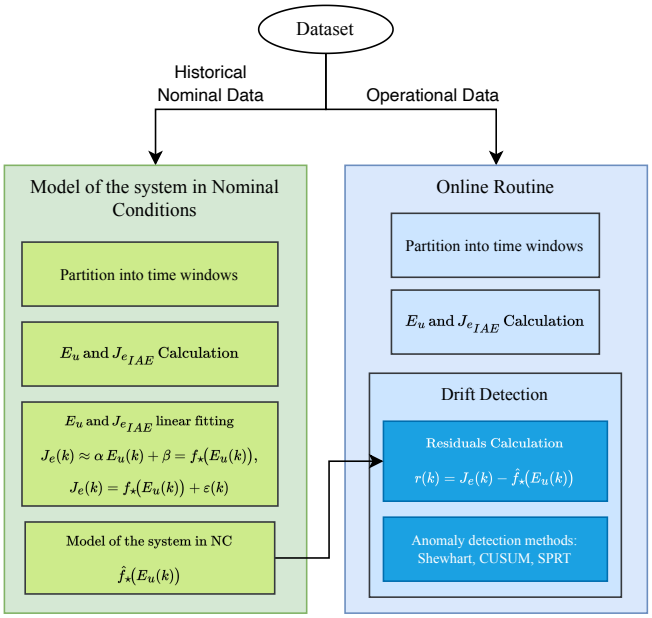

Figure 1. Overall methodology. The two branches indicate the creation of the model of the system in nominal condition and the actual operation of the asset leveraging the model of the system in nominal condition of residual calculation.

thus inducing persistent deviations from the nominal relationship.

4. MODELING FRAMEWORK

To evaluate and test the proposed strategy, a simple closed-loop control system has been implemented in Matlab Simulink. The loop architecture is reported in Figure 2. In particular, a reference command r is compared with the measured output y to produce the tracking error e, which is regulated by a Proportional Integral (PI) controller to compute the control input u. In this version, the controlled process is modeled as a linear time-invariant (LTI) state-space system as reported in Equation 10.

$$x_{n+1} = Ax_n + Bu_n, \qquad y_n = Cx_n + Du_n \tag{10}$$

To emulate an incipient drift mechanism, an additive stochastic drift term is injected at the output measurement, generated from a standard Gaussian source $\mathcal{N}(0,1)$ and shaped by the parameters μ_{drift} and σ_{drift}, which control the mean and variability of the drift contribution. The effect of the two parameters are reported in Figure 3, which illustrates how the deterministic trend and the stochastic dispersion can be tuned independently through $\mu_{\max}$ and σ_{growth}, respectively.

A key aspect of the setup is the design of the reference input r, which is intentionally not chosen as a constant set-point or typical step or square signal. Instead, r is generated as a stochastic time-varying command obtained by sampling from

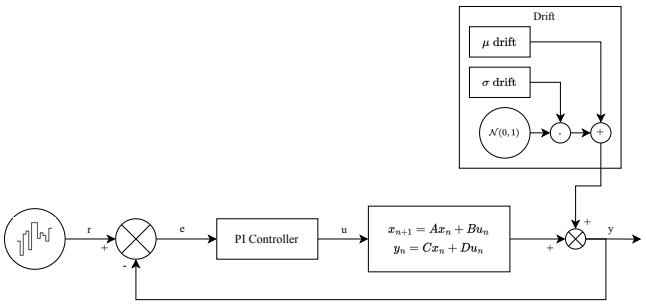

Figure 2. The closed-loop discrete-time simulation block diagram.

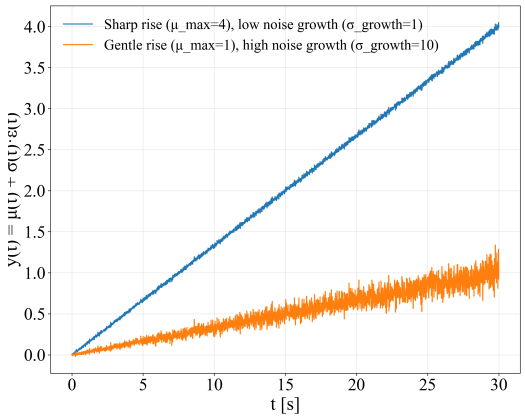

Figure 3. Comparison of two ramp-plus-noise signals, $y(t) = \mu(t) + \sigma(t)\varepsilon(t)$: a steep drift with low noise growth ($\mu_{\max} = 4,\ \sigma_{\text{growth}} = 1$) versus a milder drift with high noise growth ($\mu_{\max} = 1,\ \sigma_{\text{growth}} = 10$).

a prescribed Gaussian probability distribution ($\mathcal{N}(0.1, 1)$), thus producing a persistently evolving input signal over time. This choice is explicitly made to prevent the closed-l oop system from settling into a steady-state operating condition: while the plant output would naturally tend to converge towards the reference under PI control, the continuous set-point variations introduce new tracking demands before full convergence is reached. As a result, the controller–plant loop operates in a permanent transient regime, where the dynamics remain persistently excited and the tracking error does not collapse to a stationary equilibrium. In a long steady-state regime, the (E_u, J_e) samples would collapse around a single operating point, making the energy–performance slope poorly identifiable; this is why the reference signal is designed to provide persistent excitation. This operating condition is particularly relevant for PHM and drift-detection studies in closed-loop systems, as it ensures sustained system activity, thereby enabling the observation of subtle deviations without being masked by purely steady-state behavior.

The plant dynamics is modeled using the continuous-time matrices reported in Equation 12. The continuous-time state-space model is discretized using a Zero-Order Hold (ZOH)

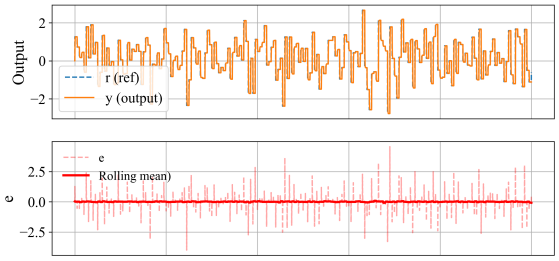

Figure 4. Reference and output signals for a nominal simulation along with the error (filtered and raw). Note that the reference signal changes very frequently to make the system operates in a permanent transient regime. This causes the error to oscillate violently at every step change.

assumption on the input which yields the corresponding discrete-time matrices at sampling time T_s.

$$A_c = \begin{bmatrix} \sigma & \omega & 0 \\ -\omega & \sigma & 0 \\ 0 & 0 & \lambda \end{bmatrix}, \qquad B_c = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \tag{11}$$

$$C_c = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}, \qquad D_c = 0, \tag{12}$$

where $\sigma = -0.2$, $\omega = 2.0$, and $\lambda = -0.5$. These values are selected to be representative of a simple yet non-trivial dynamical system, comprising one real mode and a pair of complex-conjugate modes. This configuration provides a stable closed-loop behavior with realistic oscillatory dynamics. After discretization, the resulting state transition matrix exhibits eigenvalues

$$\{\, 0.9998 \pm 0.0020i,\ \ 0.9995 \,\}, \tag{13}$$

corresponding to two lightly damped oscillatory components and one non-oscillatory (real) component.

5. EFFICACY AND EFFICIENCY INDICATORS

A set of Simulink simulations has been carried out in nominal ($\mu_{\text{drift}} = 0$ and $\sigma_{\text{drift}} = 0$) and non nominal conditions. Data has been saved and then analyzed in VSCode in a Python 3.8.18 environment. Figure 4 shows the reference signal along with the output signal in nominal conditions. Simulation sample time (T_s) is set as $0.001s$ for a total simulation time of $10s$. The sampling time of the input signal is set as $T_s \cdot 50$. The PI controller gains are tuned by trial-and-error starting from Matlab autotune routine and are equal to 300 and 100 for K_p and K_i respectively.

The raw closed-loop signals (r, e, u, y) are processed through a sliding-window analysis of duration T_w (window length). The window is then advanced by a constant shift T_s (step size).

Over each window, a set of scalar indicators are computed. In particular, we considered statistics derived from the tracking error $e(t) = r(t) - y(t)$ and the control input $u(t)$. For the remainder of this work, we focus on two indicators computed over each window, namely a control-energy index E_u (related to the system efficiency) and an error-performance index J_e (related to the efficacy), already introduced in Section 2.

To identify a robust drift indicator, each candidate metric was evaluated over the full $(\mu_{\max}, \sigma_{\text{growth}})$ simulation grid by quantifying its average sensitivity to the drift amplitude $\mu_{\max}$ while penalizing dependence on the parameter σ_{growth}. Specifically, for each fixed σ_{growth} slice, the mean range of the metric across $\mu_{\max}$ levels was computed and contrasted with the mean range across σ_{growth} levels at fixed $\mu_{\max}$. The resulting ratio Δ_μ/Δ_σ provides a score of degradation tracking capability under varying noise-growth conditions. According to this criterion, along with the control-effort index E_u, $J_{e,\text{IAE}}$ showed the best trade-off between sensitivity to $\mu_{\max}$ and robustness to σ_{growth} as reported in Table 1 and Figure 5. Hence, the final feature stream used for detection is based on the pair $\big(E_u(k),\ J_{e,\text{ITAE}}(k)\big)$, which jointly captures the effort–performance trade-off under closed-loop operation.

Table 1. Sensitivity-based ranking of J_e candidate indicators over the $(\mu_{\max}, \sigma_{\text{growth}})$ grid.

Metric	Δ_μ	Δ_σ	Δ_μ/Δ_σ
$J_{e,\text{IAE}}$	0.100326	0.000767	**130.86**
$J_{e,\text{ITAE}}$	0.158342	0.001325	119.55
$J_{e,\text{ITSE}}$	0.013958	0.000181	77.29
$J_{e,\text{ISE}}$	0.008087	0.000139	58.29

Figure 5 shows the evolution of the different metrics with the same reference signals as presented in Figure 4. Figure 5a shows simulation considering fixed σ_{growth} and changing $\mu_{\max}$, while Figure 5b showed fixed $\mu_{\max}$ and varying σ_{growth}.

6. RESULTS & DISCUSSION

6.1. Control Loop masking

Figure 6 shows the same signals of Figure 4 but in a non nominal condition. Despite the progressive perturbation introduced on the output measurement, the tracking error remains remarkably similar to Figure 4, with only marginal changes in its amplitude and rolling mean, highlighting the drift masking effect of feedback regulation. In contrast, Figure 7 reveals a clear signature in the control action u: the controller progressively increases its effort to preserve tracking as the added disturbance on y grows linearly and eventually reaches an amplitude of 10, i.e., more than five times the maximum reference command magnitude. This confirms that, while the output-level error may appear nearly unchanged, the control input provides a substantially more informative channel to expose the drift growth in closed-loop operation.

6.2. Development of the Model of the System in Nominal Condition

An extended simulation of 100 seconds under nominal conditions was performed in order to collect a sufficiently large

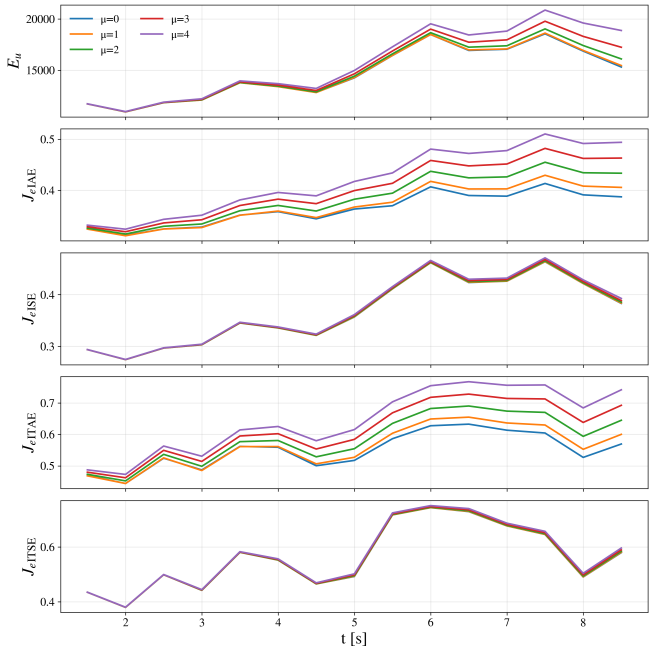

(a) $\sigma = 0$ (varying μ).

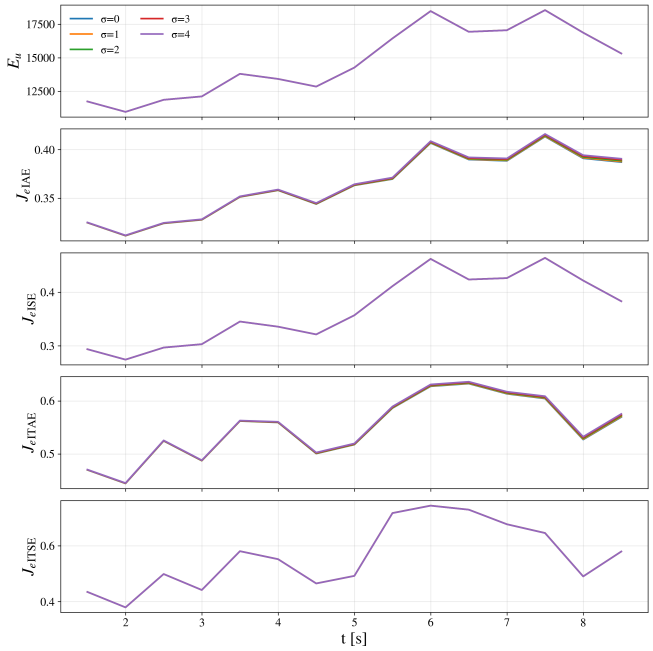

(b) $\mu = 0$ (varying σ).

Figure 5. Comparison of candidate sliding-window indicators under increasing degradation. Top: mean degradation sweep ($\sigma = 0$). Bottom: noise-growth sweep ($\mu = 0$). Among the considered error-performance formulations, $J_{e,\text{IAE}}$ shows the clearest and most consistent separation across degradation levels while containing the effect due to increasing σ_{growth}.

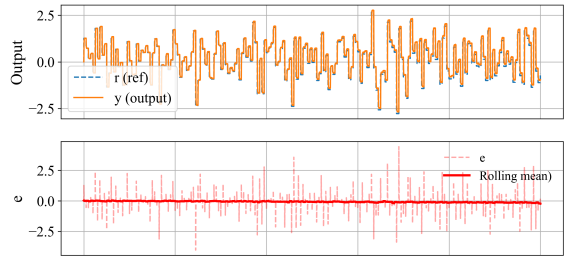

Figure 6. Reference and output signals for a non nominal simulation along with the error (filtered and raw).

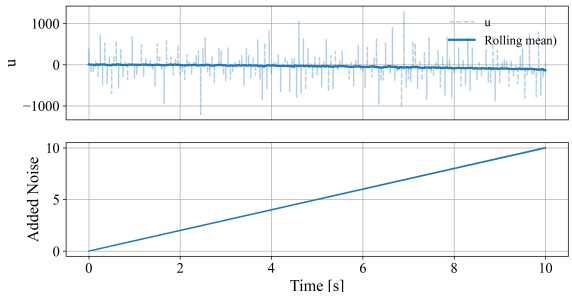

Figure 7. Control action u (raw and filtered) and added disturbance.

baseline dataset and identify the nominal energy-error relationship via linear regression. This simple model structure was selected after inspection of the data, which exhibited a clear linear tendency, thus supporting a trade-off between modeling simplicity and descriptive accuracy. The nominal relationship was estimated as $J_e = b_0 + b_1 \cdot E_u$ at a 95% confidence level yielding $b_0 = 0.175$ ($[0.166616, 0.182278]$) and $b_1 = 1.26 \times 10^{-5}$ ($[1.2, 1.31] \times 10^{-5}$), which confirms a monotonic increase of J_e with the control-energy indicator. The fitted model explains a substantial fraction of nominal variability $R^2 = 0.91548$, adjusted $R^2 = 0.91504$, with RMSE 1.06×10^{-2}. Figure 8a shows the resulting nominal trend together with the associated uncertainty bands and the deviation of non-nominal windows from the expected effort-performance trade-off.

6.3. Methodology in Operation

Figure 8b shows the three monitoring strategy applied to the nominal simulation ($\mu_{\max} = 0$, $\sigma_{\text{growth}} = 0$) showing no detection of drift: the majorities of standardized residuals are inside the PI, the CUSM does not detect any consistent drift and the SPRT shows a cumulative log likelihood ratio inside the thresholds.

If we now consider the performance of the fault detection strategy with non nominal simulations, for instance the one

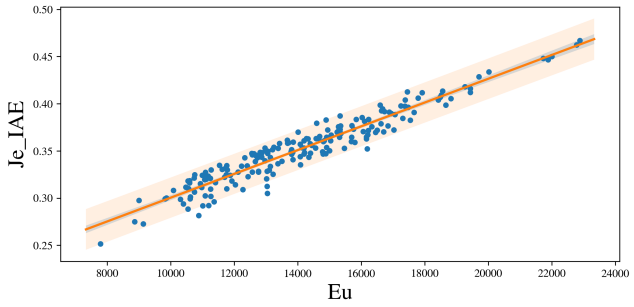

(a) Nominal energy–error relationship with linear fit and uncertainty bands (confidence band on the mean and prediction band for individual samples)

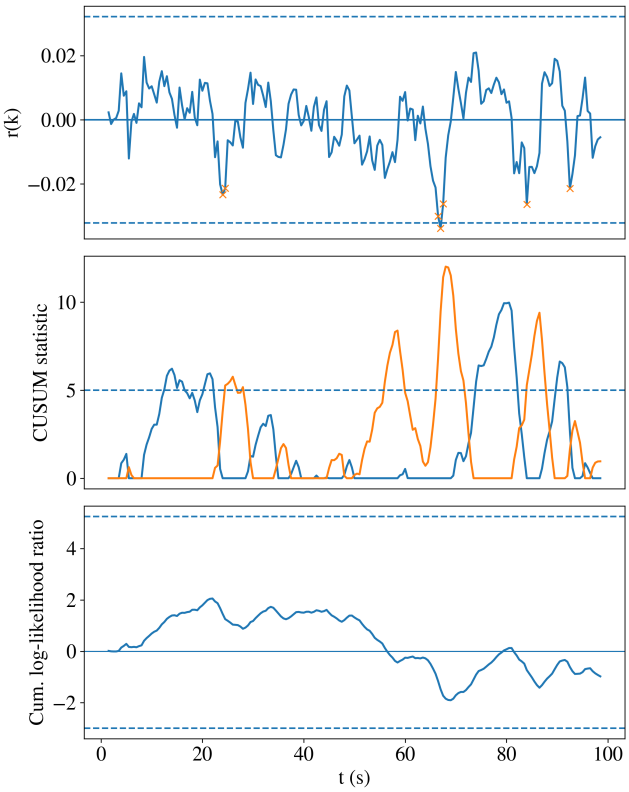

(b) Residual-based monitoring: Shewhart thresholds, CUSUM, and SPRT (shown in function of cum. log-likelihood ratio).

Figure 8. Energy–error nominal modelling and residual-based change detection for closed-loop fault monitoring.

reported in Figure 9, we can see that a small but non negligible number of points, highlighted with an orange cross in Figure 9a, begin to fall outside of the PI. In this case the injected drift corresponds to ($\mu_{\max} = 1$, $\sigma_{\text{growth}} = 0$) This is further confirmed by the three monitoring strategies on the residuals which detect a drift at around $40s$ (CUSUM) and $35/60s$ (SPRT). The methodology is then tested on a higher scale of drift, as reported in Figure 10, where ($\mu_{\max} = 3$, $\sigma_{\text{growth}} = 0$). The impact of the drift on the energy-error relationship is clearly visible in Figure 10 and the detection on the residuals

is reported in Figure 10b, where the drift is promptly detected in both CUSUM and SPRT methods. From the evolution of the energy–error relationship, it is apparent that increasing drift manifests primarily as a systematic shift of the curve, i.e., a progressive change in the intercept ($b_{0_{nonnominal}}$), while the slope remains essentially unchanged. This behaviour indicates that degradation does not fundamentally alter the effort–performance sensitivity of the closed-loop, but rather increases the baseline control action required to sustain a given level of tracking. In practical terms, as model drift develops, comparable performance can still be achieved, yet only at the expense of a higher control-energy budget.

7. CONCLUSION

Closed-loop feedback inherently compensates for system changes, thereby preserving reference tracking even in the presence of degradation. While beneficial for performance, this compensation mechanism also masks model drift in the measured outputs, often delaying detection until controller is saturated. This paper explores the possibility of a drift detection strategy based on the principle that the degradation can be tracked along with the shift from a nominal relationship between the performance of the control loop and the effort the controller makes. The proposed strategy is inspired by FDI techniques and is tested using a Matlab/Simulink model of a system in closed-loop control with different levels of model drift. Increasing degradation proved to add an overhead of required effort which does not change the slope of the energy-error relationship. The paper provides a grounded base for the development of model drift and fault detection frameworks based on these preliminary results. The idea has proven promising and is currently the subject of an in-depth study. For instance, we explore how to extend the approach to nonlinear plants and real industrial datasets. In this case, the proposed framework could be used to monitor the internal degradation of controlled systems. In addition, the robustness of the energy–error relationship under varying operating regimes and controller retuning will be investigated.

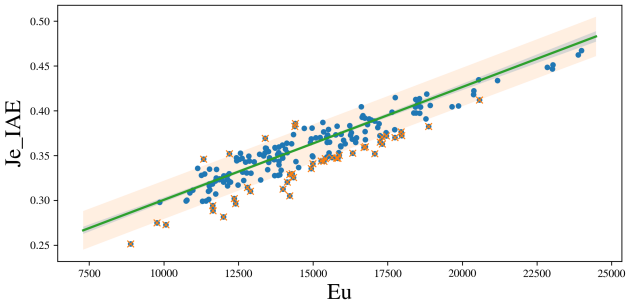

(a) Energy–error relationship with linear fit and CI/PI bands $\mu_{\max} = 1$.

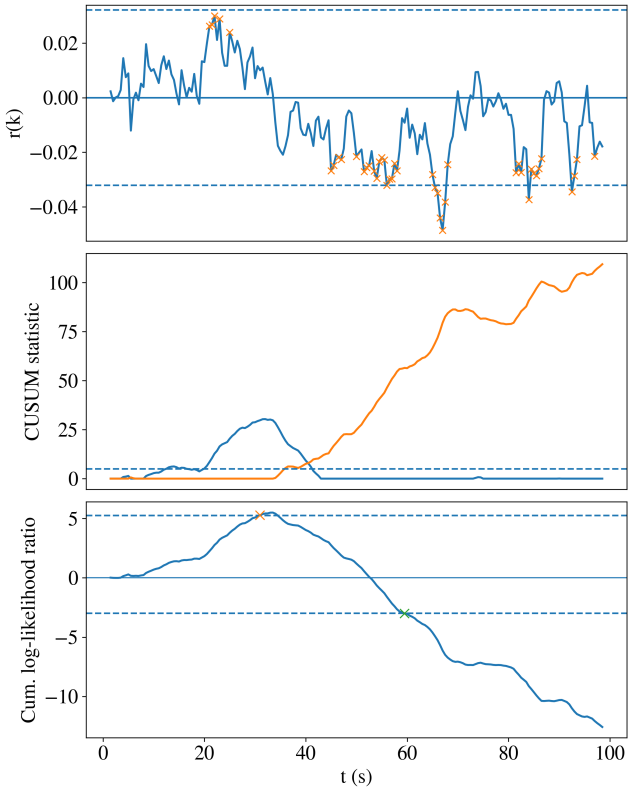

(b) Residual-based monitoring: Shewhart thresholds, CUSUM, and SPRT for $\mu_{\max} = 1$.

Figure 9. Energy–error relation and residual-based change detection for closed-loop drift monitoring in non-nominal conditions ($\mu_{\max} = 1$, $\sigma_{\text{growth}} = 0$).

8. ACKNOLEDGEMENT

The work of Jorge F. Silva is supported by Fondecyt Grant 1250098. The work of Marcos E. Orchard is supported by Fondecyt Grant 1250036 and by Advanced Center for Electrical and Electronic Engineering (AC3E) through ANID Basal Project CIA250006.

REFERENCES

Baïkeche, H., Marx, M., Maquin, D., & Ragot, J. (2006, November). On parametric and nonparametric fault detection in linear closed-loop systems. In *Proceedings of the 4th workshop on advanced control and diagnosis (acd 2006).* Nancy, France.

Bateman, F., Noura, H., & Ouladsine, M. (2008, June). Active fault detection and isolation strategy for an unmanned aerial vehicle with redundant flight control surfaces. In *Proceedings of the 16th mediterranean conference on control and automation (med 2008)* (pp. 1246–1251). Ajaccio, France.

Kordestani, M., Saif, M., Orchard, M. E., Razavi-Far, R., & Khorasani, K. (2021). Failure prognosis and applications—a survey of recent literature. *IEEE Transactions on Reliability*, *70*(2), 728-748. doi:

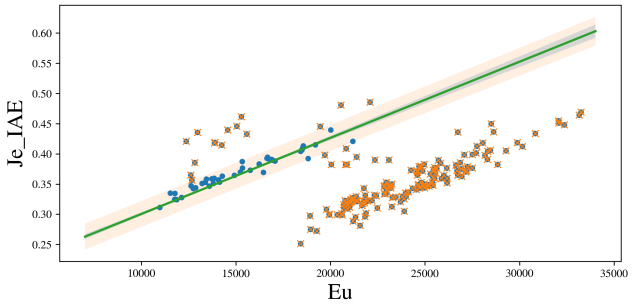

(a) Energy–error relationship with linear fit and CI/PI bands $\mu_{\max} = 3$.

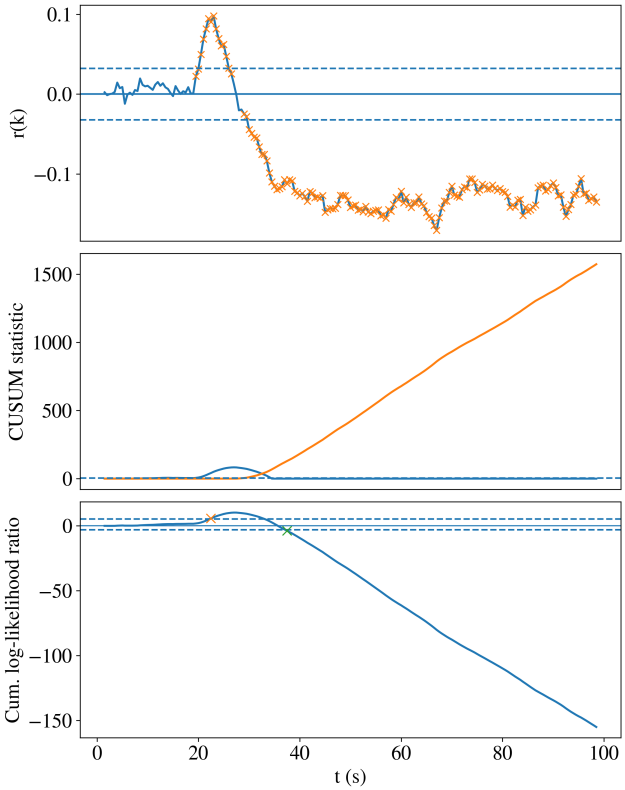

(b) Residual-based monitoring: Shewhart thresholds, CUSUM, and SPRT for $\mu_{\max} = 3$.

Figure 10. Energy–error relationship and residual-based drift detection for non nominal conditions ($\mu_{\max} = 3$, $\sigma_{\text{growth}} = 0$).

10.1109/TR.2019.2930195

Liu, R., Yang, Y., Li, L., & Han, H. (2022). Performance-oriented fault detection for nonlinear control systems via data-driven t-s fuzzy modeling technique. *IEEE Transactions on Fuzzy Systems*, *30*(1), 133-146. doi: 10.1109/TFUZZ.2020.3033140

Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2019). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, *31*(12), 2346-2363. doi: 10.1109/TKDE.2018.2876857

Marzat, J., Piet-Lahanier, H., Damongeot, F., & Walter, E. (2012a). Control-based fault detection and isolation for autonomous aircraft. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, *226*(5), 510-531. doi: 10.1177/0954410011413834

Marzat, J., Piet-Lahanier, H., Damongeot, F., & Walter, E. (2012b). Model-based fault diagnosis for aerospace systems: a survey. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, *226*(10), 1329–1360. doi: 10.1177/0954410011421716

Matei, I., Zhenirovsky, M., de Kleer, J., & Goebel, K. (2023). Joint feedback control and fault diagnosis disambiguation. In *Annual conference of the prognostics and health management society* (Vol. 15). doi: 10.36001/phmconf.2023.v15i1.13366

Niemann, H. (2006). Active fault diagnosis in closed-loop uncertain systems. In *Fault detection, supervision and safety of technical processes 2006* (Vol. 39, pp. 587–592). doi: 10.1016/B978-008044485-7/50099-3

Patton, R. J., & Chen, J. (1991). A review of parity space approaches to fault diagnosis. *IFAC Proceedings Volumes*, *24*(6), 65–81. doi: 10.1016/S1474-6670(17)51124-6

Ramírez, C., Silva, J. F., Tamssaouet, F., Rojas, T., & Orchard, M. E. (2025). Fault detection and monitoring using a data-driven information-based strategy: Method, theory, and application. *Mechanical Systems and Signal Processing*, *228*, 112403. doi: https://doi.org/10.1016/j.ymssp.2025.112403

Stoustrup, J., & Niemann, H. (2006). Active fault diagnosis by temporary destabilization. In *Fault detection, supervision and safety of technical processes 2006* (Vol. 39, pp. 563–568). doi: 10.1016/B978-008044485-7/50095-6

Sun, B., Wang, J., He, Z., Zhou, H., & Gu, F. (2019). Fault identification for a closed-loop control system based on an improved deep neural network. *Sensors*, *19*(9), 2131. doi: 10.3390/s19092131

Zaccaria, V., Sartor, D., Del Favero, S., & Susto, G. A. (2024). Fault identification enhancement with reinforcement learning (fierl). In *2024 american control conference (acc)* (p. 1494-1499). doi: 10.23919/ACC60939.2024.10644272

Zio, E. (2022). Prognostics and health management (phm): Where are we and where do we (need to) go in theory and practice. *Reliability Engineering & System Safety*, *218*, 108119. doi: 10.1016/j.ress.2021.108119