Application of Physics-Informed Models for Electrical Asset Diagnostics in Kenya

Brian K. Mwangi

Department of Engineering, University of Nairobi, Nairobi, Kenya

ABSTRACT

Condition-based maintenance of electrical engineering assets in Kenya is constrained by sparse fault data, which undermines purely data-driven diagnostic models. We formalise the diagnostic problem within a state-space representation and derive a residual-based detection statistic whose distribution is characterised under both healthy and faulty regimes. A principal contribution is the introduction of a lifecycle cost model that explicitly quantifies the trade-off between detection sensitivity and the economic consequences of false alarms and missed faults. The analysis demonstrates that physics-based residuals reduce the variance of detection statistics under distribution shift, thereby improving the reliability of fault declarations relative to purely empirical approaches. We further show that the optimal detection threshold depends on the ratio of misclassification costs and the prior probability of fault onset, which varies across Kenyan asset classes. The framework is illustrated through analytical examples grounded in transformer thermal dynamics and induction motor current signatures. The article concludes by identifying institutional and data constraints that must be addressed before the proposed methodology can be operationalised within Kenyan utility and industrial settings.

Keywords: *physics-informed diagnostics, out-of-distribution detection, lifecycle cost, condition-based maintenance, electrical assets, Kenya*

The Diagnostic Problem in Data-Sparse Electrical Networks

Electrical Utilities And Industrial Operators In Kenya Maintain Critical Assets

Electrical utilities and industrial operators in Kenya maintain critical assets such as power transformers, switchgear and rotating machines under conditions of considerable uncertainty. Scheduled maintenance remains the dominant strategy, yet it is poorly matched to the actual degradation trajectories of individual assets. The alternative, condition-based maintenance, promises to align intervention with need, but its realisation depends on the ability to infer the internal state of an asset from externally measurable signals.

This inference problem is fundamentally complicated by the fact that fault data are rare. A transformer that fails catastrophically after decades of service contributes exactly one failure event to the historical record, and that event is typically accompanied by a cascade of secondary effects that obscure the initiating mechanism. The scarcity of labelled fault examples creates a methodological tension.

The Scarcity Of Labelled Fault Examples Creates A Methodological Tension

Contemporary machine learning approaches to fault diagnosis, particularly those based on deep neural networks, achieve impressive performance when trained on large corpora of labelled examples. Such corpora do not exist for most electrical assets in Kenya, and they are unlikely to be assembled in the near term because the underlying failure processes are slow and the population of identical assets is small. Transfer learning from overseas datasets is only partially successful because Kenyan operating conditions, including ambient temperature profiles, loading patterns and maintenance histories, differ systematically from those under which foreign data were collected.

The result is that a model which performs well on training data drawn from one population may degrade unpredictably when deployed on assets that lie outside the support of that training distribution. the present analysis argues that the way forward lies not in accumulating more data, but in changing the inductive bias of the diagnostic model. Physics-informed diagnostics embed governing equations, derived from first principles, into the structure of the inference procedure. The governing physics provides a strong prior that remains valid even when the statistical properties of the operating environment shift.

A thermal model of a transformer, for example, encodes the relationship between load current, ambient temperature and hot-spot temperature. When sensor readings deviate from the predictions of this model, the deviation carries information about a change in the asset itself, rather than merely a change in the operating point. The central claim developed here is that physics-based residuals are inherently more robust to out-of-distribution conditions than purely statistical features, and that this robustness has direct economic consequences for maintenance decision-making.

State-Space Formulation of Asset Degradation

We Consider An Electrical Asset Whose Condition Is Characterised By

We consider an electrical asset whose condition is characterised by a latent state vector that evolves according to known physical laws. The observable quantities are related to the latent state through a measurement equation that includes sensor noise. This formulation is standard in model-based diagnosis (Isermann, 2006) and provides a natural language for expressing the distinction between healthy operation and incipient fault conditions. Let the discrete-time state-space model be written as follows.

$$x_{k+1} = f(x_k, u_k, \theta) + w_k \tag{1}$$

where xk denotes the latent state at time step k, uk is the known input, θ is a vector of physical parameters, and wk is process noise. The measurement equation is given by the following.

$$y_k = h(x_k, u_k, \theta) + v_k \tag{2}$$

Let The Discrete-Time State-Space Model Be Written As Follows

In this formulation, yk represents the sensor measurements and v_k is measurement noise. Under healthy operation, the parameter vector θ takes a nominal value θ_0 that is known from the asset's design specification or from commissioning tests. A fault manifests as a change in the system dynamics, which can be represented as a deviation of the parameter vector from its nominal value.

For example, an inter-turn short circuit in a transformer winding changes the effective inductance and resistance of the affected phase, thereby altering the relationship between applied voltage and circulating current. Similarly, a broken rotor bar in an induction motor introduces an asymmetry that modulates the stator current spectrum at characteristic sideband frequencies (Thomson & Fenger, 2001). The diagnostic task is to decide, on the basis of the measurement sequence, whether the asset is operating in the healthy regime or whether a fault has initiated.

A natural approach is to construct a residual signal that compares the actual measurements with the predictions of a model that assumes healthy operation. When the residual is small, the healthy model is consistent with the observations. When the residual grows, the evidence for a fault accumulates.

The residual is defined as the difference between the measured output and the output predicted by a filter or observer that runs in parallel with the physical asset (Chen & Patton, 1999). The choice of the residual generation mechanism is consequential. A Kalman filter provides the optimal state estimate under linear dynamics and Gaussian noise, and the resulting innovation sequence has well-characterised statistical properties (Anderson & Moore, 1979).

For nonlinear systems, the extended Kalman filter or the unscented Kalman filter provides an approximation, but the optimality guarantees are lost. In the context of electrical assets, the governing physics is often mildly nonlinear, and the extended Kalman filter has been applied successfully to problems such as battery state-of-charge estimation and transformer thermal modelling (Plett, 2004).

Residual Characterisation Under Distribution Shift

The Central Advantage Of Physics-Based Residuals Becomes Apparent When The

The central advantage of physics-based residuals becomes apparent when the operating conditions of the asset shift outside the range encountered during model development. Consider a purely data-driven approach in which a neural network is trained to map measurement features to a health indicator. The network learns the statistical regularities of the training distribution, including correlations that are incidental to the underlying physics.

When deployed under different ambient temperatures or loading patterns, these incidental correlations break down, and the network's predictions become unreliable. This phenomenon is known as out-of-distribution failure and is a well-documented limitation of deep learning systems (Ovadia et al., 2019). A physics-based residual, by contrast, is constructed from a model whose structure is dictated by the governing equations rather than by the statistical properties of the training data.

The residual responds to genuine deviations from the expected physical behaviour, regardless of whether those deviations arise from a fault or from a change in the operating environment. The challenge is to distinguish between these two sources of deviation. A change in ambient temperature, for example, will cause the thermal model's prediction to shift, but this shift is accounted for by the model's input u_k.

A fault, on the other hand, changes the relationship between inputs and outputs in a way that the healthy model cannot accommodate. To make this distinction precise, we consider the residual sequence generated by a filter that assumes the healthy parameter value θ_0. Under the null hypothesis of healthy operation, the residual is a zero-mean random sequence whose covariance can be computed from the filter's Riccati equation.

A Physics-Based Residual By Contrast Is Constructed From A Model

Under the alternative hypothesis of a fault, the residual acquires a deterministic component whose magnitude depends on the size of the parameter deviation and on the observability of that deviation through the available measurements. The detection problem reduces to deciding between these two hypotheses on the basis of the observed residual sequence. A standard approach is to form a chi-squared statistic from the squared residuals, normalised by their covariance.

Under healthy operation, this statistic follows a central chi-squared distribution. Under a fault, it follows a non-central chi-squared distribution whose non-centrality parameter is proportional to the energy of the fault-induced residual component. The detection threshold is then set to achieve a desired false alarm rate.

This framework is classical and well established in the fault detection literature (Basseville & Nikiforov, 1993). The key insight for the Kenyan context is that the distribution of the residual under healthy operation is largely insensitive to the operating point, provided the model is correctly specified. The filter's prediction error reflects only the process and measurement noise, whose statistics are properties of the asset and its sensors rather than of the external environment.

This stands in contrast to a purely statistical feature such as the root-mean-square value of a vibration signal, whose distribution under healthy operation depends strongly on the load and speed of the machine. A threshold set on the vibration feature at one operating point may produce an unacceptable false alarm rate at another. A threshold set on the physics-based residual does not suffer from this defect to the same degree.

Formal Detection Guarantees

We Now Formalise The Detection Problem And State The Principal

We now formalise the detection problem and state the principal theoretical result concerning the robustness of physics-based residuals. Consider the residual sequence rk generated by a filter operating under the healthy parameter value. We assume that the filter has reached steady state so that the residual covariance is constant in time.

Definition 3.1. The detection statistic at time N is defined as follows.

$$S_N = \sum_{k=1}^{N} r_k^T \Sigma^{-1} r_k \tag{3}$$

where Σ is the steady-state covariance of the residual. Under the null hypothesis of healthy operation, the residual is a zero-mean Gaussian sequence with covariance Σ, and SN follows a central chi-squared distribution with N·m degrees of freedom, where m is the dimension of the measurement vector. Theorem 3.2.

Suppose the asset operates under a distribution shift that changes the input sequence uk but leaves the physical parameters at their healthy values θ_0. If the filter is correctly specified, then the residual sequence remains zero-mean with covariance Σ, and the detection statistic SN retains its central chi-squared distribution under the null hypothesis. Proof.

Definition 3 1

Under the healthy parameter value, the filter's state estimate converges to the true state regardless of the input sequence, provided the system is observable and the filter gain has converged. The innovation sequence, which is the residual, is therefore a martingale difference sequence with the claimed covariance structure. This result follows from the orthogonality principle of the Kalman filter (Anderson & Moore, 1979).

Since the distribution of SN depends only on the first two moments of the residual, and these are invariant to the input sequence, the claim follows. Theorem 3.2 captures the essential robustness property. A purely data-driven feature, whose distribution under healthy operation depends on the operating point, would not satisfy the invariance property.

The physics-based residual, by virtue of being constructed from a model that explicitly accounts for the input, is immune to this source of distribution shift. Corollary 3.3. Under the conditions of Theorem 3.2, the false alarm rate of a threshold test on SN is invariant to changes in the operating environment.

Consequently, a threshold calibrated during commissioning remains valid throughout the asset's operational life, provided the physical parameters do not drift. Remark 3.4. The invariance property degrades gracefully when the model is misspecified.

If the filter's model omits a physical effect that is excited by certain operating conditions, the residual will acquire a bias under those conditions, and the false alarm rate will increase. This observation motivates careful model validation prior to deployment, as emphasised in the broader literature on model-based diagnosis (Gertler, 1998).

Lifecycle Cost Modelling for Maintenance Decisions

The Statistical Properties Of The Detection Statistic Determine The Operating

The statistical properties of the detection statistic determine the operating characteristic of the diagnostic system, but they do not, by themselves, determine the optimal detection threshold. The threshold should be chosen to minimise the expected lifecycle cost, which includes the cost of false alarms, the cost of missed faults, and the cost of maintenance interventions. This economic perspective is often neglected in the technical literature, yet it is decisive in resource-constrained environments where maintenance budgets are limited.

Let CF denote the cost of a false alarm, which includes the cost of dispatching a maintenance crew, the cost of downtime during inspection, and the opportunity cost of diverting resources from other activities. Let CM denote the cost of a missed fault, which includes the cost of catastrophic failure, unplanned outage, collateral damage to adjacent equipment, and the social cost of supply interruption. In the Kenyan context, CM is typically much larger than CF because the consequences of transformer failure include prolonged outages and expensive emergency replacement.

The expected cost per decision epoch depends on the detection threshold τ, the prior probability of fault onset p, and the operating characteristic of the test, which maps the threshold to the false alarm probability $\alpha(\tau)$ and the detection probability $\beta(\tau)$. The expected cost is given by the following expression.

$$J(\tau) = (1-p)\cdot\alpha(\tau)\cdot C_F + p\cdot(1-\beta(\tau))\cdot C_M \quad (4)$$

Let C F Denote The Cost Of A False Alarm

The optimal threshold minimises this expected cost. The first-order condition balances the marginal reduction in expected missed-fault cost against the marginal increase in expected false-alarm cost. The structure of the solution is intuitive: when the cost ratio CM/CF is large, the optimal threshold is low, favouring early detection at the expense of more frequent false alarms.

When the prior probability of fault p is small, the optimal threshold is high, reflecting the fact that most alarms will be false. The trade-off between sensitivity and specificity is fundamental and cannot be eliminated by improved signal processing alone. A more sensitive detection statistic, such as one that integrates information over a longer window, shifts the operating characteristic in a favourable direction, but it does not remove the need to choose a threshold.

The contribution of the physics-based approach is to make the operating characteristic more stable across operating conditions, so that a threshold chosen at commissioning remains near-optimal throughout the asset's life. The lifecycle cost model also illuminates the value of diagnostic information. A diagnostic system that provides early warning of an incipient fault allows the operator to plan the intervention, to procure replacement parts, and to schedule the outage during periods of low demand.

The economic benefit of this lead time is substantial. In the Kenyan context, where transformer lead times can exceed twelve months, the ability to detect a fault early and to order a replacement before the asset fails can reduce the duration of an outage by an order of magnitude.

Asset-Specific Analytical Examples

The General Framework Is Best Understood Through Concrete Examples That

The general framework is best understood through concrete examples that illustrate the interaction between physics, statistics and economics. We consider two asset classes that are critical to Kenyan electricity infrastructure: power transformers and induction motors. For a power transformer, the dominant ageing mechanism is thermal degradation of the paper insulation.

The hot-spot temperature, which drives the ageing process, can be estimated from the load current and ambient temperature using the thermal model described in the IEC loading guide (IEC 60076-7, 2018). A physics-based diagnostic monitors the consistency between the measured top-oil temperature and the temperature predicted by the thermal model. A persistent positive residual indicates that the transformer is running hotter than expected, which may signal a deterioration in cooling efficiency, a partial blockage of oil flow, or an increase in stray losses.

The residual is robust to changes in load because the thermal model explicitly accounts for the load-dependent heat generation. The economic analysis for transformers in Kenya is dominated by the high cost of failure. A catastrophic transformer failure at a substation can interrupt supply to an entire industrial zone, with economic losses that far exceed the replacement cost of the transformer itself.

The optimal detection threshold is therefore low, favouring early intervention even at the cost of occasional false alarms. The cost of a false alarm, which involves an oil analysis and possibly a dissolved gas test, is small relative to the cost of an unplanned failure. For induction motors, a common fault is a broken rotor bar, which produces a characteristic sideband in the stator current spectrum.

For A Power Transformer The Dominant Ageing Mechanism Is Thermal

The physics-based diagnostic computes the expected current spectrum under healthy operation and compares it with the measured spectrum. The residual energy in the sideband frequencies is the detection statistic. The robustness property is particularly valuable here because the sideband amplitudes depend on the motor load.

A purely statistical approach that does not account for load would produce a high false alarm rate when the motor operates at varying loads, as is common in industrial applications such as pumps and compressors. The economic analysis for motors differs from that for transformers because the cost of failure is lower and the asset is more readily replaceable. A failed motor can often be replaced within days, whereas a failed transformer requires months of lead time.

The optimal detection threshold for motors is therefore higher, reflecting the lower cost ratio CM/CF. The diagnostic system is still valuable, but its primary benefit is in avoiding secondary damage, such as damage to the driven equipment or to the motor's bearings, rather than in avoiding the cost of the motor replacement itself.

Institutional and Data Constraints in Kenya

The Theoretical Framework Developed In The Present Analysis Is Necessary

The theoretical framework developed in the present analysis is necessary but not sufficient for the realisation of physics-informed diagnostics in Kenya. Several institutional and data-related constraints must be addressed before the methodology can be deployed at scale. These constraints are not merely practical inconveniences; they shape the feasibility of the entire approach.

The first constraint concerns the availability of accurate asset data. The physics-based approach requires knowledge of the asset's design parameters, including winding resistances, thermal time constants and cooling system characteristics. This information is typically available from the manufacturer's documentation, but in Kenya, as in many developing countries, asset records are often incomplete.

Transformers that have been in service for decades may have lost their original nameplates, and maintenance records may be fragmented across multiple organisations, particularly where assets have been transferred between utilities or from the national grid to county-level distributors. The second constraint concerns the quality and continuity of sensor data. The physics-based residual is only as good as the measurements that feed it.

Many Kenyan substations lack the sensor infrastructure to measure load current, oil temperature and ambient temperature at the temporal resolution required for continuous monitoring. Retrofitting this infrastructure is a capital-intensive undertaking that must be justified by the expected reduction in lifecycle cost. The economic model developed in the present analysis provides a framework for this justification, but the inputs to the model, including the cost of sensors and the probability of fault onset, must be estimated from local data.

The First Constraint Concerns The Availability Of Accurate Asset Data

The third constraint concerns the institutional capacity to act on diagnostic information. A diagnostic system that generates alarms is of no value if the utility lacks the maintenance crews, the replacement parts and the managerial processes to respond to those alarms. In many Kenyan organisations, maintenance is reactive rather than planned, and the introduction of condition-based maintenance requires a cultural shift as much as a technical one.

This observation is consistent with the broader literature on maintenance management in developing countries, which emphasises the importance of organisational factors alongside technical ones (Pintelon & Parodi-Herz, 2008). The fourth constraint concerns the validation of the physics-based models under Kenyan conditions. The thermal models that underpin transformer diagnostics were developed primarily for temperate climates, and their parameters may require adjustment for the high ambient temperatures and strong solar radiation experienced in many parts of Kenya.

Model validation requires access to instrumented assets that can be monitored over extended periods, which in turn requires a commitment from utilities to support research activities. This is a classic collective action problem: the benefits of validated models accrue to all utilities, but the costs of validation fall on the utility that hosts the instrumented assets.

The Choice Between Physics-Based And Purely Data-Driven Diagnostic Approaches Is

The choice between physics-based and purely data-driven diagnostic approaches is not merely a technical preference; it has direct economic consequences. The lifecycle cost model developed earlier provides a framework for comparing the two approaches in terms of their expected costs over the asset's operational life. A purely data-driven approach, such as a neural network classifier trained on historical vibration or current data, has an apparent advantage: it does not require an explicit physical model and can, in principle, learn arbitrary patterns from data.

However, this flexibility comes at a cost. The model must be trained on data that are representative of the deployment conditions, and the scarcity of fault data in the Kenyan context means that the model will be trained primarily on healthy examples. The model may learn to associate certain operating conditions with health, and when those conditions change, the model's output becomes unreliable.

The out-of-distribution problem is not a minor technical nuisance; it is a fundamental limitation of purely empirical approaches in data-sparse settings. The literature on domain adaptation and distribution shift has documented numerous cases where models that perform well on their training distribution degrade catastrophically when deployed under shifted conditions (Quionero-Candela et al., 2009). The physics-based approach mitigates this problem by anchoring the model to the governing equations, which remain valid regardless of the operating environment.

The economic comparison can be made concrete through the lifecycle cost model. Consider two diagnostic systems for a population of distribution transformers. The physics-based system has a stable operating characteristic, with a false alarm probability αP and a detection probability βP that do not vary with operating conditions.

A Purely Data-Driven Approach Such As A Neural Network Classifier

The data-driven system has an operating characteristic that is well calibrated under the training conditions but degrades under distribution shift: the false alarm probability αD increases when the operating conditions deviate from the training distribution, while the detection probability βD may decrease. The expected lifecycle cost of the data-driven system is therefore higher than that of the physics-based system, even if the data-driven system has a superior operating characteristic under the training distribution. The magnitude of the difference depends on the degree of distribution shift and on the cost parameters.

In the Kenyan context, where seasonal and diurnal temperature variations are significant and where loading patterns are subject to rapid change as the economy grows, the distribution shift is likely to be substantial. This analysis suggests that the optimal investment strategy for Kenyan utilities is to prioritise the development of physics-based diagnostic capabilities, while using data-driven methods as a complement rather than a substitute. Data-driven methods can be valuable for anomaly detection in the early stages of a monitoring programme, when the physical models have not yet been validated.

As the monitoring programme matures and the physical models are refined, the diagnostic system should transition towards a physics-based approach.

Conclusion: A Path Towards Reliable Condition Monitoring

The Present Analysis Has Developed A Theoretical Framework For Physics-Informed

The present analysis has developed a theoretical framework for physics-informed diagnostics that addresses the specific constraints of electrical engineering asset management in Kenya. The central contribution is the demonstration that physics-based residuals offer a formal robustness guarantee against out-of-distribution operating conditions, a property that purely data-driven approaches cannot match in data-sparse settings. This robustness has direct economic consequences, as it enables the calibration of detection thresholds that remain valid throughout the asset's operational life.

The lifecycle cost model provides a principled basis for choosing detection thresholds and for justifying investments in diagnostic infrastructure. The model makes explicit the trade-off between false alarms and missed faults, and it shows how the optimal threshold depends on the cost ratio and the prior probability of fault onset. The framework is sufficiently general to accommodate a wide range of asset classes, from power transformers to rotating machines, and it can be extended to include more sophisticated cost structures, such as time-varying costs of failure or the cost of reduced asset life due to delayed intervention.

The theoretical results are necessarily limited by the assumptions of the analysis. The robustness guarantee of Theorem 3.2 depends on the correct specification of the physical model, and the economic analysis depends on the accurate estimation of cost parameters. Both of these requirements point to the need for empirical validation in the Kenyan context.

The Lifecycle Cost Model Provides A Principled Basis For Choosing

The next step is to instrument a small number of critical assets, to collect high-quality sensor data over an extended period, and to validate the thermal and electrical models that underpin the diagnostic framework. This validation effort is a prerequisite for the widespread deployment of physics-informed diagnostics in

Kenya. The argument advanced here is not that physics-based diagnostics are universally superior to data-driven approaches.

Rather, the argument is that the relative value of the two approaches depends on the data environment. In settings where labelled fault data are abundant and the deployment distribution matches the training distribution, purely data-driven methods may be entirely adequate. In settings such as Kenya, where fault data are scarce and operating conditions are diverse and changing, the inductive bias provided by physical knowledge is not a limitation but an asset.

The physics constrains the hypothesis space to models that are consistent with the known behaviour of electrical assets, and this constraint is precisely what enables reliable diagnosis under uncertainty.

References

Anderson, B. D. O., & Moore, J. B. (1979). Optimal filtering. Prentice-Hall.

Basseville, M., & Nikiforov, I. V. (1993). Detection of abrupt changes: Theory and application. Prentice-Hall.

Chen, J., & Patton, R. J. (1999). Robust model-based fault diagnosis for dynamic systems. Kluwer Academic Publishers.

Gertler, J. (1998). Fault detection and diagnosis in engineering systems. Marcel Dekker.

IEC 60076-7. (2018). Power transformers – Part 7: Loading guide for mineral-immersion transformers. International Electrotechnical Commission.

Isermann, R. (2006). Fault-diagnosis systems: An introduction from fault detection to fault tolerance. Springer.

Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Seabrook, M., Murphy, K., & Snoek, J. (2019). Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. In Advances in Neural Information Processing Systems 32 (pp. 13991–14002).

Pintelon, L., & Parodi-Herz, A. (2008). Maintenance: An evolutionary perspective. In K. A. H. Kobbacy & D. N. P. Murthy (Eds.), Complex system maintenance handbook (pp. 21–48). Springer.

Plett, G. L. (2004). Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 2. Modeling and identification. Journal of Power Sources, 134(2), 262–276.

Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., & Lawrence, N. D. (Eds.). (2009). Dataset shift in machine learning. MIT Press.

Thomson, W. T., & Fenger, M. (2001). Current signature analysis to detect induction motor faults. IEEE Industry Applications Magazine, 7(4), 26–34.