A Methodology for Progressive Physics Integration in Data-Driven Anomaly Detection - Application to Circuit Breaker Monitoring

Melvin Fernandes Novo[1,2], Augustin Cathignol[1], Diego Alberto[1], Maya Alalam[1], Benoit Iung[2],Alexandre Voisin[2], Phuc Do[3]

[1] *Schneider Electric, 38000 Grenoble, France*
melvin.fernandesnovo@se.com

[2] *Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France*
melvin.fernandes-novo@univ-lorraine.fr

[3] *SyCoIA, IMT Mines Ales, Ales, France*
phuc.do@mines-ales.fr

Abstract

Physics-informed machine learning has emerged as a promising paradigm for industrial health monitoring, yet practical guidance on when and how to integrate domain knowledge into detection pipelines remains limited. This paper proposes a structured methodology for progressive physics integration in unsupervised anomaly detection, organised into three levels of increasing depth. Level 0 refers to purely data-driven models operating on raw signals. Level 1 injects operational covariates such as temperature or equipment subtype through stratification or conditioning. Level 2 integrates physical knowledge about the internal structure of the signal—its segmentation into electromechanical phases of differing diagnostic relevance—through phase-aware representations, scoring, or end-to-end architectural integration. The methodology is applied systematically to three method families—statistical envelopes, isolation forests and variational autoencoders—for monitoring medium-voltage circuit breaker coil currents, where the breaker's protection-switching function makes condition monitoring critical. At the deepest level, a physics-informed conditional VAE (PicVAE) injects domain knowledge through phase-segmented inputs, FiLM-conditioned architecture, and a phase-weighted reconstruction loss. Validated on real operational data with expert-labelled anomalies, the results reveal two findings: (i) operational conditioning at Level 1 consistently improves detection across all three method families; (ii) structural physics injection at Level 2 has a method-dependent impact, yielding clear benefits for phase-aware representation learning while introducing trade-offs for simpler models. The PicVAE achieves the best overall performance (AUC-ROC = 0.951, Youden J = 0.854). The proposed methodology provides a reproducible template for integrating domain knowledge into anomaly detection pipelines.

1. Introduction

Medium-voltage (MV) circuit breakers are critical assets for the reliability of electrical power distribution networks. Their primary function is to interrupt fault currents on demand, thereby protecting downstream equipment and ensuring service continuity. Switching is achieved by an electromagnetic actuator: dedicated opening and closing coils, when energised by a control voltage, generate the magnetic flux that drives the mechanical operating mechanism through its stroke. The mechanical operating mechanism is responsible for the largest share of breaker failures, with wear and ageing accounting for over 40% of major failure causes (Ito, Richter, le Roux, & Pepper, 2023). Condition monitoring based on coil current analysis has emerged as a practical, non-intrusive solution: the coil current waveform reflects the coupled electromagnetic and mechanical behaviour of the actuator and is sensitive to friction increase, lubrication loss and mechanical wear (Johal & Mousavi, 2008).

Existing diagnostic approaches span a wide range of complexity, from expert-defined characteristic points and statistical envelopes (Rao, Huang, Hu, & Xiao, 2009; Natti & Kezunovic, 2011) to data-driven machine learning (Strachan, McArthur, Stephen, McDonald, & Campbell, 2007; Razi-Kazemi, Vakilian, Niayesh, & Lehtonen, 2015) and deep learning methods—including variational autoencoders (VAEs) trained on healthy operations to detect deviations through reconstruction error (An & Cho, 2015; Jakubowski, Stanisz, Bobek, & Nalepa, 2022). Their common objective is to enable condition-based and predictive maintenance of installed MV equipment.

However, these methods share a common limitation: each integrates domain knowledge in an *ad hoc* manner, without a principled framework for deciding *what* knowledge to inject, *where* in the pipeline, and *how deeply*. Feature-based methods exploit physical knowledge at the input level but cannot learn complex patterns. Statistical envelopes capture global signal shape but ignore the physical meaning of different signal regions. Deep autoencoders learn powerful representations but treat the signal as a generic time series, discarding the rich physical structure embedded in the waveform.

Meanwhile, physics-informed machine learning (PIML) has shown promising results across PHM applications (Karniadakis et al., 2021; Nascimento, Corbetta, Kulkarni, & Viana, 2021), and recent reviews have proposed taxonomies of *what* and *how* to inform data-driven models with physics (Deng, Nguyen, Medjaher, Gogu, & Morio, 2024; Wu, Sicard, & Gadsden, 2024). These reviews, however, predominantly describe *what has been done* rather than providing actionable guidance on *how to proceed* for a new application where qualitative expert knowledge exists but no quantitative degradation equations are available. This is precisely the situation encountered in circuit breaker coil current monitoring: domain experts can identify the physically meaningful regions of the signal and the mechanical phenomena they reflect, but no closed-form model relates these regions to degradation indicators.

This paper addresses this gap by proposing a structured methodology for **progressive physics integration** into unsupervised anomaly detection, organised into three levels of increasing depth, and by empirically illustrating under which conditions each level translates into detection gains. The contributions are:

- A **four-stage hybrid physics–AI methodology** for unsupervised anomaly detection (signal preparation, progressive physics integration, anomaly scoring, deployment), with a **three-level taxonomy** at its core: data-driven baseline (Level 0), operational conditioning (Level 1), and structural physics injection (Level 2).
- A **physics-informed conditional VAE** (PicVAE) integrating domain knowledge simultaneously at the input, architectural, and loss levels.
- A **systematic ablation study** instantiating the taxonomy across three method families—statistical envelopes, isolation forests, and variational autoencoders—and the three integration levels.
- An **empirical illustration** on real operational data from an accelerated ageing campaign: operational conditioning (Level 1) consistently improves detection, whereas structural physics injection (Level 2) shows a method-dependent impact, yielding clear gains for phase-aware representation learning while introducing trade-offs for simpler models.

The remainder of the paper is organised as follows. Section 2 presents the experimental dataset and the physical analysis of the coil current signal. Section 3 details the three-level methodology and its instantiation across the three method families. Section 4 describes the evaluation protocol. Section 5 reports the detection results and the ablation analysis. Section 6 discusses practical implications and limitations, and Section 7 concludes.

2. EXPERIMENTAL DATA

2.1. Studied system

The system studied in this work is a medium-voltage vacuum circuit breaker designed by Schneider Electric for primary distribution networks. Its main characteristics are summarised in Table 1 and identify it as a representative MV breaker with industrial-grade ratings. Its overall structure, including the opening and closing coils that drive the actuation mechanism, is shown in Fig. 1.

Figure 1. Overall structure of the studied medium-voltage vacuum circuit breaker

Table 1. Main characteristics of the studied medium-voltage vacuum circuit breaker.

Characteristic	**Value**
Breaker type	MV vacuum circuit breaker
Rated voltage	up to 24 kV
Rated continuous current	630–2500 A
Rated short-circuit current	up to 40 kA
Operating mechanism	Stored-energy spring drive
Standard temperature range	$-5\,°C$ to $+40\,°C$

2.2. Accelerated ageing campaign and dataset

The data originate from an accelerated ageing campaign conducted on the system described above, combining intensive mechanical cycling with extreme climatic conditions to reproduce the progressive degradation of the operating mechanism in a controlled environment (El Khoury et al., 2025). Both opening and closing coils were operated continuously through automated switching cycles under temperature variations from 25 °C to 70 °C and elevated relative humidity levels, with operating ranges deliberately exceeding the standard envelope (Table 1) to accelerate ageing phenomena affecting coil resistance, lubrication viscosity, spring elasticity and mechanical clearances. The campaign was conducted until mechanical failure, yielding approximately 27 000 coil current operations over three months. For each operation, the ambient temperature θ and the coil type $c \in \{\text{Opening}, \text{Closing}\}$ are recorded. Temperature is a critical covariate: it directly affects coil resistance and therefore the steady-state current (El Khoury et al., 2025), introducing waveform variability that is operationally induced rather than degradation-related.

2.3. Anomaly labelling

Domain experts labelled the final ten days of operation preceding each coil's failure. The resulting test set comprises 1 579 operations, of which 38 are anomalous (2.4%): 20 core jamming events (the actuator fails to complete its stroke), 13 excessive-slowness events (the actuator completes its stroke with abnormally long operation time), and 5 unclassified anomalies. As reported in (El Khoury et al., 2025), corrosion and oxidation of the electromagnet and its frame were identified as the primary failure mechanism. Examples of healthy and anomalous waveforms across the three labelled categories are shown in Fig. 2. All operations preceding the labelling window are treated as healthy training data.

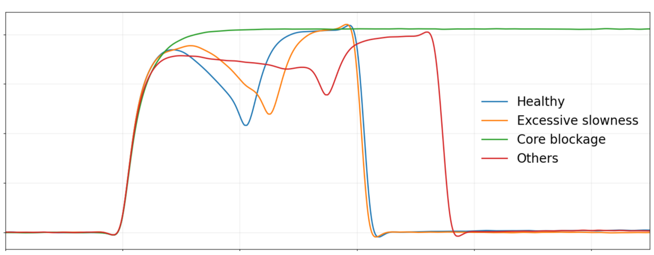

Figure 2. Examples of healthy and anomalous coil current waveforms from the ageing campaign.

3. HYBRID PHYSICS-AI FRAMEWORK

At the core of this paper is a structured framework that guides *how* and *where* domain knowledge can be injected into anomaly-detection pipelines. Rather than focusing on a specific model, the proposed methodology organises physics integration as a progressive process, in which physical information is introduced at increasing depths of the pipeline. This framework, illustrated in Fig. 3, structures the analysis from signal preprocessing to deployment considerations. The different integration levels and their concrete instantiations are detailed in the following sections.

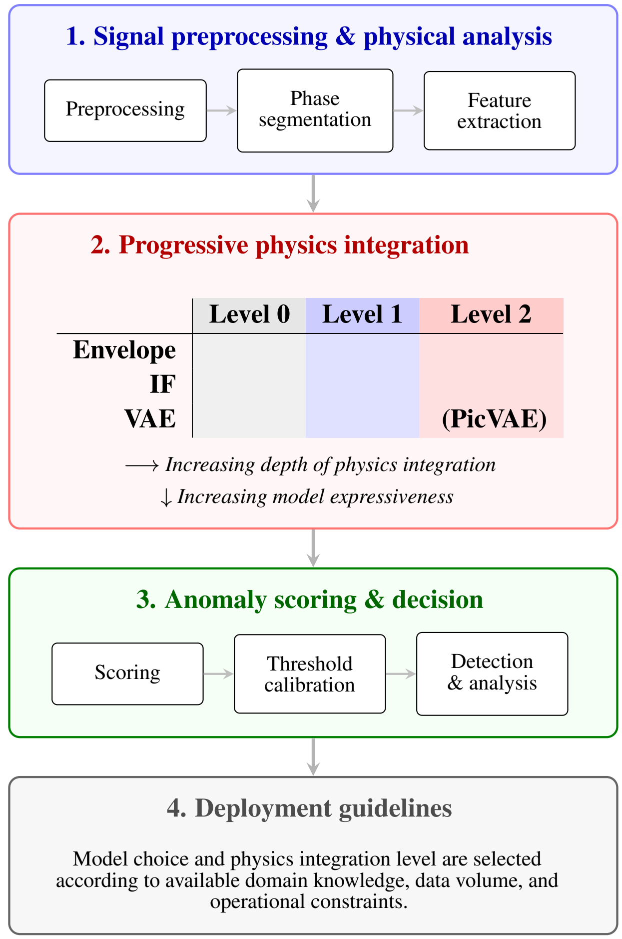

Figure 3. Overview of the proposed methodology. Method families in stage 2 are ordered by structural flexibility (top to bottom); the PicVAE corresponds to the deepest configuration of the most flexible family.

3.1. Signal preprocessing and physical analysis

Before any learning or physics integration, all signals undergo a common preprocessing stage, shown in Fig. 3. This stage is deliberately separated from the integration levels, as it corresponds to signal conditioning rather than diagnostic modelling.

First, raw coil current waveforms are temporally aligned using the command triggering time, ensuring a consistent reference across operations. The signals are then smoothed to reduce measurement noise while preserving the electromechanical dynamics of interest. Finally, a phase segmentation is applied based on expert knowledge of the actuator mechanism. Importantly, this preprocessing does not use any anomaly labels and does not introduce diagnostic decisions. It provides a physically interpretable representation of the signal that can be exploited by subsequent integration levels.

3.2. Progressive physics integration

This section introduces the proposed progressive physics integration strategy and its instantiation across the considered method families. Three integration levels are defined, corresponding to the increasing depths of domain knowledge incorporated into the anomaly-detection pipeline. Each level can be instantiated using different modelling paradigms such as statistical envelopes, isolation forests, and variational autoencoders.

Table 2. Evaluated configurations grouped by physics integration level. Across all methods, Level 2 relies on the same physical information: four-phase segmentation of the coil-current waveform, six characteristic points, and expert-defined phase relevance. The difference lies in how this information is absorbed by each method.

Configuration	**Physics integration and modelling role**
Reference (no learning)	
Characteristic points	Direct comparison of expert-selected signal landmarks against predefined thresholds; purely rule-based diagnosis
Level 0 — Data-driven baseline	
Envelope	Global statistical tolerance band capturing average signal shape without distinguishing operating conditions or physical phases
IF	Unsupervised anomaly detection on generic time-domain descriptors treating the signal as an unstructured entity
VAE	End-to-end reconstruction of the raw signal, learning a generic latent representation of nominal behaviour
Level 1 — Operational conditioning	
Envelope	Condition-specific tolerance bands isolating normal variability due to temperature and coil type
IF	Separate unsupervised models per operating regime to avoid mixing distinct nominal behaviours
VAE	Conditioned representation learning enabling operating-mode-aware reconstruction of the signal
Level 2 — Structural physics injection	
Envelope	Phase-weighted anomaly scoring emphasising diagnostically relevant actuation phases using expert-defined importance factors
IF	Aggregation of phase-specific statistics and characteristic points reflecting the physical actuation sequence (weights applied implicitly)
VAE	Phase-aware representation learning guided by physically meaningful signal regions and phase-weighted reconstruction

A steady-state thermal variant (Envelope L1+) is additionally reported in Section 3.2.3.

Table 2 provides a unified overview of the evaluated configurations, summarising, for each method family and integration level, how physical knowledge is incorporated. The following subsections detail the rationale and implementation of each level, starting from a physics-only baseline, and progressing from purely data-driven models to deep structural physics integration.

3.2.1. Feature expert baseline

Before introducing the three integration levels, we consider a purely physics-based method that serves as a lower-bound reference against which all learning-based approaches are evaluated. This approach relies exclusively on expert knowledge and exploits extracted signal landmarks directly, without any statistical learning. For each operation, the amplitudes and temporal positions of these landmarks are compared against expert-defined thresholds to flag abnormal behaviour. Because this baseline requires no training data, it can be deployed immediately on newly installed equipment. It therefore provides a meaningful reference for assessing whether data-driven models effectively leverage physics beyond explicit expert rules.

3.2.2. Level 0 — Data-driven baseline

Level 0 defines a purely data-driven reference in which no domain knowledge is injected beyond standard signal preprocessing. The signal is treated as a generic time series, and anomaly detection relies solely on statistical or representation-learning mechanisms to identify deviations from nominal behaviour. This level establishes a baseline against which the impact of deeper physics integration can be quantified.

At this level, different modelling paradigms can be used without modification, including statistical envelopes, isolation forests, and autoencoder-based architectures. Their respective implementations are summarised in Table 2.

3.2.3. Level 1 — Operational conditioning

Level 1 incorporates knowledge about the operating conditions under which signals are acquired, while still treating the signal as a generic time series. The objective is to explain and absorb benign variability caused by known covariates—such as temperature or coil type—that affect the waveform but are not indicative of degradation.

Conditioning the model on these variables enables different nominal operating regimes to be distinguished and reduces false positives caused by normal operational variations. Depending on the method family, this conditioning can be implemented through stratification or through learned modulation mechanisms (Perez, Strub, de Vries, Dumoulin, & Courville, 2018).

For envelope-based methods, Level 1 is realised by constructing separate statistical tolerance bands for discrete operating-condition strata. In practice, the training data are partitioned into a small number of temperature bins per coil type, chosen to balance physical relevance and statistical robustness. An enhanced variant, denoted L1+, further restricts the training set to steady-state thermal conditions, yielding tighter envelopes and improved discrimination by filtering out transient thermal effects.

For learning-based approaches, operating conditions are incorporated either by training separate models per regime or by conditioning a single model through learned embeddings, as summarised in Table 2.

3.2.4. Level 2 — Structural physics injection

Level 2 integrates qualitative knowledge about the internal physical structure of the coil-current signal. From an electromechanical standpoint, the actuation sequence can be decomposed into a small number of consecutive phases corresponding to distinct physical phenomena (e.g. energisation, motion initiation, latch engagement and steady-state conduction). Domain experts can further assess which of these phases are most sensitive to mechanical degradation.

At this level, all methods are provided with the same physical information: a segmentation of the signal into electromechanical phases, a set of characteristic points extracted at key transitions, and expert-defined relevance of each phase for diagnosis. The difference between methods lies in how this information is exploited.

For envelope-based approaches, Level 2 modifies the anomaly score by emphasising deviations occurring in physically informative phases, thereby preserving the simplicity of statistical modelling while introducing phase awareness. For isolation forests, phase information is incorporated through the construction of phase-specific descriptors and aggregated statistics reflecting the temporal structure of the actuation sequence.

For the VAE, Level 2 corresponds to a genuinely phase-aware modelling strategy, as illustrated in Fig. 4. Expert features and phase embeddings are injected into the encoder and decoder, enabling the latent representation to jointly encode signal morphology and physical context.

Concretely, the phase labels, characteristic points, and operating conditions are combined into a conditioning vector $\mathbf{v}_i$ that is injected at three points of the network. In the encoder and decoder, conditioning is performed through Feature-wise Linear Modulation (FiLM) (Perez et al., 2018):

$$\mathrm{FiLM}(\mathbf{h}, \mathbf{v}_i) = \gamma(\mathbf{v}_i) \odot \mathbf{h} + \beta(\mathbf{v}_i), \tag{1}$$

where $\mathbf{h}$ is a layer activation and $\gamma(\cdot), \beta(\cdot)$ are learnt affine transformations of $\mathbf{v}_i$. At the latent level, conditioning is performed by concatenating $\mathbf{v}_i$ to the sampled latent code before it is passed to the decoder, providing explicit access to the physical context during reconstruction.

In parallel, the reconstruction objective term (Eq. 2) is adapted such that errors in diagnostically relevant phases contribute more strongly to the anomaly score.

$$\mathcal{L}_{\text{phys}} = \frac{1}{T}\sum_{t=1}^{T} w(z_i(t)) \cdot (x_i(t) - \hat{x}_i(t))^2, \tag{2}$$

Here $z_i(t)$ denotes the electromechanical phase at time step t and $w(\cdot)$ assigns a diagnostic relevance weight to each phase. These weights were defined based on expert knowledge of the actuation mechanism and further refined empirically to reflect their relative sensitivity to mechanical degradation. This design explicitly guides the representation learning process toward physically meaningful signal regions, while retaining the end-to-end nature of the model. The full architecture of the PicVAE is shown in Fig. 4.

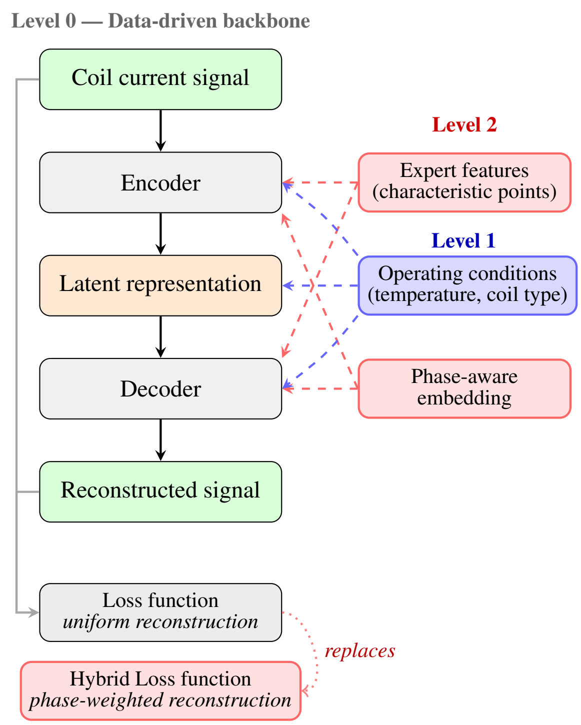

Figure 4. Three-level progressive physics integration, illustrated on the VAE family. The vertical backbone (gray) is the Level 0 data-driven model, trained with a uniform reconstruction loss (bottom). Level 1 (blue, dashed) conditions the encoder, latent, and decoder on operational covariates such as temperature and coil type. Level 2 (red, dashed) injects expert physical knowledge: handcrafted features at the input, phase-aware information at the architectural level, and a phase-weighted reconstruction loss replacing the Level 0 loss.

4. EVALUATION PROTOCOL

Following the unsupervised anomaly detection paradigm, all models are trained exclusively on healthy operations—no anomalous samples are used during training. Models are evaluated on the expert-labelled test set described in Section 2.

Performance is assessed using the AUC-ROC as the primary threshold-independent metric. At the operating-point level, the Youden index $J = \mathrm{TPR} - \mathrm{FPR}$ is maximised over a sweep of 24 percentiles (30th to 99.99th) applied to a calibration set consisting of the last 30% of training data, ordered chronologically. This ensures that the calibration data reflects the most recent ageing state, closest to test-time conditions. Detection recall is additionally reported per anomaly type to assess whether physics integration selectively improves the detection of specific failure modes.

The purely rule-based Characteristic Points baseline is included as a reference that requires no training data and whose performance is by definition invariant to the training set size.

5. RESULTS

5.1. Overall detection performance

Table 3 reports the detection performance at the Youden-optimal operating point for all evaluated configurations. Physics integration does not lead to monotonic performance gains across all methods and levels; instead, its impact depends on both the nature of the injected knowledge and the way it is absorbed by each method.

Table 3. Detection performance at the Youden-optimal operating point. Best value per column in bold. Within each method family, the level of physics integration is indicated by the suffix L0/L1/L2.

Configuration	AUC	J	TPR	FPR
Reference (no learning)				
Char. Points	0.913	0.759	0.842	0.083
Envelope-based methods				
Envelope L0	0.917	0.791	0.868	**0.078**
Envelope L1	0.947	0.822	0.921	0.099
Envelope L1+	0.950	0.841	0.921	0.080
Envelope L2	0.942	0.799	0.895	0.096
Isolation Forest (IF)				
IF L0	0.898	0.773	0.895	0.121
IF L1	0.914	0.792	0.947	0.156
IF L2	0.928	0.720	0.868	0.149
Variational autoencoder-based methods				
VAE L0	0.940	0.802	0.895	0.093
VAE L1	0.942	0.833	0.921	0.088
VAE L2 (PicVAE)	**0.951**	**0.854**	**0.974**	0.119

For envelope-based methods, incorporating operational information (Level 1) yields the largest improvement, with the Youden index increasing from 0.791 (L0) to 0.841 (L1+). In contrast, the phase-aware Level 2 formulation does not provide additional benefit for envelopes and slightly degrades performance, indicating that deeper structural constraints are not systematically advantageous for this family. Isolation Forest methods exhibit a similar trend: while IF L2 achieves a higher AUC (0.928), its Youden index decreases relative to IF L1, reflecting an unfavorable trade-off between sensitivity and false positives.

For the VAE family, performance consistently improves with the depth of physics integration. Moving from Level 0 to Level 2 yields progressive gains in both AUC and Youden index, with the proposed PicVAE achieving the best overall performance (AUC = 0.951, J = 0.854).

These results indicate that physics integration is most beneficial when it is matched to the modelling paradigm. While shallow integration (Level 1) effectively improves performance across all method families, deeper structural integration (Level 2) proves advantageous only when the model can exploit this information through an appropriate representation and scoring mechanism as shown in Fig. 5. All results should be interpreted in the context of the considered experimental setting, notably the limited number of annotated anomalies.

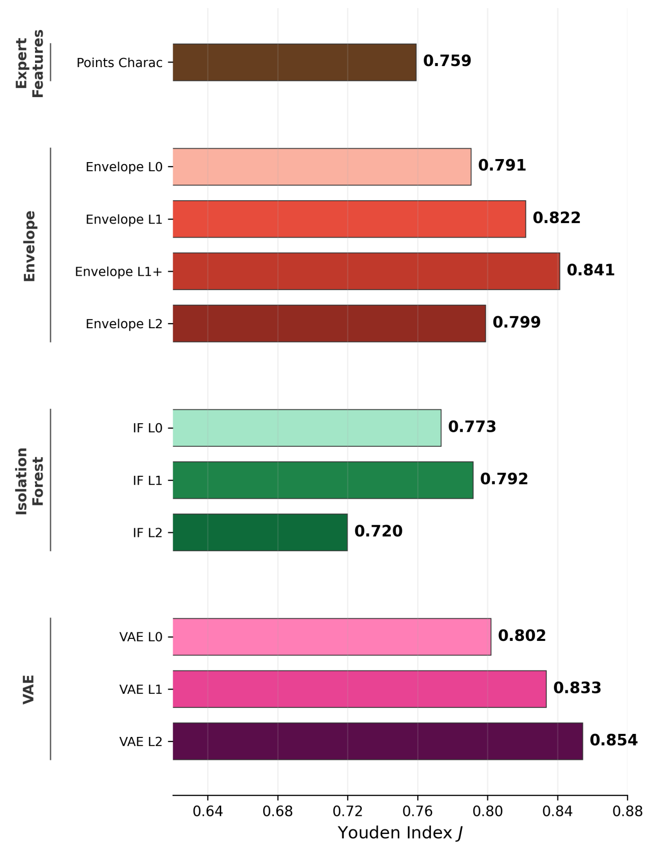

Figure 5. Youden index J by configuration. The effect of physics integration depends on both the integration level and the modelling approach: operational conditioning (Level 1) improves performance across all method families, while deeper structural integration (Level 2) leads to method-dependent trade-offs.

5.2. Detection by anomaly type

Figure 6 reports the recall achieved for each anomaly category at the Youden-optimal operating point.

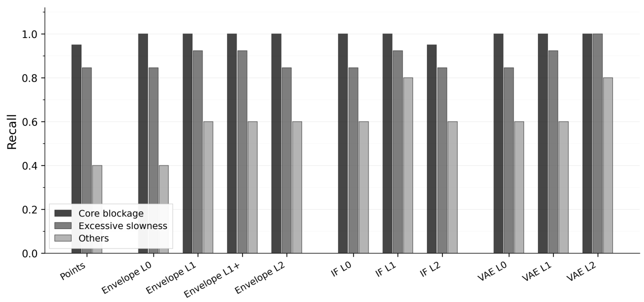

Figure 6. Detection recall by anomaly type. The PicVAE achieves perfect recall on blockages and excessive slowness, and the highest recall on the Others category (0.80 vs. 0.40 for the Characteristic Points baseline).

All configurations, except for the feature expert baseline, detect core blockage events perfectly (recall = 1.0). These faults induce large and abrupt deviations in the coil-current waveform, which are easily captured regardless of the underlying modelling approach.

Differences between configurations emerge for more subtle degradation patterns. For excessive slowness anomalies, most data-driven methods achieve recall values between 0.85 and 0.92, while the PicVAE is the only configuration achieving perfect recall (1.0). The Characteristic Points baseline reaches a recall of 0.85, indicating that expert-defined rules alone may miss some gradual degradation signatures (Fig. 6).

The largest variability across methods is observed for the rare *Others* category. Recalls range from 0.40 for simpler approaches (Characteristic Points, Envelope L0) up to 0.80 for the PicVAE, corresponding to a twofold improvement. This suggests that explicitly incorporating phase-aware physical structure improves sensitivity to localized or atypical anomalies that are difficult to capture with global or unstructured representations. However, given the limited number of samples in this category, these differences should be interpreted cautiously.

5.3. Interpretability through physics integration

Beyond quantitative metrics, the physics-informed architecture provides inherent interpretability absent from data-driven approaches. The phase segmentation identifies *where* in the signal an anomaly manifests (e.g. the latch engagement phase for mechanical degradation). The characteristic points indicate *which* mechanical events are affected (e.g. delayed engagement, reduced actuation force). The phase-weighted reconstruction error quantifies the *severity* relative to physically meaningful signal regions. This structured interpretability is particularly relevant in industrial contexts, where maintenance decisions require understanding the nature and origin of a detected anomaly, not only its presence.

6. DISCUSSION

6.1. Practical guidance for physics integration

The experimental results demonstrate that progressive physics integration is feasible across a wide range of modelling approaches and can improve anomaly-detection performance when appropriately matched to the modelling paradigm and the nature of the injected knowledge. Rather than advocating a single optimal method, the proposed framework provides practitioners with a structured way to reason about how domain knowledge can be incorporated into monitoring pipelines. Based on the presented results, the following practical guidelines can be derived:

1. *Start with a physics-only baseline* when little or no training data are available. Expert-defined rules applied to physically meaningful signal features provide immediate monitoring capability and establish a reference performance. In the considered case, this baseline achieved an AUC of 0.913 without any learning.
2. *Incorporate operational covariates* that explain benign variability unrelated to degradation, such as temperature or equipment subtype. Level 1 integration, through stratification or conditioning, proved effective across all method families by improving performances.
3. *Integrate structural physical knowledge with care*. Level 2 allows all considered methods to exploit information about the internal physical structure of the signal, such as phase segmentation and diagnostic relevance. However, the impact of this deeper integration depends on how the method absorbs and uses this information. Simple approaches benefit from phase awareness at the scoring or feature aggregation stage, whereas more flexible models can incorporate it in an end-to-end manner.
4. *Align integration depth with modelling goals*. Deeper physics integration does not systematically yield better performance and may even be counterproductive when the modelling mechanism is not well suited to exploit structural constraints. Consequently, the desired level of interpretability, robustness, and deployment complexity should guide the integration strategy rather than maximizing integration depth alone.

These guidelines are transferable to other electromechanical monitoring applications where domain knowledge is available, such as motor current analysis, relay diagnostics, or actuator monitoring.

6.2. Limitations and future work

The presented evaluation is subject to several limitations. The dataset originates from a single circuit breaker model with two coil types (El Khoury et al., 2025). It includes a limited number of annotated anomalous operations, which restricts the statistical power of per-category analyses. As a result, differences between configurations—particularly for rare anomaly types—should be interpreted cautiously.

Several research directions are currently under investigation to address these limitations. Ongoing work focuses on validating the proposed methodology on additional endurance-test campaigns and device variants, which will allow a more robust assessment of generalisation across operating conditions and equipment designs. We also investigate the sensitivity of the results to data availability, exploring different training-set sizes and temporal subsampling strategies to derive practical deployment guidelines for newly installed equipment.

Finally, extending the framework toward deeper and more quantitative physics integration represents a key direction for future work. This includes incorporating simplified electromagnetic or mechanical models as constraints or regularisation terms, with the objective of further strengthening the inductive bias while preserving applicability in industrial settings.

7. CONCLUSION

This paper introduced a structured methodology for progressive physics integration in unsupervised anomaly detection and applied it to medium-voltage circuit breaker coil current monitoring. Three integration levels were defined and instantiated across multiple modelling paradigms, complemented by a purely physics-based reference derived from expert knowledge.

The experimental results highlight that integrating physics is not a one-size-fits-all process. While incorporating operational covariates (Level 1) consistently improves detection performance across all method families, deeper structural physics integration (Level 2) yields benefits only when the modelling approach can appropriately absorb and exploit this information. In particular, phase-aware representation learning enabled the proposed PicVAE to achieve the best overall performance on the considered dataset, while simpler methods exhibited trade-offs between sensitivity and false positives at the deepest level.

Beyond performance, the proposed framework provides structured interpretability by linking detected anomalies to physically meaningful regions and events in the actuation sequence. This capability is especially relevant in industrial settings, where understanding the origin and nature of detected anomalies is essential for maintenance decision-making.

Overall, the proposed three-level methodology offers a reproducible and pragmatic template for incorporating qualitative domain knowledge into data-driven monitoring systems. Future work will validate these findings on additional endurance-test campaigns and explore the integration of quantitative electromagnetic or mechanical models as inductive biases, broadening the framework's applicability across the wider class of electromechanical monitoring problems.

NOMENCLATURE

AUC	Area Under Curve
FPR	False Positive Rate
IF	Isolation Forest
J	Youden Index
MV	Medium-voltage
$PicVAE$	Physics-informed conditional VAE (or VAE L2)
ROC	Receiver Operating Characteristic
TPR	True Positive Rate
VAE	Variational AutoEncoder

REFERENCES

An, J., & Cho, S. (2015). Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, *2*(1), 1–18.

Deng, W., Nguyen, K. T. P., Medjaher, K., Gogu, C., & Morio, J. (2024). Physics-informed machine learning in prognostics and health management: State of the art and challenges. *Applied Mathematical Modelling*, *124*, 325–352. doi: 10.1016/j.apm.2023.07.011

El Khoury, D., Meley, J.-P., Alberto, D., Bernard, Z., Alalam, M., & Cathignol, A. (2025). Accelerated aging tests for smart mv circuit breaker predictive maintenance. In (Vol. 2025, p. 3208-3212). doi: 10.1049/icp.2025.2287

Ito, H., Richter, M., le Roux, R., & Pepper, D. (2023). CIGRE fourth reliability survey on switching equipment. In *Cigre session 2023.* (Paper No. 105)

Jakubowski, J., Stanisz, P., Bobek, S., & Nalepa, G. J. (2022). Anomaly detection in asset degradation process using variational autoencoder and explanations. *Sensors*, *22*(1), 291. doi: 10.3390/s22010291

Johal, H., & Mousavi, M. J. (2008). Coil current analysis method for predictive maintenance of circuit breakers. In *Ieee/pes transmission and distribution conference and exposition* (pp. 1–6). doi: 10.1109/TDC.2008.4517278

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, *3*(6), 422–440. doi: 10.1038/s42254-021-00314-5

Nascimento, R. G., Corbetta, M., Kulkarni, C. S., & Viana, F. A. (2021). Hybrid physics-informed neural networks for lithium-ion battery modeling and prognosis. *Journal of Power Sources*, *513*, 230526. doi:

10.1016/j.jpowsour.2021.230526

Natti, S., & Kezunovic, M. (2011). Assessing circuit breaker performance using condition-based data and bayesian approach. *Electric Power Systems Research*, *81*(9), 1796-1804. doi: https://doi.org/10.1016/j.epsr.2011.04.010

Perez, E., Strub, F., de Vries, H., Dumoulin, V., & Courville, A. (2018). FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 32). doi: 10.1609/aaai.v32i1.11671

Rao, P., Huang, J., Hu, X., & Xiao, Y. (2009). Testing of circuit breakers using coil current characteristics analysis. In *Ieee international conference on control and automation* (pp. 185–189). doi: 10.1109/ICCA.2009.5410549

Razi-Kazemi, A. A., Vakilian, M., Niayesh, K., & Lehtonen, M. (2015). Data mining of online diagnosed waveforms for probabilistic condition assessment of circuit breakers. *IEEE Transactions on Power Delivery*, *30*(3), 1354–1362. doi: 10.1109/TPWRD.2015.2399454

Strachan, S. M., McArthur, S. D. J., Stephen, B., McDonald, J. R., & Campbell, A. (2007). Providing decision support for the condition-based maintenance of circuit breakers through data mining of trip coil current signatures. *IEEE Transactions on Power Delivery*, *22*(1), 178–186. doi: 10.1109/TPWRD.2006.883001

Wu, Y., Sicard, B., & Gadsden, S. A. (2024). Physics-informed machine learning: A comprehensive review on applications in anomaly detection and condition monitoring. *Expert Systems with Applications*, *255*, 124678. doi: 10.1016/j.eswa.2024.124678