Progressive Physics-AI Hybrid Methodology for Unsupervised Anomaly Detection in Electromechanical Systems

Melvin Fernandes Novo[1,2], Augustin Cathignol[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: data-driven baselines (Level 0), operational conditioning (Level 1), and structural physics injection (Level 2). The methodology is designed for systems where qualitative expert knowledge is available but no quantitative degradation model exists. It is applied systematically across three method families—statistical envelopes, principal component analysis, isolation forests, deterministic autoencoders, and variational autoencoders—for the monitoring of medium-voltage circuit breaker coil currents. At the highest integration level, a physics-informed conditional variational autoencoder (PicVAE) incorporates 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: operational conditioning at Level 1 consistently improves detection across all method families, while structural physics injection at Level 2 has a method-dependent impact, yielding clear gains for phase-aware representation learning while introducing trade-offs for simpler models. The PhD plan extends this work along three axes: quantitative physics integration on a second use case, cross-domain validation on public benchmarks through a human-in-the-loop knowledge acquisition protocol, and consolidation of the methodology into a transferable deployment framework.

1. PROBLEM STATEMENT

Physics-informed machine learning (PIML) has emerged as a powerful paradigm for industrial health monitoring, combining the flexibility of data-driven methods with the robustness of domain knowledge. Comprehensive reviews have proposed taxonomies describing *what* and *how* to inform models with physics (Deng, Nguyen, Medjaher, Gogu, & Morio, 2024; Wu, Sicard, & Gadsden, 2024; Karniadakis et al., 2021). However, a fundamental gap persists between these general taxonomies and their practical application: existing works describe what has been achieved on specific systems, but provide limited *actionable guidance* on how to proceed when facing a new monitoring application—particularly when qualitative expert knowledge is available but no quantitative degradation model exists.

This gap is especially acute for electromechanical systems such as medium-voltage circuit breakers, where the operating mechanism is responsible for over 40% of major failures (Ito, Richter, le Roux, & Pepper, 2023). The coil current waveform—a non-intrusive, low-cost diagnostic signal—reflects the coupled electromagnetic and mechanical dynamics of the actuator and is sensitive to friction, lubrication loss, and mechanical wear (Johal & Mousavi, 2008). Rich expert knowledge exists about the physical meaning of different signal regions, yet this knowledge is rarely exploited in modern detection pipelines. Feature-based methods use it at the input level but cannot learn complex patterns (Rao, Huang, Hu, & Xiao, 2009). Deep autoencoders learn powerful representations but treat the signal as a generic time series (An & Cho, 2015). No approach has systematically investigated *how deeply* physics should be integrated, *where* in the pipeline, and *under what modelling assumptions* the integration provides genuine benefit.

2. EXPECTED NOVEL CONTRIBUTIONS

This PhD aims to establish a general methodology for progressive physics integration in unsupervised anomaly detection, validated on real industrial data and designed to be transferable across electromechanical monitoring applications. The expected contributions are:

1. **A multi-level physics integration methodology**, organising the incorporation of domain knowledge at increasing depth: Level 0 (data-driven baseline), Level 1 (operational conditioning), Level 2 (structural physics injection), Level 3 (quantitative physics integration via analytical models).
2. **A physics-informed deep generative model** (PicVAE) integrating qualitative domain knowledge through phase-aware architectural conditioning and a phase-weighted reconstruction loss.
3. **Empirical characterisation of the method-dependent value of physics integration**, based on a systematic ablation across several method families and four integration levels.
4. **Cross-domain validation through expert knowledge acquisition**, leveraging a human-in-the-loop protocol on public benchmarks.
5. **Deployment guidelines** accounting for available physical knowledge, model complexity, and data conditions.

3. RESEARCH PLAN

3.1. Completed work

Framework instantiation across three method families. Levels 0 through 2 of the taxonomy have been implemented across statistical envelopes, principal component analysis, isolation forests, deterministic autoencoders, and variational autoencoders (Fig. 1). At Level 1, operational conditioning is performed through temperature and coil-type stratification or Feature-wise Linear Modulation (FiLM) (Perez, Strub, de Vries, Dumoulin, & Courville, 2018). At Level 2, physical structure is injected through phase-aware scoring, phase-specific feature engineering, or direct architectural conditioning, depending on the model family. At the deepest level, the PicVAE embeds phase segmentation, characteristic-point features, and a phase-weighted reconstruction loss into a conditional variational autoencoder architecture.

Application to circuit breaker coil currents. The methodology has been developed and validated on a first industrial use case: unsupervised anomaly detection in medium-voltage circuit breaker coil current signals. An accelerated ageing campaign yielded approximately 27 000 time series from two coil types under temperature variations from 25°C to 70°C until mechanical failure. A physical analysis of the signal morphology identified distinct electromechanical phases and characteristic points. Domain experts labelled 38 anomalies (20 core blockages, 13 excessive slowness events, 5 other defects) in the final days of operation preceding each coil's failure.

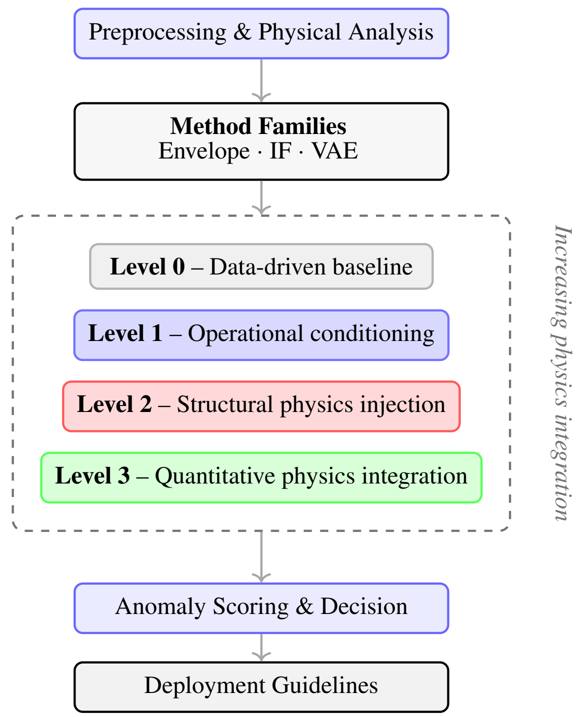

Figure 1. Progressive physics integration pipeline. The four integration levels (L0–L3), grouped within the dashed frame, represent increasing incorporation of domain knowledge across all method families.

Systematic ablation results. Experiments conducted across the three method families and three integration levels reveal two methodologically important findings. Results for three representative families are reported in Table 1; . First, *operational conditioning at Level 1 consistently improves detection* across every method family, confirming that adapting models to known sources of benign variability is universally beneficial. Second, *structural physics injection at Level 2 has a method-dependent impact*: it provides clear gains for the VAE family (which can exploit phase-aware representation learning), while introducing trade-offs for envelope-based and partition-based methods, whose internal mechanisms cannot fully absorb structural priors. The PicVAE achieves the highest overall detection performance (AUC = 0.951, J = 0.854), with a doubling of recall on rare anomalies relative to a purely expert-based reference (0.80 versus 0.40).

Table 1. 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

3.2. Ongoing work

Cross-equipment validation. A second accelerated ageing campaign on a separate batch of tripping coils is underway. This independent dataset will assess whether the methodology and the PicVAE architecture generalise beyond the original equipment instance.

Journal extension. An extended version of the methodological work is being prepared for submission to a PHM-oriented journal, incorporating additional architectural analyses (notably the deterministic-versus-variational comparison at increasing integration depth) and a refined characterisation of the interplay between regularisation and physics injection.

3.3. Future work

Level 3: quantitative physics integration. A second industrial use case, focused on thermal anomaly detection in the cooling system of power converters, provides access to a validated quantitative thermal model. This will enable an extension of the framework to Level 3, incorporating analytical equations as inductive constraints, and a direct comparison between qualitative (Level 2) and quantitative (Level 3) physics integration.

Cross-domain validation through human-in-the-loop knowledge acquisition. To assess transferability beyond electromechanical systems, the framework will be applied to public time-series anomaly detection benchmarks (UCR Anomaly Archive, NASA SMAP/MSL, Server Machine Dataset). For each benchmark, expert knowledge about the signal structure and operating context will be elicited through a structured human-in-the-loop protocol, allowing the same taxonomy to be instantiated on heterogeneous domains.

Methodology consolidation. The ultimate objective is to consolidate the multi-level methodology into a transferable decision framework, validated across at least two industrial applications with different types of available physics and across cross-domain benchmarks, providing actionable guidance to PHM practitioners.

3.4. Timeline

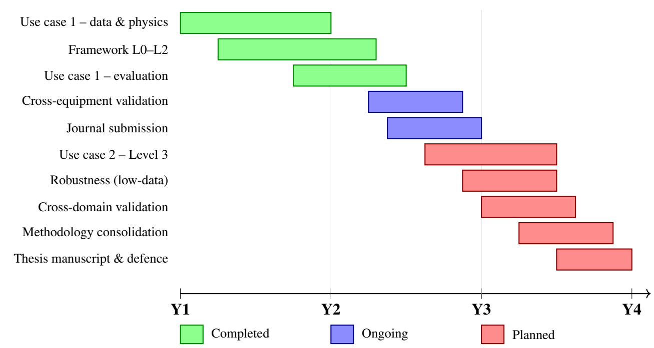

Figure 2. PhD timeline including methodological development, robustness analysis under limited data, and cross-domain validation.

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

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

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

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

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