



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Physics-Based Deep Learning for Online and Onboard Prediction of Insulation Degradation Under Variable Stress Conditions

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Seri, P., Trotta, A., Schafberger, T., Reil, C., Cavallini, A. (2025). Physics-Based Deep Learning for Online and Onboard Prediction of Insulation Degradation Under Variable Stress Conditions. IEEE [10.1109/CEIDP61707.2025.11218381].

Availability:

This version is available at: <https://hdl.handle.net/11585/1036334> since: 2026-01-12

Published:

DOI: <http://doi.org/10.1109/CEIDP61707.2025.11218381>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

Physics-Based Deep Learning for Online and Onboard Prediction of Insulation Degradation Under Variable Stress Conditions

P. Seri¹, A. Trotta¹, T. Schafberger², C. Reil², A. Cavallini¹

¹ University of Bologna, Bologna, Italy

² University of Applied Sciences Amberg-Weiden, Amberg, Germany
paolo.seri2@unibo.it

Abstract— This study presents a novel approach combining physics-informed models with machine learning to predict the residual life of transformer insulation. By integrating partial differential equations (PDEs) to simulate degradation processes and employing Convolutional Neural Networks (CNNs) trained on synthetic data generated through Monte Carlo simulations, the model achieves an accuracy within $\pm 25\%$ error margin. While promising, initial results suggest room for improvement by refining simulation parameters and exploring hybrid models.

Index Terms— Insulation Lifetime Prediction, Machine Learning, Convolutional Neural Networks

I. INTRODUCTION

The degradation of insulation systems is influenced by simultaneous stresses: electrical, thermal, mechanical, and environmental. Physics-based models excel at simulating these interactions over time to understand their impact on material properties. The accuracy of those models often come at the expense of higher complications, and modelling materials under different circumstances often imply both a high number of modelling parameters, and tests necessary to estimate them, when unavailable from literature.

Deep learning could be seen as an alternative to this approach, which enhances the capabilities of a model by processing complex datasets reflecting real-world conditions, uncovering hidden patterns and correlations within large volumes of data.

In the realm of power systems, the potential of machine learning (ML) is being investigated for predictive maintenance strategies [1],[2]. By analyzing historical data and real-time sensor inputs, ML algorithms can anticipate equipment failures before they occur, reducing downtime and extending asset lifespans. These capabilities are further extended to improve load forecasting accuracy, which is critical for efficient resource allocation and minimizing operational costs. Additionally, as renewable energy sources become more prevalent in the grid, machine learning helps manage their variable outputs by predicting fluctuations in generation from solar or wind farms, facilitating smoother grid operations and better energy storage strategies.

Smart grids represent a significant leap forward in electrical engineering, integrating digital communication technologies

with traditional power systems. Here, machine learning is essential for developing advanced monitoring, control, and optimization capabilities. For instance, anomaly detection algorithms can swiftly identify irregularities in grid operations, allowing for quick responses to prevent outages or blackouts. Machine learning also enhances demand response programs by dynamically balancing supply and demand through real-time data analysis, which is crucial as the energy landscape evolves. The rise of electric vehicles has prompted innovations in intelligent charging systems where machine learning optimizes charging schedules to balance load across the grid effectively. By predicting user behavior and vehicle requirements, ML models can minimize peak demand impacts and improve the efficiency of charging stations. Predictive analytics also play a key role in EV battery management, ensuring optimal performance and longevity by monitoring health indicators and recommending proactive maintenance actions.

Machine learning is also transforming fault detection and diagnosis into electrical engineering. By analyzing sensor data for deviations from normal conditions, ML can identify faults early, helping prevent failures and reduce maintenance costs. However, challenges remain, such as the need for large training datasets and the difficulty of modeling complex, nonlinear physical behaviors. Ongoing research is tackling these issues with advanced algorithms that better capture the nuances of insulation degradation.

A key challenge in applying machine learning to electrical engineering is ensuring model interpretability, especially with complex methods like deep neural networks. While these models offer high prediction accuracy, their lack of transparency can hinder trust—particularly in safety-critical areas like power grid management.

Looking ahead, hybrid approaches that blend machine learning with physics-based models offer a promising path. This work takes a step in that direction, aiming to improve predictive maintenance by integrating both methodologies. The goal is to develop robust, adaptive systems that can forecast insulation degradation under varying conditions, even with limited data, while reducing reliance on extensive empirical testing.

I. PHYSICS-BASED DEEP LEARNING

Physics-based deep learning represents a transformative approach in predicting insulation degradation by integrating traditional physical models with advanced machine learning techniques. This synergy combines the interpretability and reliability inherent in physics-driven approaches with the adaptive capabilities of modern neural networks, offering enhanced predictive accuracy crucial for applications like electrical insulating materials.

At its core, this methodology ensures that predictions are not only statistically plausible but also physically viable. A key component is the integration of established physical laws into neural networks, enhancing their ability to make reliable forecasts. Physics-informed neural networks (PINNs) exemplify this approach by embedding differential equations governing physical phenomena directly into the training process, showing promise in predicting material behavior under various stressors such as electrical fields and temperature variations[3]-[5].

The foundation of any effective ML application is high-quality data. In the realm of electrical engineering, data often comes from sensors and historical records, which can be noisy or incomplete. This noise and the presence of outliers in complex environments pose a significant challenge. Engineers must implement rigorous data collection protocols to minimize these issues. Moreover, the diversity of data types—from time-series sensor data to structured logs—necessitates meticulous preprocessing efforts. This includes cleaning, normalization, and feature engineering, all of which demand substantial resources and domain expertise. The goal is to reproduce synthetic data already into a format that accurately captures the relevant features for model training, effectively leveraging physics to bootstrap the training phase of the ML model.

II. SYNTHETIC DATASET GENERATION

Authentic data regarding the residual life time of dielectric substrates of power modules is a commodity hard to access from OEMs. Hence, this work developed a method to reproduce a high number of synthetic data points, to be used for the training of machine learning models. The following will elucidate the methodology employed.

The residual life of a dielectric is often described in literature by physical models, usually considering thermal (T) and electrical (E) stresses as main players in their aging process. The expected life (L) can be estimated using the following equation, involving various material-related parameters:

$$L = \frac{1}{R} = \frac{L_0}{\left(\frac{E}{E_0}\right)^n \exp(B \cdot cT)} \quad (1)$$

Where:

- R is the reaction rate,
- E_0 is a reference electrical stress,
- T_0 is a reference temperature,
- L_0 , the reference life, i.e., the life of the insulation system subjected to E_0 and T_0 ,
- cT is the conventional thermal stress $cT = 1/T_0 - 1/T$,

- n is a model parameter indicated as voltage endurance coefficient (VEC), estimated by values found in literature
- B is the ratio of the activation energy of the dominant chemical reaction leading to degradation to the Boltzmann constant.

As stresses T, E are random variables, it is necessary to address the impact of their stochastic nature in the proposed physical model replacing the reaction rate R with its expected value $\langle R \rangle$. A Monte Carlo method was applied here, randomly sampling these parameters to construct a joint probability distribution function reflecting the workload on individual modules.

$$\langle R \rangle = R_0 \iint \left(\frac{E}{E_0}\right)^n \exp(B \cdot cT) f\left(\frac{E}{E_0}, cT\right) dE dcT \quad (2)$$

Where:

- $R_0 = 1/L_0$ and
- $f(E/E_0, cT)$ is the joint probability density of the normalized electric field and the conventional thermal stress.

Once the distributions characterizing stresses are known, the expected life can be derived from Equation (2). As a result, an estimation of the residual life, RL , can also be derived, as shown in the schematic representation in Figure 1.

When this process is iterated over time, at different time steps, a time series of residual life over time can be obtained, as shown in Figure 2. As can be seen, the prediction over time features higher variability at the beginning, when the predictors related to the distributions of variables characterizing the insulation life are more uncertain. However, after a relatively low operational time of about 500 hours, the prediction evolves with a lower variance over time.

This process can be repeated across multiple devices and platforms (e.g., electric vehicles, or battery chargers) achieving a plurality of distributions of temperature and electrical stress, yielding a large and comprehensive dataset of estimated or residual lifetimes, based on a preliminary set of model parameters (e.g. derived from literature or accelerated testing).




Figure 1 - Schematic representation of the proposed process for estimation of residual life over time.

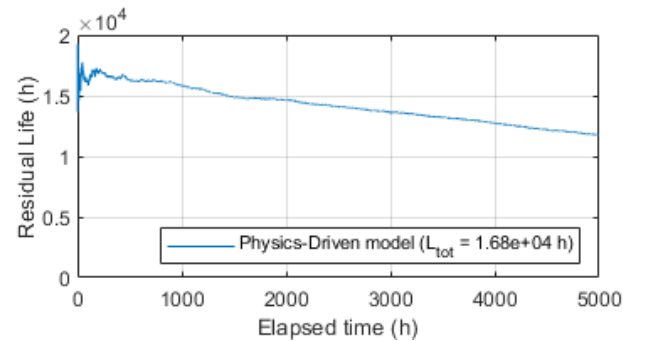


Figure 2 – Trend of residual life over time, after the statistical-physical model proposed.

Additionally, after a suitable number of failures have been collected, a more general life model based on e.g. CNNs can be trained and re-deployed to achieve more accurate estimates. The spirit of Equation (2) can also be further extended, including the variability of *all* the parameters in Equation (1), which are usually considered constants for the sake of simplicity. In reality those depend on a number of factors related to the nature of the dielectric material and its production process. In order to try and better replicate a realistic scenario, the final form of the model used in this work for the production of a synthetic dataset will be considering also this aspect, resorting to a Monte Carlo approach that more easily introduces more complex joint probability density functions into simulations:

$$f(E, E_0, T, T_0, n, B) \quad (3)$$

A schematic representation of the process is shown in Figure 3. An example of the statistical distributions of temperature, electric field and life for a pool of 1000 devices estimated by this algorithm is provided by histograms in Figure 4. In the next chapters we will see if the produced set of residual life predictions can be used to successfully train a convolutional neural network (CNN) for this kind of continuous life prediction.

III. DATASET PROCESSING

CNNs are a type of deep learning model particularly well-suited for analyzing visual data. Inspired by the structure of the human visual cortex, CNNs are designed to automatically and adaptively learn spatial hierarchies of features through a series of convolutional layers. These networks take as input multi-dimensional arrays.

The convolutional layers act as filters that slide over the input data, detecting patterns such as edges, textures, or more complex shapes in deeper layers. These learned patterns help the network extract meaningful features from raw input, reducing the need for manual feature engineering.

As the data passes through the network, it's progressively transformed into more abstract and informative representations. Ultimately, these representations are passed to fully connected layers, which function similarly to traditional neural networks and are used to make predictions.

The previous chapter covered a way of producing a number of synthetic time series of length M for the temperatures, electric fields and residual life, relative to the working conditions and life of a random device. Those are arrays that will be respectively named $[T]$, $[E]$ and $[RL]$. The arrays $[T]$, $[E]$ can be seen as a series of real measurements performed by, for example, the management system of the device itself.

From those, it is possible to estimate the metadata relevant to the statistical distributions those variables belong to, typically in the form of a cumulative average.

The same applies to $[E]$. The multidimensional array composed of $\langle [T] \rangle$, $\langle var[T] \rangle$, $\langle [E] \rangle$, $\langle var[E] \rangle$ was chosen as the input to be fed to the CNN, which prediction should reveal coincident with the synthetic result $[RL]$.

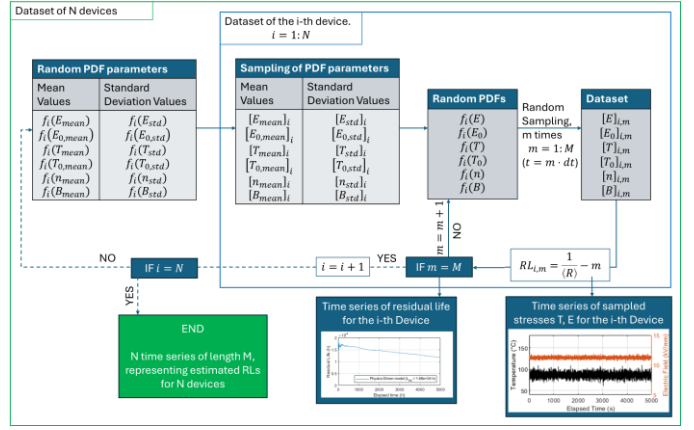


Figure 3 - Block diagram illustrating the key steps in the residual life prediction algorithm.

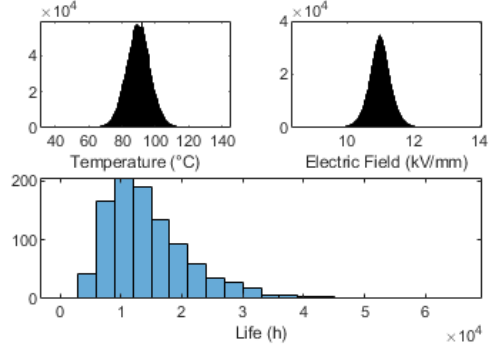


Figure 4 - Histograms showing the statistical distributions of temperature (top left), electric field (top right), and insulation life for a population of 1000 devices.

For example, the expected value and variance of $[T]$ after k observations can be obtained from:

$$\langle [T] \rangle_k = \frac{\sum_{j=1}^k [T]_j}{k} \quad (4)$$

$$\langle var[T] \rangle_k = \frac{\sum_{j=1}^k ([T]_j - \langle [T] \rangle_k)^2}{k - 1} \quad (5)$$

The model learns to make accurate predictions by adjusting its parameters, called weights, to optimize its prediction to the correct value to be forecast. The difference between the predicted output and the true value is measured using a loss function, which gives a numerical value representing how wrong the prediction is. This loss is used to guide the adjustment of the weights throughout the network, eventually allowing the CNN to gradually learn which features are important and how to combine them to make better predictions. In this work, training is done on a subset of 80% of the devices simulated with the technique described in the previous chapter, providing the training session with the multidimensional array input described above, and the relevant outputs $[RL]$. Testing is done on the remaining 20%.

IV. RESULTS AND DISCUSSION

A dataset for 1000 devices was produced by means described in Chapter III, simulating an observation period of 5000 hours, with 1 sampling cycle per hour. As mentioned, results for 800 devices were used for training, while 200 were used for testing the resulting CNN.

An overview of results of life predictions after 5000h is shown in Figure 5, while a more detailed plot of RL over time for one random testing device can be found in Figure 6.

A satisfactory result was reached, as predictions generally are contained within a $\pm 25\%$ error from values predicted by the physical model (dashed lines) in Figure 5.

Figure 6 shows a residual life prediction over time with higher variability at the beginning, stabilizing with a lower variance, as already observed for predictions from the physics driven model. It can also be noticed that the discrepancy of prediction is an overall offset of the plot from what should constitute the real values, rather than a random oscillation around them. This characteristic is observed for all results, possibly related to the regularity of the random distributions used to sample values of T, E over time, eventually resembling a pseudo-gaussian distribution.

While the CNN demonstrated promising accuracy when trained only on physics-informed synthetic data, an open challenge remains in applying the model to real-world scenarios, where operating conditions and stress profiles might differ from the simulated ones. However, since the data generation process already embeds the main physical degradation mechanisms, we believe that the model could represent a good starting point for real applications. In this perspective, only a limited amount of real-world data might be required to fine-tune the model and correct possible deviations, reducing the effort usually needed for extensive data collection in the field.

V. CONCLUSIONS

This study successfully demonstrated the feasibility of integrating physics-informed models with machine learning techniques to predict the residual life of transformer insulation. By leveraging PDE-based simulations and Monte Carlo methods, we generated synthetic data that allowed us to train a CNN effectively. The results showed promising accuracy, with predictions within $\pm 20\%$ error margin, indicating the model's capability to generalize well from the training data.

Some limitations were identified, as the initial higher variability in predictions suggests that the model may require further refinement or additional features to capture the complexities of insulation degradation more accurately. Additionally, the observed systematic offset in predictions highlights potential biases in the synthetic data generation process or simplifications made in the physical model.

Several avenues for improvement present themselves, as incorporating more realistic simulation parameters and environmental factors into the PDE model could better reflect real-world conditions, potentially reducing the observed prediction offset.

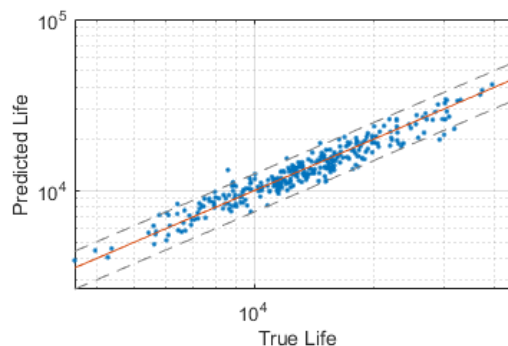


Figure 5 – Visualization of residual life predictions obtained after 5000h from the CNN (Predicted Life) VS from the physics-driven model (True Life). Grey dashed lines delimit a $\pm 25\%$ prediction error.

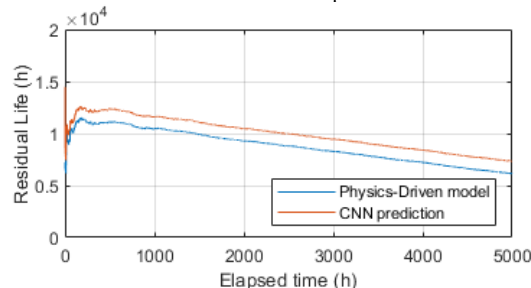


Figure 6 - Time-dependent residual life prediction with stabilized variance over time, demonstrating the offset discrepancy between predicted and actual values.

REFERENCES

- [1] B. Huang and J. Wang, "Applications of Physics-Informed Neural Networks in Power Systems - A Review," in *IEEE Transactions on Power Systems*, vol. 38, no. 1, pp. 572-588, Jan. 2023, doi: 10.1109/TPWRS.2022.3162473
- [2] S. Baltazar, C. Li, H. Daniel and J. Valente de Oliveira, "A Review on Neurocomputing Based wind Turbines Fault Diagnosis and Prognosis," 2018 Prognostics and System Health Management Conference (PHM-Chongqing), Chongqing, 2018, pp. 437-443, doi: 10.1109/PHM-Chongqing.2018.00081.
- [3] F. Bragone, K. Oueslati, T. Laneryd, M. Luvisotto and K. Morozovska, "Physics-Informed Neural Networks for Modeling Cellulose Degradation in Power Transformers," 2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA), Nassau, Bahamas, 2022, pp. 1365-1372, doi: 10.1109/ICMLA55696.2022.00216
- [4] K. Oueslati, N. Dhahbi-Megriche, F. Bragone, K. Morozovska, T. Laneryd and M. Luvisotto, "Physics-Informed Neural Networks for modelling insulation paper degradation in Power Transformers," 2022 IEEE International Conference on Electrical Sciences and Technologies in Maghreb (CISTEM), Tunis, Tunisia, 2022, pp. 1-6, doi: 10.1109/CISTEM55808.2022.10043884
- [5] Zhang, G., Li, Y., Jiang, W., & Shu, L. (2022). A fault diagnosis method for wind turbines with limited labeled data based on balanced joint adaptive network. *Neurocomputing*, 481, 133–153. doi: 10.1016/j.neucom.2022.01.067