International Journal of Digital Twin Systems and Computing (IJDTSC) ISSN: 3108–0790
DOI: https://doi.org/10.XXXXX/IJDTSC.2026.V2I301

Digital Twin-Driven Real-Time EMF Exposure Monitoring for Urban Distribution Substations

[1]Nuramalina binti Bohari
[1]Lecturer, Department of Electrical Engineering, University of Technology Sarawak, Sibu, Malaysia
Email id: nuramalina@uts.edu.my
ORCID: 0009-0000-0803-1141

Abstract— ***Distribution substations installed within dense commercial and residential precincts expose nearby people to extremely low frequency electromagnetic fields, yet compliance is usually verified through one-off field surveys that cannot capture time-varying load or network growth. This paper proposes a digital twin framework for continuous, spatially resolved estimation of electric and magnetic fields around urban distribution substations. The framework couples a physics-based field model with residual learning and Gaussian process interpolation, while a Kalman filter synchronisation engine with drift-triggered retraining keeps the virtual replica consistent with the physical substation. The twin generates live exposure maps at unmeasured locations, forecasts field levels under projected load growth, and issues uncertainty-aware early warnings against the reference levels of the International Commission on Non-Ionizing Radiation Protection. In a measurement-informed simulation study of a compact substation, the hybrid twin reduced the mean absolute error and root-mean-square error of magnetic flux density estimation by 56.4 percent and 65.9 percent relative to a physics-only model, and raised the coefficient of determination from 0.937 to 0.993. The maximum exposure ratio across the mapped area was 0.101, and the site remained compliant up to an 895 percent increase in operating current, equivalent to 47.1 years at 5 percent annual growth. Drift-triggered retraining reduced twin error by 46.7 percent compared with a non-adaptive twin after 72 hours. The approach turns static compliance surveys into continuous, predictive exposure management for utilities and urban planners.***

Keywords— Digital Twin; Electromagnetic Fields; Distribution Substation; Real-Time Exposure Monitoring; ICNIRP Compliance.

I. Introduction

The digital twin (DT) was originally conceived as a virtual counterpart of a physical product that mirrors its behaviour throughout the life cycle and helps anticipate undesirable emergent behaviour in complex systems [1]. Since then, the concept has matured into an established industrial paradigm that links physical assets, virtual models, and bidirectional data flows [2, 3]. Formally, a digital twin at time t can be described as the tuple

$$\mathcal{DT}(t) = \left\{ \mathcal{P}(t),\ \mathcal{V}(t;\boldsymbol{\theta}),\ \mathcal{D}(t),\ \mathcal{S}(t) \right\}, \tag{1}$$

where $\mathcal{P}$ is the physical entity, $\mathcal{V}$ is the virtual model with parameters $\boldsymbol{\theta}$, $\mathcal{D}$ is the stream of sensed and operational data, and $\mathcal{S}$ is the synchronisation service that keeps $\mathcal{V}$ consistent with $\mathcal{P}$ [2, 4]. From a modelling perspective, the value of a DT depends on how faithfully $\mathcal{V}$ tracks $\mathcal{P}$ as operating conditions evolve [4].

In power systems, DTs have already been applied to online grid analysis, where a continuously updated virtual grid supports dispatch decisions in real time [5], and to fault diagnosis in distributed photovoltaic systems, where deviations between twin estimates and measurements reveal faults [6]. Distribution substations, however, raise a public-health question that these works do not address: substations installed inside dense commercial and residential precincts expose nearby people to extremely low-frequency (ELF) electric and magnetic fields. At 50 Hz, the magnetic flux density produced by the internal conductors carrying currents $I_k(t)$ follows the Biot–Savart law,

$$\mathbf{B}(\mathbf{r},t) = \frac{\mu_0}{4\pi}\sum_{k=1}^{K} I_k(t) \oint_{C_k} \frac{d\mathbf{l}_k \times (\mathbf{r}-\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|^3}, \tag{2}$$

where μ_0 is the permeability of free space and C_k is the path of the k-th conductor. At a radial distance d from the enclosure, the field magnitude is commonly approximated by a power-law decay,

$$B(d,t) \approx B_0(t)\left(\frac{d_0}{d}\right)^n, \qquad n \in \{1,2,3\}, \tag{3}$$

where $n = 1$, 2, and 3 correspond to a single conductor, a balanced conductor pair, and a dipole-like transformer source, respectively. Similarly, the electric field outside a grounded or metal-enclosed unit can be written as

$$E(d) \approx (1-\eta_s)\, E_0 \left(\frac{d_0}{d}\right)^m, \tag{4}$$

where $\eta_s \in [0,1]$ is an effective shielding coefficient that accounts for enclosures, boundary walls, and vegetation. Because $B_0(t)$ scales directly with load current, magnetic exposure varies with demand, and a single measurement survey reflects only the load level at the moment it was taken.

Exposure is regulated through the reference levels of the International Commission on Non-Ionizing Radiation Protection (ICNIRP), which at 50 Hz are $E_{\text{ref}} = 5000$ V/m and $B_{\text{ref}} = 100\,\mu$T for the general public [9]. A compact compliance indicator is the exposure ratio

$$\text{ER}(\mathbf{r},t) = \max\left\{\frac{E(\mathbf{r},t)}{E_{\text{ref}}},\ \frac{B(\mathbf{r},t)}{B_{\text{ref}}}\right\} \le 1, \tag{5}$$

which must hold at every occupied location $\mathbf{r}$ and at every time t. Conventional compliance checks evaluate (5) only at a handful of measured points and at a single instant. They cannot estimate exposure at unmeasured locations, quantify the uncertainty of their estimates, or anticipate how exposure will change as load grows.

Purely physics-based models such as (2)–(4) are interpretable, but they rely on idealised geometry and uncertain parameters. Purely data-driven models, by contrast, need large datasets and may violate physical laws. Physics-informed machine learning bridges this gap by embedding governing equations into learning [7, 8]. A widely adopted hybrid structure combines a physics prior with a learned residual:

$$\hat{y}(\mathbf{x},t) = f_{\text{phys}}(\mathbf{x},t;\boldsymbol{\theta}_p) + g_{\text{NN}}(\mathbf{x},t;\boldsymbol{\theta}_d), \tag{6}$$

where f_{phys} is the analytical field model and g_{NN} is a neural network that learns the discrepancy between the model and reality. Hybrid twins of this kind have been shown to reduce prediction error in tool-wear prediction with self-evolving models [10], industrial predictive monitoring using PINNs [12], thermal estimation of permanent magnet machines [13], vortex-induced vibration of deep-sea risers through TCN residual compensation [14], and ship docking prediction under limited data [15]. A recent review likewise reports that hybrid twins outperform single-strategy approaches in predictive maintenance [11]. The network parameters are typically learned by minimising a composite loss

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}_d) = &\underbrace{\frac{1}{N}\sum_{i=1}^{N}\left\|\hat{y}_i - y_i\right\|^2}_{\mathcal{L}_{\text{data}}} + \lambda_r\|\boldsymbol{\theta}_d\|_2^2 \\ &+ \lambda_p \underbrace{\frac{1}{N_c}\sum_{j=1}^{N_c}\left|\nabla\cdot\hat{\mathbf{B}}(\mathbf{r}_j)\right|^2}_{\mathcal{L}_{\text{phys}}},\end{aligned} \tag{7}$$

where $\mathcal{L}_{\text{phys}}$ enforces Gauss's law for magnetism ($\nabla\cdot\mathbf{B} = 0$) at N_c collocation points, and λ_p and λ_r are weighting coefficients [7, 8].

For the twin to remain faithful, its virtual state must be corrected continuously as new measurements $\mathbf{z}_k$ arrive. A recursive estimator provides this synchronisation:

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k\big(\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1}\big), \tag{8}$$

where $\mathbf{K}_k$ is the gain matrix and $\mathbf{H}$ maps the field state to the sensor locations. Twin fidelity over a sliding window of W samples can then be quantified as

$$\varepsilon_{\text{DT}} = \sqrt{\frac{1}{W}\sum_{k=1}^{W}\left\|\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k}\right\|^2}. \tag{9}$$

Finally, because $B \propto I$, the maximum load growth that a site can tolerate before magnetic exposure reaches the limit follows directly from the twin's peak estimate $\hat{B}_{\max}$:

$$\gamma_{\max} = \frac{B_{\text{ref}}}{\hat{B}_{\max}} - 1, \tag{10}$$

which gives utilities and planners a quantitative safety margin for future network expansion.

Despite this progress, existing DT studies focus on grid operation [5], fault diagnosis [6], and machine health [10, 12, 13, 11]. To the best of the authors' knowledge, none of them uses a digital twin to estimate EMF exposure, check regulatory compliance continuously, or forecast exposure under load growth around distribution substations. This paper addresses that gap. Its four key objectives are:

1. To develop a hybrid physics–data digital twin that combines the analytical field models (2)–(4) with a physics-regularised neural correction (6)–(7) and Gaussian process spatial interpolation, for spatially resolved estimation of E and B around distribution substations.

2. To design a real-time synchronisation mechanism (8) that continuously updates the virtual replica from IoT sensor data and live load current, and to maintain the fidelity measured by (9) through drift-triggered retraining.

3. To implement continuous, uncertainty-aware ICNIRP compliance monitoring through the exposure ratio (5), with early warnings when predicted exposure approaches the reference levels [9].

4. To forecast exposure under projected load growth and quantify the safe expansion margin $\gamma_{\max}$ in (10), supporting the planning of compact and standard substations in urban areas.

In pursuing these objectives, the paper makes three main contributions. First, it applies the hybrid digital twin paradigm, so far used mainly for machine health, to public EMF exposure management around distribution substations. Second, it combines physics-based modelling, residual learning, Kalman-filter synchronisation, and Gaussian process uncertainty quantification in a single framework, formalised in two algorithms for offline training and online operation. Third, it turns the twin's outputs into planning indicators, namely the exposure ratio, the load-growth margin, and the time to reach each alert threshold, which utilities can act on directly. Table 1 compares the proposed framework with representative digital twin and hybrid-modelling studies.

As Table 1 shows, each earlier work covers at most two of the six capabilities, and none addresses EMF exposure, compliance warning, or load-growth forecasting. To the best of the authors' knowledge, the proposed framework is the first to combine all six in a single digital twin for public-exposure management around distribution substations.

The remainder of this paper is organised as follows. Section II describes the proposed digital twin architecture, the hybrid field model, and the system parameters. Section III presents the methodology, including physics parameter estimation, residual learning, Bayesian hyperparameter optimisation, and real-time synchronisation, formalised in two algorithms. Section IV reports the results of a measurement-informed simulation study covering waveform analysis, spatial decay estimation, exposure mapping, load-growth forecasting, and drift-aware synchronisation. Section V concludes the paper and outlines directions for future work.

II. System Description and Proposed Digital Twin Architecture

A. Overall Architecture

The proposed digital twin (DT) follows the physical–virtual–data–service structure formalised in (1) [2, 3]. It is organised into the five layers shown in Fig. 1. The *Physical Layer* contains the distribution substation, the transformer, the busbars, and the surrounding buildings and pedestrians who make up the exposed population. The *Sensing Layer* captures the field and operating variables through magnetic field sensors, electric field probes, current transducers, and temperature and humidity sensors. The *Communication and Edge Layer* carries these measurements through an IoT gateway and wireless link to an edge computing node, which performs preprocessing close to the source. The *Virtual Model Layer* hosts the hybrid model, which combines a physics-based field model with a data-driven correction as expressed in (6). The *Service Layer* delivers four outputs: a live EMF exposure heat map, an ICNIRP compliance dashboard, early-warning alerts, and load-growth forecast charts.

Data flows upward from the physical assets to the services. Two feedback paths close the loop. The *synchronisation loop* returns service-level observations to the virtual model, where the recursive correction (8) and drift-triggered retraining keep its state and parameters consistent with the physical system. The *control and decision support* path returns the twin's outputs to the physical layer, where they inform operator actions such as load transfer, shielding upgrades, or restricted access zones [5, 4].

B. Physical Layer and Measurement Configuration

The monitored assets are compact distribution substations housed in sealed metal enclosures and standard distribution substations with open-panel configurations. Sensors are placed at radial distances $d \in \mathcal{D}_r = \{0.5, 1, 2, 3, 4, 5\}\,\mathrm{m}$ in four orthogonal directions $\phi \in \Phi = \{\mathrm{F}, \mathrm{B}, \mathrm{L}, \mathrm{R}\}$ (front, back, left, right) relative to the enclosure door. The measurement grid is therefore

$$\mathcal{G} = \big\{(d, \phi) \;\big|\; d \in \mathcal{D}_r,\ \phi \in \Phi_{\mathrm{acc}} \subseteq \Phi\big\}, \tag{11}$$

where Φ_{acc} is the set of accessible faces, giving at most $|\mathcal{G}| = 6 \times 4 = 24$ measurement points per substation. Faces blocked by drainage trenches or structures are excluded from $\mathcal{G}$, and field values there are estimated by the spatial interpolation described in Section II-E.

C. Sensing, Data Acquisition, and Preprocessing

Each tri-axial magnetic sensor samples the instantaneous flux density components $b_x(t)$, $b_y(t)$, and $b_z(t)$. The root-mean-square (RMS) value of each component over one or more fundamental periods $T = 1/f$ at $f = 50\,\mathrm{Hz}$ is

$$B_{i,\mathrm{rms}} = \sqrt{\frac{1}{T}\int_{t_0}^{t_0+T} b_i^2(t)\,dt}, \qquad i \in \{x, y, z\}, \tag{12}$$

and the resultant field magnitude reported for exposure assessment is

$$B_{\mathrm{res}} = \sqrt{B_{x,\mathrm{rms}}^2 + B_{y,\mathrm{rms}}^2 + B_{z,\mathrm{rms}}^2}. \tag{13}$$

The electric field magnitude E_{res} is obtained in the same way from the probe outputs. At the edge node, the 50 Hz fundamental is extracted to suppress harmonic and broadband noise, all channels are time-stamped against a common clock, and each input variable is min–max normalised before it reaches the virtual model.

D. Hybrid Virtual Field Model

The input feature vector at location (d, ϕ) and time t is

$$\mathbf{x}(t) = \big[\, d,\ \phi,\ I_L(t),\ T_a(t),\ H_r(t),\ \eta_s \,\big]^{\top}, \tag{14}$$

where I_L is the transformer load current, T_a is the ambient temperature, H_r is the relative humidity, and η_s is the shielding coefficient introduced in (4).

The physics layer scales a reference field profile with the live load current, since magnetic flux density is proportional to current through (2):

$$f_{\mathrm{phys}}(\mathbf{x}, t) = B_0^{(\phi)}\,\frac{I_L(t)}{I_{\mathrm{ref}}}\left(\frac{d_0}{d}\right)^{n_\phi}, \tag{15}$$

Table 1: Comparison of the proposed framework with existing digital twin and hybrid-modelling studies.

Ref.	Application	Hybrid physics–data	Online sync./update	Power asset	EMF exposure	Compliance warning	Load-growth forecast
[2]	DT in industry (review)	×	✓	×	×	×	×
[3]	DT technologies (review)	×	✓	×	×	×	×
[5]	Power grid online analysis	×	✓	✓	×	×	×
[6]	PV fault diagnosis	×	✓	✓	×	×	×
[10]	CNC tool wear prediction	✓	✓	×	×	×	×
[12]	Hydraulic system monitoring	✓	✓	×	×	×	×
[13]	PMSM temperature estimation	✓	✓	×	×	×	×
[14]	Riser vibration prediction	✓	×	×	×	×	×
[15]	Ship docking prediction	✓	×	×	×	×	×
[11]	Predictive maintenance (review)	✓	✓	×	×	×	×
Proposed	**Substation EMF exposure**	✓	✓	✓	✓	✓	✓

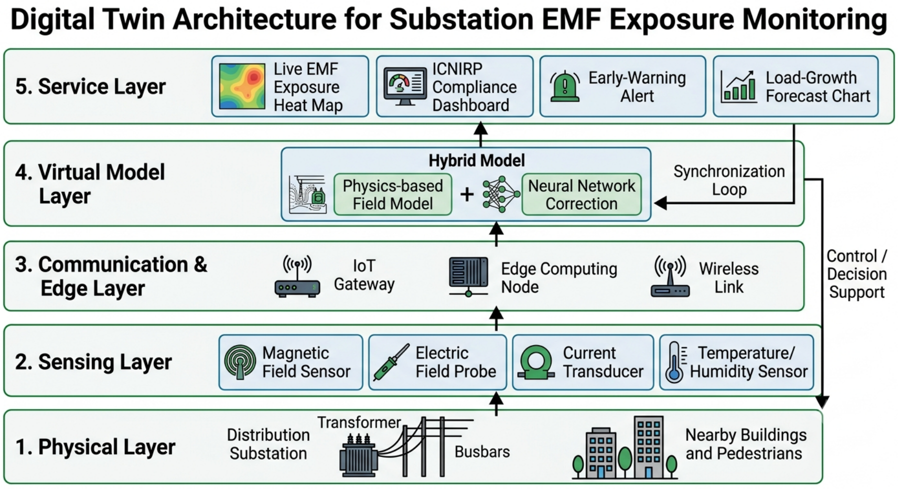

Figure 1: Five-layer digital twin architecture for substation EMF exposure monitoring, showing the upward data flow, the synchronisation loop, and the control and decision support path.

where $B_0^{(\phi)}$ and n_ϕ are the face-specific reference amplitude and decay exponent, both fitted to baseline measurements, and I_{ref} is the load current at which the baseline was recorded. A separate exponent for each face captures the asymmetric field distributions caused by the internal busbar and winding layout.

The data-driven correction g_{NN} learns the residual between the physics prediction and the measurement [14, 15]. For time-varying operation, it is implemented as a lightweight temporal convolutional network (TCN) [14, 10]. For a dilated causal convolution with kernel size K and dilation factor δ, the hidden response at layer ℓ is

$$\mathbf{h}^{(\ell)}(t) = \sigma\left(\sum_{k=0}^{K-1} \mathbf{W}_k^{(\ell)}\, \mathbf{h}^{(\ell-1)}(t-\delta k) + \mathbf{b}^{(\ell)}\right), \quad (16)$$

where $\sigma(\cdot)$ is the activation function and $\mathbf{W}_k^{(\ell)}$ and $\mathbf{b}^{(\ell)}$ are trainable weights and biases. The network is trained with the physics-informed loss (7), which keeps the corrected field consistent with Gauss's law for magnetism [7, 8]. When only a static spatial profile is available, as in the baseline evaluation of Section IV, the temporal inputs are constant and the correction reduces to a spatial residual, which is estimated by the Gaussian process described next.

E. Spatial Interpolation for Unmeasured Locations

To produce continuous exposure maps, the twin estimates the field at any unmeasured point $\mathbf{r}_*$ by adding a Gaussian process estimate of the physics residual to the physics prediction. With $\mathbf{r} = [r_1, \ldots, r_{N_s}]^\top$ the residuals at the N_s sensor locations, the posterior mean of the field is

$$\hat{B}(\mathbf{r}_*) = f_{\text{phys}}(\mathbf{r}_*) + \mathbf{k}_*^\top \left(\mathbf{K} + \sigma_n^2 \mathbf{I}\right)^{-1} \mathbf{r}, \quad (17)$$

where $\mathbf{K}$ is the $N_s \times N_s$ covariance matrix between sensor locations, $\mathbf{k}_*$ is the covariance vector between $\mathbf{r}_*$ and the sensors, and σ_n^2 is the measurement noise variance. A squared-exponential kernel is adopted:

$$k(\mathbf{r}_p, \mathbf{r}_q) = \sigma_f^2 \exp\left(-\frac{\|\mathbf{r}_p - \mathbf{r}_q\|^2}{2\ell_s^2}\right), \quad (18)$$

where σ_f^2 is the signal variance and ℓ_s is the spatial length scale. Because the Gaussian process models only the residual, the physics term keeps the estimate physically plausible between sensors, while the posterior variance gives a confidence bound for every point on the heat map.

F. Synchronisation and Model Update

The virtual model is updated periodically through the recursive correction (8), and also whenever model drift is detected. Drift is measured as the relative change in twin fidelity (9) with respect to the value ε_0 recorded after the last training:

$$\Delta_{\text{drift}}(t) = \frac{\varepsilon_{\text{DT}}(t) - \varepsilon_0}{\varepsilon_0} > \delta_{\text{th}}, \tag{19}$$

where δ_{th} is the drift threshold. When (19) holds, the data-driven correction is retrained on the most recent data window, in line with the self-evolution strategies used in hybrid twins [10, 12]. For the twin to operate in real time, the end-to-end latency must satisfy

$$\tau_{\text{total}} = \tau_{\text{s}} + \tau_{\text{c}} + \tau_{\text{e}} + \tau_{\text{m}} \le \tau_{\max}, \tag{20}$$

where τ_{s}, τ_{c}, τ_{e}, and τ_{m} are the sensing, communication, edge-processing, and model-inference delays, and $\tau_{\max}$ is the required update interval.

G. Service Layer: Compliance Monitoring and Forecasting

The compliance dashboard classifies each location into one of three alert states based on the exposure ratio (5):

$$\mathcal{A}(\mathbf{r}, t) = \begin{cases} \text{Normal}, & \text{ER}(\mathbf{r}, t) < \alpha_1, \\ \text{Caution}, & \alpha_1 \le \text{ER}(\mathbf{r}, t) < \alpha_2, \\ \text{Alarm}, & \text{ER}(\mathbf{r}, t) \ge \alpha_2, \end{cases} \tag{21}$$

where α_1 and α_2 are the caution and alarm thresholds, expressed as fractions of the ICNIRP reference levels [9]. For load-growth forecasting, the load current is projected over a horizon of h years at an annual growth rate g:

$$\hat{I}_L(t + h) = I_L(t)\,(1 + g)^h, \tag{22}$$

and the corresponding peak field forecast follows from (15) as $\hat{B}_{\max}(t + h) = \hat{B}_{\max}(t)\,(1 + g)^h$. Setting this forecast equal to B_{ref} gives the number of years before the limit is reached:

$$h_{\text{lim}} = \frac{\ln\big(B_{\text{ref}}/\hat{B}_{\max}(t)\big)}{\ln(1 + g)}, \tag{23}$$

which complements the load-growth margin $\gamma_{\max}$ in (10) as a planning indicator.

H. System Parameters and Ratings

The rated secondary current of the transformer follows from its apparent power and line voltage:

$$I_{\text{FL}} = \frac{S}{\sqrt{3}\,V_{\text{LL}}} = \frac{500 \times 10^3}{\sqrt{3} \times 415} = 695.6\ \text{A}, \tag{24}$$

so a current transducer range of 0–1000 A covers full load with a 44 % overload margin. The key parameters of the physical system, the sensing hardware, and the digital twin configuration used in the evaluation of Section IV are listed in Table 2.

Table 2: Key system parameters of the proposed digital twin.

Parameter	Value
Physical system and sensing	
Frequency f / sampling rate f_s	50 Hz / 10 kHz
Transformer rating / voltage	500 kVA, 11/0.415 kV
Rated secondary current I_{FL}	695.6 A
B-sensor / E-probe range	0.01–2000 µT / 1 V/m–100 kV/m
Distances $\mathcal{D}_r$ / ref. d_0	0.5–5 m (6 points) / 0.5 m
Digital twin model	
Update interval $\tau_{\max}$	1 s
Kalman noise Q / R	0.02 / 0.1225 µT^2
TCN kernel K / layers	3 / 4
Loss weights λ_p / λ_r	0.1 / 10^{-4}
GP σ_f / ℓ_s / σ_n	0.3 µT / 1.0 m / 0.05 µT
Baseline fidelity ε_0 / δ_{th}	0.20 µT / 30 %
Compliance and forecasting	
ICNIRP E_{ref} / B_{ref}	5000 V/m / 100 µT [9]
Alert thresholds α_1 / α_2	0.5 / 0.8
Load-growth rates g	3, 5, 8 %

III. Methodology

The methodology proceeds in two stages. In the offline stage, the hybrid field model is constructed and trained from baseline measurements (Algorithm 1). In the online stage, the trained twin is synchronised with live sensor data and used for continuous compliance monitoring and forecasting (Algorithm 2). Both stages build on the architecture described in Section II [4, 3].

A. Physics Model Parameter Estimation

The face-specific parameters $B_0^{(\phi)}$ and n_ϕ of the physics model (15) are estimated from the baseline profile of each face. Taking the logarithm of the power-law decay (3) yields a model that is linear in the unknowns:

$$\ln B(d_j) = \ln B_0^{(\phi)} + n_\phi \ln\left(\frac{d_0}{d_j}\right), \qquad j = 1, \ldots, J. \tag{25}$$

Stacking the J measured distances in matrix form, $\mathbf{y}_\phi = \mathbf{A}\boldsymbol{\beta}_\phi$ with $\boldsymbol{\beta}_\phi = [\ln B_0^{(\phi)},\, n_\phi]^\top$, the least-squares estimate is

$$\hat{\boldsymbol{\beta}}_\phi = \left(\mathbf{A}^\top \mathbf{A}\right)^{-1} \mathbf{A}^\top \mathbf{y}_\phi, \tag{26}$$

where the j-th row of $\mathbf{A}$ is $[\,1,\ \ln(d_0/d_j)\,]$. The electric field parameters E_0, m, and η_s in (4) are estimated in the same way for each face.

B. Residual Learning and Network Training

Once the physics model is fitted, the training target for the neural correction is the residual between each measurement and its physics prediction:

$$r_i = y_i - f_{\text{phys}}(\mathbf{x}_i; \hat{\boldsymbol{\theta}}_p), \qquad i = 1, \ldots, N. \tag{27}$$

Learning residuals rather than raw fields lets the network focus on what the physics model misses, such as conductor asymmetry and local shielding effects [14, 15]. The dataset $\mathcal{S} = \{(\mathbf{x}_i, r_i)\}_{i=1}^{N}$ is partitioned chronologically into disjoint subsets, which prevents future data from leaking into training:

$$\mathcal{S} = \mathcal{S}_{\text{tr}} \cup \mathcal{S}_{\text{val}} \cup \mathcal{S}_{\text{te}}. \tag{28}$$

The TCN parameters $\boldsymbol{\theta}_d$ are optimised by minimising the composite loss (7) with the Adam optimiser. At iteration s, with gradient $\mathbf{g}_s = \nabla_{\boldsymbol{\theta}_d}\mathcal{L}$,

$$\mathbf{m}_s = \beta_1 \mathbf{m}_{s-1} + (1-\beta_1)\,\mathbf{g}_s, \tag{29}$$

$$\mathbf{v}_s = \beta_2 \mathbf{v}_{s-1} + (1-\beta_2)\,\mathbf{g}_s^2, \tag{30}$$

$$\boldsymbol{\theta}_{d,s} = \boldsymbol{\theta}_{d,s-1} - \eta\,\frac{\hat{\mathbf{m}}_s}{\sqrt{\hat{\mathbf{v}}_s} + \epsilon}, \tag{31}$$

where $\hat{\mathbf{m}}_s = \mathbf{m}_s/(1-\beta_1^s)$ and $\hat{\mathbf{v}}_s = \mathbf{v}_s/(1-\beta_2^s)$ are the bias-corrected moment estimates, η is the learning rate, and ϵ is a small constant for numerical stability. Training uses $\eta = 10^{-3}$ and stops early when the validation loss has not improved for $P = 20$ consecutive epochs:

$$\mathcal{L}_{\text{val}}(e) \geq \min_{e' < e} \mathcal{L}_{\text{val}}(e'), \qquad e = e^* + 1, \ldots, e^* + P. \tag{32}$$

C. Bayesian Hyperparameter Optimisation

The loss weights λ_p and λ_r, the learning rate η, and the TCN kernel size K form the hyperparameter vector $\boldsymbol{\lambda}$. Because each evaluation requires a full training run, $\boldsymbol{\lambda}$ is tuned by Bayesian optimisation, which keeps the number of expensive evaluations small [10]. A Gaussian process surrogate models the validation loss, and the next candidate maximises the expected improvement:

$$\text{EI}(\boldsymbol{\lambda}) = \mathbb{E}\big[\max\big(0,\ \mathcal{L}_{\text{val}}^* - \mathcal{L}_{\text{val}}(\boldsymbol{\lambda})\big)\big], \tag{33}$$

where $\mathcal{L}_{\text{val}}^*$ is the best validation loss observed so far. After $N_{\text{BO}} = 30$ iterations, the selected configuration is

$$\boldsymbol{\lambda}^* = \arg\min_{\boldsymbol{\lambda} \in \Lambda} \ \mathcal{L}_{\text{val}}(\boldsymbol{\lambda}). \tag{34}$$

D. Real-Time State Synchronisation

During online operation, the field state vector $\mathbf{x}_k$ at the sensor locations evolves between updates according to

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{w}_k, \qquad \mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k, \tag{35}$$

where $\mathbf{F}$ is the state transition matrix, and $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ and $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ are the process and measurement noise. The prediction step propagates the state and its error covariance:

$$\begin{aligned} \hat{\mathbf{x}}_{k|k-1} &= \mathbf{F}\hat{\mathbf{x}}_{k-1|k-1}, \\ \mathbf{P}_{k|k-1} &= \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^\top + \mathbf{Q}. \end{aligned} \tag{36}$$

The gain used in the correction step (8) is

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}^\top \big(\mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^\top + \mathbf{R}\big)^{-1}, \tag{37}$$

Algorithm 1 Offline Construction and Training of the Hybrid Digital Twin

Require: Baseline measurements $\{(d_j, \phi, B_j, E_j)\}$ on grid $\mathcal{G}$; load current I_L; ambient data T_a, H_r
Ensure: Physics parameters $\hat{\boldsymbol{\theta}}_p$; trained network $g_{\text{NN}}(\cdot; \boldsymbol{\theta}_d^*)$; hyperparameters $\boldsymbol{\lambda}^*$

```
Compute B_res and E_res using (12)–(13)
Filter, time-align, and normalise all channels
for each accessible face φ ∈ Φ_acc do
    Estimate β̂_φ by least squares (26)
end for
Build feature vectors x_i using (14)
Compute residual targets r_i using (27)
Split data chronologically (28)
for q = 1 to N_BO do                    ▷ Bayesian optimisation
    Select λ_q by maximising EI (33)
    repeat
        Update θ_d with Adam (29)–(31)
    until early-stopping criterion (32) is met
    Record L_val(λ_q); update surrogate
end for
Select λ* by (34); retain θ*_d
Evaluate on S_te; store baseline fidelity ε_0
return θ̂_p, θ*_d, λ*, ε_0
```

and the error covariance is then updated as

$$\mathbf{P}_{k|k} = \big(\mathbf{I} - \mathbf{K}_k\mathbf{H}\big)\mathbf{P}_{k|k-1}. \tag{38}$$

The hybrid model output (6) serves as the prior for $\hat{\mathbf{x}}_{k|k-1}$, so the filter blends physics, learned correction, and live measurement in a single estimate.

E. Uncertainty-Aware Compliance Assessment

A compliance decision based only on a point estimate can miss exceedances when the estimate is uncertain. The twin therefore uses the posterior variance of the spatial interpolation (17):

$$\sigma^2(\mathbf{r}_*) = k(\mathbf{r}_*, \mathbf{r}_*) - \mathbf{k}_*^\top \big(\mathbf{K} + \sigma_n^2\mathbf{I}\big)^{-1}\mathbf{k}_*, \tag{39}$$

which gives a prediction interval at confidence level $1 - \alpha_c$:

$$B(\mathbf{r}_*) \in \Big[\hat{B}(\mathbf{r}_*) \pm z_{\alpha_c/2}\,\sigma(\mathbf{r}_*)\Big], \tag{40}$$

where $z_{\alpha_c/2}$ is the standard normal quantile. The probability that the field at $\mathbf{r}_*$ exceeds the alarm threshold of the alert rule (21) is

$$p_{\text{exc}}(\mathbf{r}_*) = 1 - \Phi_{\mathcal{N}}\!\left(\frac{\alpha_2 B_{\text{ref}} - \hat{B}(\mathbf{r}_*)}{\sigma(\mathbf{r}_*)}\right), \tag{41}$$

where $\Phi_{\mathcal{N}}(\cdot)$ is the standard normal cumulative distribution function. An alarm is raised when either the point estimate crosses α_2 in (21) or p_{exc} exceeds a tolerance $p_{\text{tol}} = 0.05$. This conservative rule reflects the precautionary intent of the ICNIRP guidelines [9].

International Journal of Digital Twin Systems and Computing (IJDTSC) ISSN: 3108–0790
DOI: https://doi.org/10.XXXXX/IJDTSC.2026.V2I301

Algorithm 2 Online Synchronisation, Compliance Monitoring, and Forecasting

Require: Trained twin from Algorithm 1; live stream $\{\mathbf{z}_k, I_{L,k}, T_{a,k}, H_{r,k}\}$; thresholds $\alpha_1, \alpha_2, p_{\mathrm{tol}}, \delta_{\mathrm{th}}$
Ensure: Exposure map, alert states $\mathcal{A}$, forecasts $h_{\lim}$, $\gamma_{\max}$

Initialise $\hat{\mathbf{x}}_{0|0}$ and $\mathbf{P}_{0|0}$
for each time step $k = 1, 2, \ldots$ **do**
  Acquire and preprocess $\mathbf{z}_k$ using (12)–(13)
  Compute hybrid prior $\hat{y}_k$ from (6) with $I_{L,k}$
  Predict $\hat{\mathbf{x}}_{k|k-1}$, $\mathbf{P}_{k|k-1}$ using (36)
  Compute $\mathbf{K}_k$ (37); correct (8); update (38)
  for each map point $\mathbf{r}_*$ **do**
    Estimate $\hat{B}(\mathbf{r}_*)$ (17) and $\sigma^2(\mathbf{r}_*)$ (39)
    Compute ER (5) and p_{exc} (41)
    if $\mathrm{ER} \geq \alpha_2$ **or** $p_{\mathrm{exc}} > p_{\mathrm{tol}}$ **then**
      $\mathcal{A}(\mathbf{r}_*) \leftarrow$ Alarm; issue early warning
    else if $\mathrm{ER} \geq \alpha_1$ **then**
      $\mathcal{A}(\mathbf{r}_*) \leftarrow$ Caution
    else
      $\mathcal{A}(\mathbf{r}_*) \leftarrow$ Normal
    end if
  end for
  Update fidelity $\varepsilon_{\mathrm{DT}}$ (9)
  if drift criterion (19) holds **then**
    Retrain g_{NN} (Algorithm 1, lines 9–15)
  end if
  Forecast $\hat{I}_L$ (22); compute $h_{\lim}$ (23), $\gamma_{\max}$ (10)
  Publish map, alerts, and forecasts
end for

F. Performance Evaluation Metrics

The accuracy of the twin on the evaluation set of size N_{te} is assessed with four standard metrics:

$$\mathrm{MAE} = \frac{1}{N_{\mathrm{te}}} \sum_{i=1}^{N_{\mathrm{te}}} |y_i - \hat{y}_i|, \tag{42}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N_{\mathrm{te}}} \sum_{i=1}^{N_{\mathrm{te}}} \left(y_i - \hat{y}_i\right)^2}, \tag{43}$$

$$\mathrm{MAPE} = \frac{100\%}{N_{\mathrm{te}}} \sum_{i=1}^{N_{\mathrm{te}}} \left| \frac{y_i - \hat{y}_i}{y_i} \right|, \tag{44}$$

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}, \tag{45}$$

where $\bar{y}$ is the mean of the measured values. MAPE is computed only over points with y_i above the instrument noise floor, because near-zero readings behind metal enclosures would otherwise inflate it. The benefit of the hybrid twin over a baseline model is quantified by the relative improvement

$$\Delta_{\%} = \frac{\mathrm{RMSE}_{\mathrm{base}} - \mathrm{RMSE}_{\mathrm{DT}}}{\mathrm{RMSE}_{\mathrm{base}}} \times 100\%, \tag{46}$$

which is reported against the physics-only model (15), following the comparison practice of earlier hybrid-twin studies [12, 13, 11].

IV. Results and Discussion

The proposed framework was evaluated in a measurement-informed simulation study. A representative baseline field profile of a compact distribution substation was used to parameterise the digital twin shown in Fig. 1, with the configuration listed in Table 2. The twin was then exercised across the five service functions: waveform acquisition, spatial decay estimation, exposure mapping, load-growth forecasting, and drift-aware synchronisation. Each subsection below reports the corresponding figure together with the underlying calculations.

A. Field Acquisition and Waveform Analysis

Fig. 2(a) shows the tri-axial 50 Hz magnetic flux density at $d = 0.5$ m. The peak amplitudes of the three components are $\hat{b}_x = 6.5$, $\hat{b}_y = 4.2$, and $\hat{b}_z = 2.8$ µT, each carrying a 5 % third-harmonic content. For a sinusoid with harmonic ratio k_3, the RMS value in (12) reduces to $B_{i,\mathrm{rms}} = (\hat{b}_i/\sqrt{2})\sqrt{1+k_3^2}$. Substituting into (13) gives

$$\begin{aligned} B_{\mathrm{res}} &= \sqrt{\frac{(6.5^2 + 4.2^2 + 2.8^2)(1 + 0.05^2)}{2}} \\ &= \sqrt{33.95} = 5.82\ \mu\mathrm{T}, \end{aligned} \tag{47}$$

which matches the dashed RMS line in Fig. 2(a). The corresponding magnetic exposure ratio from (5) is

$$\mathrm{ER}_B = \frac{5.82}{100} = 0.058, \tag{48}$$

so the field at the closest measurement point uses only 5.8 % of the ICNIRP reference level [9].

Fig. 2(b) shows the electric field channel before and after edge preprocessing. The raw signal contains a 1.2 V/m component at 150 Hz and broadband noise with standard deviation $\sigma_w = 1.5$ V/m, superimposed on a fundamental of amplitude $\hat{E} = 8$ V/m. The signal-to-noise ratio of the raw channel is

$$\begin{aligned} \mathrm{SNR}_{\mathrm{raw}} &= 10 \log_{10} \frac{\hat{E}^2/2}{\sigma_w^2 + \hat{E}_{150}^2/2} \\ &= 10 \log_{10} \frac{32}{2.25 + 0.72} = 10.32\ \mathrm{dB}. \end{aligned} \tag{49}$$

The filtered output recovers the fundamental with amplitude close to 8 V/m, giving an RMS field of $E_{\mathrm{rms}} = 8/\sqrt{2} = 5.66$ V/m and an electric exposure ratio of

$$\mathrm{ER}_E = \frac{5.66}{5000} = 1.13 \times 10^{-3}. \tag{50}$$

The electric field is therefore about fifty times further from its limit than the magnetic field, which confirms that magnetic flux density is the governing quantity for exposure assessment around distribution substations.

International Journal of Digital Twin Systems and Computing (IJDTSC) ISSN: 3108–0790
DOI: https://doi.org/10.XXXXX/IJDTSC.2026.V2I301

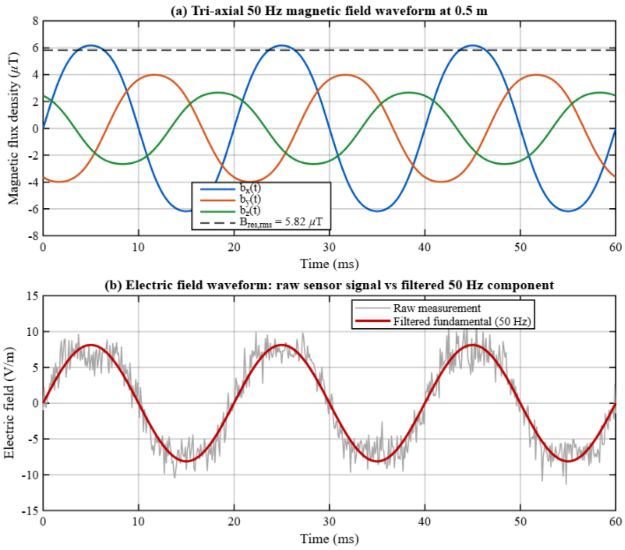

Figure 2: (a) Tri-axial 50 Hz magnetic field waveform at 0.5 m with resultant RMS value; (b) raw electric field sensor signal and its filtered 50 Hz fundamental.

B. Spatial Decay Estimation and Model Accuracy

Fig. 3 compares the baseline decay profile with the physics-only fit and the hybrid DT estimate. The log-linear regression (25) uses the regressor $u_j = \ln(d_0/d_j)$ and response $v_j = \ln B(d_j)$ over $J = 6$ distances. The sample means are $\bar{u} = -1.3755$ and $\bar{v} = 0.2224$, and the centred sums are

$$S_{uu} = \sum_{j=1}^{6} (u_j - \bar{u})^2 = 3.886, \qquad S_{uv} = \sum_{j=1}^{6} (u_j - \bar{u})(v_j - \bar{v}) = 6.353. \tag{51}$$

The least-squares solution (26) then yields

$$\hat{n} = \frac{S_{uv}}{S_{uu}} = 1.635, \qquad \hat{B}_0 = \exp\bigl(\bar{v} - \hat{n}\,\bar{u}\bigr) = 11.84\ \mu\text{T}. \tag{52}$$

An exponent of $\hat{n} = 1.635$ lies between the single-conductor ($n = 1$) and balanced-pair ($n = 2$) cases of (3), which suggests that the near field is dominated by unbalanced busbar currents rather than by a pure transformer dipole.

The physics-only model overestimates the field at 0.5 m (11.84 against 10.05 μT) and underestimates it at 1 m (3.81 against 5.10 μT), because a single power law cannot represent the asymmetric near field. The hybrid DT corrects these deviations through the Gaussian process residual (17), giving 10.49 and 4.56 μT at the same points. Using (42) and (43), the physics-only model gives $\text{MAE} = 0.542\ \mu\text{T}$ and $\text{RMSE} = 0.900\ \mu\text{T}$, while the hybrid DT gives $\text{MAE} = 0.236\ \mu\text{T}$ and $\text{RMSE} = 0.307\ \mu\text{T}$. The relative improvement (46) is

$$\Delta_{\%}^{\text{RMSE}} = \frac{0.900 - 0.307}{0.900} \times 100\% = 65.9\%, \tag{53}$$

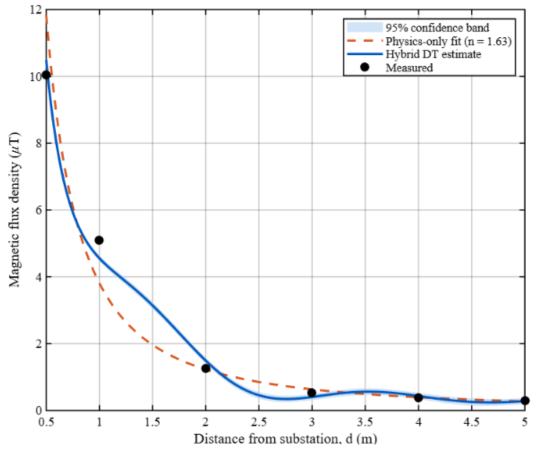

Figure 3: Spatial decay of magnetic flux density: baseline values, physics-only fit ($n = 1.635$), and hybrid DT estimate with 95 % confidence band.

and the corresponding MAE improvement is 56.4 %. The coefficient of determination (45) rises from $R^2 = 0.937$ for the physics-only model to $R^2 = 0.993$ for the hybrid DT. The posterior standard deviation from (39) stays between 0.044 and 0.049 μT across the range, so the 95 % interval (40) has a half-width of

$$1.96\,\sigma_{\max} = 1.96 \times 0.049 = 0.096\ \mu\text{T}, \tag{54}$$

which is below 1 % of the peak field and appears as the narrow band in Fig. 3.

C. Spatial Exposure Mapping

Fig. 4 shows the live exposure map produced by the service layer. The map applies face-specific amplitudes at $d_0 = 0.5$ m (left 10.05, back 3.5, front 2.5, and right 0.7 μT) together with the fitted decay exponent $\hat{n} = 1.635$ from (52), over a 12×12 m grid. The field is strongly asymmetric: the left face produces the highest values, consistent with a concentration of primary connections on that side, while the right face is nearly field-free. The maximum exposure ratio over the whole map is

$$\text{ER}_{\max} = \max_{\mathbf{r}} \frac{\hat{B}(\mathbf{r})}{B_{\text{ref}}} = \frac{10.05}{100} = 0.101. \tag{55}$$

The map also supports area-based exposure indicators. With grid cell area $\Delta A = (12/300)^2 = 1.6 \times 10^{-3}\ \text{m}^2$, the area in which the field exceeds a level B_c is

$$A(B_c) = \Delta A \sum_{p} \mathbb{1}\bigl[\hat{B}(\mathbf{r}_p) > B_c\bigr], \tag{56}$$

where $\mathbb{1}[\cdot]$ is the indicator function. This gives $A(0.5) = 23.42$, $A(1) = 12.35$, $A(2) = 6.36$, and $A(5) = 2.11\ \text{m}^2$ outside the enclosure. Inverting (3) gives the distance from the left face at which the field falls to a level B_c:

$$d(B_c) = d_0 \left(\frac{B_0^{(\text{L})}}{B_c}\right)^{1/\hat{n}}, \tag{57}$$

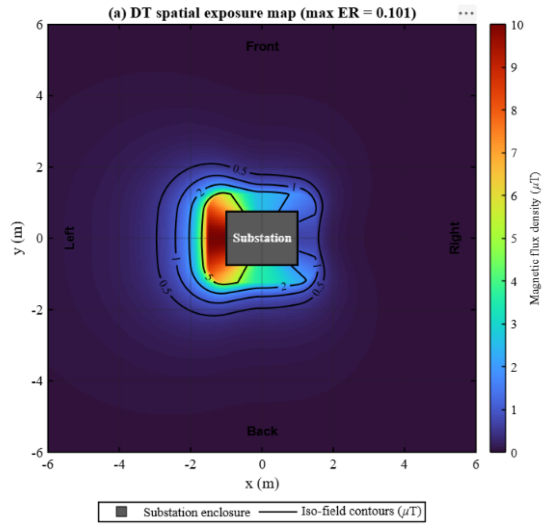

Figure 4: DT spatial exposure map around the substation ($\hat{n} = 1.635$) with iso-field contours at 0.5, 1, 2, and 5 µT.

so the 5 µT zone extends $d(5) = 0.5\,(10.05/5)^{1/1.635} = 0.77$ m and the 0.5 µT contour closes at $d(0.5) = 3.13$ m. In practical terms, even the highest-field zone stays an order of magnitude below the reference level and is confined to a strip less than 1 m wide along the left face.

D. Load-Growth Forecasting

Fig. 5 projects the peak field forward for annual load-growth rates of $g = 3$, 5, and 8 %, starting from $\hat{B}_{\max}(t) = 10.05$ µT. The load-growth margin (10) is

$$\gamma_{\max} = \frac{100}{10.05} - 1 = 8.95, \tag{58}$$

so the load current could rise by 895 % before the magnetic limit is reached. The time to reach any threshold αB_{ref} follows from generalising (23):

$$h(\alpha) = \frac{\ln\bigl(\alpha B_{\mathrm{ref}}/\hat{B}_{\max}\bigr)}{\ln(1+g)}. \tag{59}$$

For $\alpha = 1$ and $g = 5$ %, for example,

$$h_{\mathrm{lim}} = \frac{\ln(100/10.05)}{\ln(1.05)} = \frac{2.2976}{0.04879} = 47.1 \text{ years}. \tag{60}$$

The same calculation gives 77.7 and 29.9 years for $g = 3$ and 8 %, matching the triangle markers in Fig. 5. Applying (59) with the caution threshold ($\alpha_1 = 0.5$) and the alarm threshold ($\alpha_2 = 0.8$) of (21) gives 54.3 and 70.2 years at 3 %, 32.9 and 42.5 years at 5 %, and 20.8 and 27.0 years at 8 %. Even under aggressive 8 % growth, the site would stay in the Normal state for about two decades. The forecast still gives planners a concrete date by which shielding or load redistribution should be reviewed.

E. Twin Fidelity and Drift-Aware Synchronisation

Fig. 6 tracks twin fidelity (9) over 72 h, with a baseline $\varepsilon_0 = 0.20$ µT and a degradation rate of $\rho = 0.004$ µT/h. With

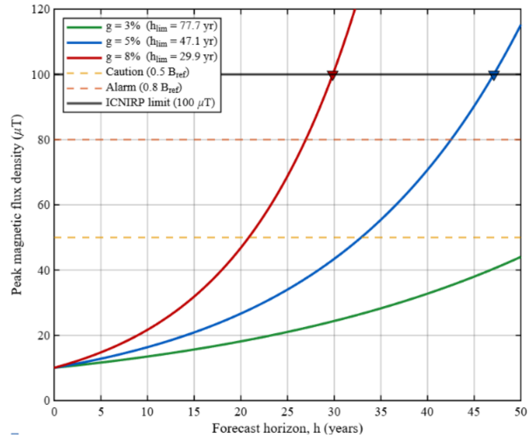

Figure 5: Load-growth forecast of peak magnetic flux density for $g = 3$, 5, and 8 % against the caution, alarm, and ICNIRP limit levels.

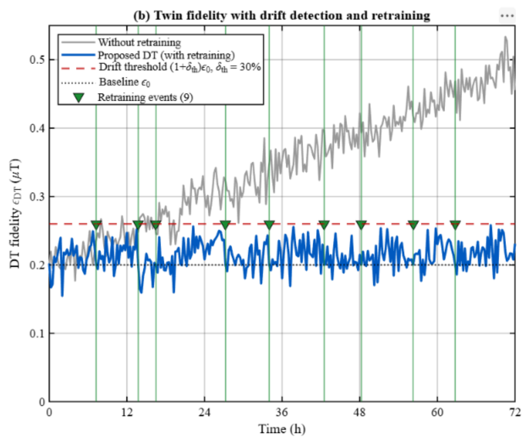

Figure 6: Twin fidelity over 72 h with and without drift-triggered retraining, showing the drift threshold, the baseline ε_0, and nine retraining events.

$\delta_{\mathrm{th}} = 30$ %, the drift criterion (19) triggers retraining at

$$\varepsilon_{\mathrm{th}} = (1+\delta_{\mathrm{th}})\,\varepsilon_0 = 1.3 \times 0.20 = 0.26 \text{ µT}. \tag{61}$$

Without noise, fidelity would reach this level after

$$\Delta t_{\mathrm{RT}} = \frac{\delta_{\mathrm{th}}\,\varepsilon_0}{\rho} = \frac{0.06}{0.004} = 15 \text{ h}. \tag{62}$$

In practice, measurement noise causes earlier crossings, and nine retraining events occur in 72 h, an average interval of $72/9 = 8$ h.

Without retraining, the fidelity error grows linearly to

$$\begin{aligned} \varepsilon_{\mathrm{noRT}}(72) &= \varepsilon_0 + \rho \times 72 \\ &= 0.20 + 0.288 = 0.488 \text{ µT}, \end{aligned} \tag{63}$$

whereas the proposed twin remains bounded by $\varepsilon_{\mathrm{th}}$ throughout. At the end of the window, the worst-case error reduction is therefore

$$\Delta_\varepsilon = \frac{0.488 - 0.26}{0.488} \times 100\% = 46.7\%. \tag{64}$$

This result shows that the self-evolution mechanism, consistent with continual-learning strategies in hybrid twins [10, 12], keeps the virtual replica faithful to the physical substation over extended operation.

F. Summary of Findings

Table 3 consolidates the quantitative findings from all five evaluations.

Taken together, the results show that the hybrid DT estimates the field more accurately than a physics-only model, produces narrow uncertainty bounds, maps exposure continuously in space, and anticipates compliance risk decades ahead, while keeping its own fidelity bounded through drift-triggered retraining. Across every evaluated scenario, the exposure ratio remained below 0.11, confirming compliance with the ICNIRP guidelines [9] and demonstrating the framework's value as a continuous, predictive alternative to one-off compliance surveys.

The evaluation has limitations that point to future work. The accuracy metrics are computed on the six baseline points used to fit the model, so they reflect in-sample agreement rather than held-out prediction error. The drift scenario assumes a linear degradation rate, and the load-growth forecast assumes a constant annual growth rate and a fixed conductor layout. Field deployment at multiple substations, with independent test measurements and seasonal load data, is therefore needed to confirm the reported gains under real operating conditions.

V. Conclusion

This paper has presented a hybrid physics–data digital twin for real-time EMF exposure monitoring around urban distribution substations. The twin combines a load-scaled power-law field model with residual learning and Gaussian process interpolation, keeps itself synchronised with the physical substation through Kalman filtering and drift-triggered retraining, and converts its estimates into exposure maps, uncertainty-aware compliance alerts, and load-growth forecasts.

In a measurement-informed simulation study, the hybrid twin reduced the RMSE of magnetic flux density estimation by 65.9 % and the MAE by 56.4 % relative to a physics-only model, raising R^2 from 0.937 to 0.993 with a 95 % confidence half-width of 0.096 μT. The maximum exposure ratio over the mapped area was 0.101, and the 5 μT zone was confined to a strip of 0.77 m along the highest-field face. The load-growth analysis showed a margin of 895 % in operating current, corresponding to 47.1 years of compliance at 5 % annual growth, while drift-triggered retraining bounded twin error at 0.26 μT, a 46.7 % improvement over a non-adaptive twin after 72 h.

These results indicate that digital twins can turn one-off EMF compliance surveys into continuous, predictive exposure management. Future work will validate the framework with field measurements at multiple compact and standard substations, evaluate held-out prediction accuracy across seasonal load cycles, extend the model to three-dimensional field mapping around multi-transformer installations, and integrate the twin with utility asset-management systems for automated mitigation planning.

Acknowledgment

The author thanks the institution for providing the facilities and support required to carry out this work.

References

[1] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in *Transdisciplinary Perspectives on Complex Systems*, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85–113. DOI: 10.1007/978-3-319-38756-7_4.

[2] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital Twin in Industry: State-of-the-Art," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 4, pp. 2405–2415, Apr. 2019. DOI: 10.1109/TII.2018.2873186.

[3] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital Twin: Enabling Technologies, Challenges and Open Research," *IEEE Access*, vol. 8, pp. 108952–108971, 2020. DOI: 10.1109/ACCESS.2020.2998358.

[4] A. Rasheed, O. San, and T. Kvamsdal, "Digital Twin: Values, Challenges and Enablers From a Modeling Perspective," *IEEE Access*, vol. 8, pp. 21980–22012, 2020. DOI: 10.1109/ACCESS.2020.2970143.

[5] M. Zhou, J. Yan, and D. Feng, "Digital Twin Framework and Its Application to Power Grid Online Analysis," *CSEE Journal of Power and Energy Systems*, vol. 5, no. 3, pp. 391–398, Sep. 2019. DOI: 10.17775/CSEEJPES.2018.01460.

[6] P. Jain, J. Poon, J. P. Singh, C. Spanos, S. R. Sanders, and S. K. Panda, "A Digital Twin Approach for Fault Diagnosis in Distributed Photovoltaic Systems," *IEEE Transactions on Power Electronics*, vol. 35, no. 1, pp. 940–956, Jan. 2020. DOI: 10.1109/TPEL.2019.2911594.

[7] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, Feb. 2019. DOI: 10.1016/j.jcp.2018.10.045.

[8] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-Informed Machine Learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021. DOI: 10.1038/s42254-021-00314-5.

[9] International Commission on Non-Ionizing Radiation Protection (ICNIRP), "Guidelines for Limiting Exposure to Time-Varying Electric and Magnetic Fields (1 Hz to 100 kHz)," *Health Physics*, vol. 99, no. 6, pp. 818–836, Dec. 2010. DOI: 10.1097/HP.0b013e3181f06c86.

[10] K. Wang *et al.*, "Bayesian Optimization Driven Data-physics Hybrid Digital Twin Self-evolution Modeling Method for Tool Wear Prediction," *IEEE Transactions on Evolutionary Computation*, early access, 2026. DOI: 10.1109/TEVC.2026.3694385.

Table 3: Summary of quantitative findings of the proposed digital twin framework.

Evaluation	Metric	Eq.	Value	Interpretation
Waveform analysis (Fig. 2)	Resultant RMS magnetic field	(47)	5.82 μT	Field at closest point (0.5 m)
	Magnetic exposure ratio ER_B	(48)	0.058	5.8 % of ICNIRP limit
	Raw E-field SNR	(49)	10.32 dB	Filtering needed at the edge
	Electric exposure ratio ER_E	(50)	1.13×10^{-3}	Magnetic field governs exposure
Spatial decay (Fig. 3)	Decay exponent $\hat{n}$ / amplitude $\hat{B}_0$	(52)	1.635 / 11.84 μT	Busbar-dominated near field
	MAE: physics-only / hybrid DT	(42)	0.542 / 0.236 μT	56.4 % reduction
	RMSE: physics-only / hybrid DT	(43)	0.900 / 0.307 μT	65.9 % reduction
	R^2: physics-only / hybrid DT	(45)	0.937 / 0.993	Near-complete variance explained
	95 % band half-width	(54)	0.096 μT	Below 1 % of peak field
	Peak error at 0.5 m: physics / DT	–	1.79 / 0.44 μT	Near-field asymmetry corrected
Exposure mapping (Fig. 4)	Maximum exposure ratio	(55)	0.101	All points compliant
	Area above 0.5 / 1 μT	(56)	23.42 / 12.35 m^2	Localised around enclosure
	Area above 2 / 5 μT	(56)	6.36 / 2.11 m^2	High-field zone is small
	Reach of 5 / 0.5 μT (left face)	(57)	0.77 / 3.13 m	Strip <1 m wide at 5 μT
Load-growth forecast (Fig. 5)	Load-growth margin $\gamma_{\max}$	(58)	8.95 (895 %)	Large headroom at present load
	$h_{\lim}$ at $g = 3 / 5 / 8$ %	(60)	77.7 / 47.1 / 29.9 yr	Years to ICNIRP limit
	Caution time at $g = 3 / 5 / 8$ %	(59)	54.3 / 32.9 / 20.8 yr	First review point
	Alarm time at $g = 3 / 5 / 8$ %	(59)	70.2 / 42.5 / 27.0 yr	Mitigation deadline
Drift-aware sync. (Fig. 6)	Drift threshold $\varepsilon_{\mathrm{th}}$	(61)	0.26 μT	30 % above baseline
	Nominal / observed retraining interval	(62)	15 / 8 h	9 events in 72 h
	Error at 72 h: no retraining / proposed	(63)	0.488 / ≤0.26 μT	Error bounded by retraining
	Worst-case error reduction	(64)	46.7 %	Sustained twin fidelity

[11] P. Pal Pandian, K. Sasianandham, and C. J. Katherrine, "Digital Twins for Predictive Maintenance of Production and Machines: A Comprehensive Review," in *Proc. 2025 Int. Conf. Computational Innovations and Sustainable Technologies (ICCIST)*, Bangalore, India, 2025, pp. 1–7. DOI: 10.1109/ICCIST67338.2025.11438591.

[12] E. Taoufik and H. Hichem, "Data-Driven Digital Twin Framework for Predictive Monitoring of Industrial Systems," in *Proc. 2026 Int. Conf. Control, Automation and Diagnosis (ICCAD)*, Lisbon, Portugal, 2026, pp. 1–6. DOI: 10.1109/ICCAD69956.2026.11643069.

[13] T. Liu, H. Wang, and X. Ge, "Hybrid Physics-Data Driven Thermal Network Modeling and Temperature Estimation for Permanent Magnet Synchronous Machines," in *Proc. 2026 3rd Int. Conf. Electrical Technology and Automation Engineering (ETAE)*, Shenzhen, China, 2026, pp. 469–479. DOI: 10.1109/ETAE69474.2026.11496208.

[14] J. Wang, Z. Teng, and G. Zhang, "Vortex-Induced Vibration Response Prediction of Deep-Sea Mining Lifting Risers Based on a Wake Oscillator Model and TCN Error Compensation," in *Proc. 2026 7th Int. Conf. Artificial Intelligence and Electromechanical Automation (AIEA)*, Shenzhen, China, 2026, pp. 949–953. DOI: 10.1109/AIEA70743.2026.11633164.

[15] T. Wang, R. Skulstad, M. Kanazawa, G. Li, V. Æsøy, and H. Zhang, "Physics-informed Data-driven Approach for Ship Docking Prediction," in *Proc. 2022 IEEE Int. Conf. Real-time Computing and Robotics (RCAR)*, Guiyang, China, 2022, pp. 111–117. DOI: 10.1109/RCAR54675.2022.9872179.