

A Structured Neural ODE Approach for Real-Time Evaluation of AC Losses in 3D Superconducting Tapes

R. Basei ^{*†} F. Pase [‡] F. Lucchini [†] F. Toso [‡] R. Torchio ^{†§}

Abstract

Efficient modeling of High Temperature Superconductors (HTS) is crucial for real-time quench monitoring; however, full-order electromagnetic simulations remain prohibitively costly due to the strong nonlinearities. Conventional projection-based reduced-order modeling pipelines for nonlinear problems, such as Proper Orthogonal Decomposition (POD)–Discrete Empirical Interpolation Method (DEIM), alleviate this cost but often require intrusive access to the Full Order Model (FOM) operators and a substantial number of interpolation points for hyperreduction. This work investigates reduced-order strategies for Integral Equation Method (IEM) of HTS systems. We present the first application of POD–DEIM to IEM-based HTS models, and introduce a Structured Neural Ordinary Differential Equation (Neural ODE) approach that learns nonlinear dynamics directly in the reduced space. The benchmark results show that Neural ODE outperforms POD–DEIM both in efficiency and accuracy, highlighting its potential for real-time simulations of superconductors.

Keywords: High Temperature Superconductor, AC Losses, Integral Equation Method, Proper Orthogonal Decomposition, Discrete Empirical Interpolation Method, Neural ODE

1 Introduction

The second generation (2G) of Rare-Earth Barium Copper Oxide (REBCO)-based HTS has emerged as a promising technology for a wide range of electrical engineering applications. These include power distribution networks, high-field magnets for experimental magnetic confinement fusion devices [1], and particle accelerators [2]. HTS tapes can be configured into several cable architectures, such as the twisted stacked-tape conductor [3], the ROEBEL cable [4], and the Conductor on Round Core (CORC[®]) cable [5]. Among these, the CORC[®] design has attracted significant attention due to its advantageous mechanical strength and electromagnetic performance [5].

Although HTS offer the major advantage of operating at significantly higher temperatures than conventional Low Temperature Superconductors (LTS), monitoring the quench phenomenon remains a critical challenge. This difficulty arises from the relatively slow reduction of the critical current I_c with increasing temperature [6, 7]. The sudden transition from the superconducting state to the normal resistive state generates localized hot spots through the Joule effect, which, if not properly controlled, can lead to irreversible damage of the conductor. Reliable quench detection and mitigation are therefore essential for the safe and continuous operation of future fusion demonstrator power plants, such as the European DEMO [8].

Modeling quench phenomena requires solving multiphysics problems that couple electromagnetics (EM), thermal behavior, and fluid dynamics [9]. While lumped-element models have been proposed, capturing the underlying physics with sufficient accuracy typically demands finite element analysis. For superconducting structures, the EM problem can be addressed using several numerical formulations, most notably the Finite Element Method (FEM) with magnetic field (H) formulations [10], or the T – A approach [11]. An alternative is IEM, which is traditionally used in open-domain problems such as antenna radiation, but has also proven effective in computing AC losses in superconductors [12, 13].

Due to the strong nonlinearities in the resistivity model, both FEM and IEM simulations require substantial computational resources to capture the transient, nonlinear dynamics of realistic 3D superconducting structures. This high computational cost renders them impractical for real-time monitoring.

^{*}riccardo.basei@phd.unipd.it (Corresponding Author)

[†]Department of Industrial Engineering, University of Padua

[‡]Newtven

[§]Department of Information Engineering, University of Padua

To address this limitation, several studies have explored the use of surrogate models to accelerate the EM solution, particularly in the context of estimating AC loss [14, 15]. A first attempt to apply POD to HTS systems within the IEM framework was reported in [16]. However, in this work, no hyperreduction strategy was adopted, as the authors argued that the sparsity of the operator \mathbf{R} kept the assembly cost manageable. This is true for relatively small problems, but as the mesh dimension increases, the assembly quickly becomes considerably more demanding. In such cases, Empirical Interpolation Method (EIM) [17] and its discrete counterpart [18] can be applied to alleviate this cost, and these techniques have already been successfully used in electromagnetic problems [19]. Nevertheless, DEIM suffers from two critical drawbacks in the present context: the strong nonlinearity and spatial locality of superconducting material laws require a large number of interpolation points, which reduces efficiency, and its intrusive nature demands full access to the high-fidelity model assembly routines, which limits practical applicability. Furthermore, since the governing equations are highly nonlinear, the reduced-order model must still be solved with an implicit time-integration scheme, which entails a large number of Newton iterations and repeated Jacobian evaluations. As a result, even with DEIM, the simulation cost remains significant.

These challenges have motivated a growing interest in Machine Learning (ML)-based modeling strategies, which are typically non-intrusive and data-driven. Here, we provide only a brief overview of the main approaches, without aiming to be exhaustive; several recent reviews give a comprehensive account of this rapidly evolving area [20, 21, 22].

Neural Operators such as DeepONet [23], Fourier Neural Operator [24], and graph neural operator variants [25] aim to learn mappings from function spaces, rather than between finite-dimensional vectors. This formulation enables them to generalize across families of input functions, achieving mesh-independent predictions and fast inference once trained. Their flexibility and non-intrusivity make them particularly attractive for parametric problems and large-scale simulations. However, they also present drawbacks: training requires substantial data, the networks can be very large, and their ability to handle complex geometries is limited. Graph-based neural operators alleviate this to some extent, but at the price of predicting the full-order field directly, which implies very large network outputs and nontrivial computational costs.

Physics-Informed Neural Networks (PINNs) [26] incorporate the residual of the physical equation into the training loss, enabling the solution of forward and inverse problems even when data are scarce. PINNs are especially powerful for inverse problems and for settings where measurements are sparse. However, they tend to struggle with complex geometries, and their formulation is not well-suited for input-driven or control problems: a PINN is typically trained for a fixed set of inputs or excitations, and adapting it to new ones generally requires retraining the network. Additionally, their implementation requires formulating the governing equations as differentiable residuals, which can be cumbersome when the full-order model is complex or not available in closed form.

Operator Inference (OpInf) [27, 28] offers a physics-aware, non-intrusive alternative. Instead of learning black-box surrogates, OpInf identifies reduced operators directly from trajectory data projected onto a reduced basis. This approach has the advantage of being parsimonious, data-efficient, and interpretable, while remaining consistent with the structure of the Reduced Order Model (ROM). Nevertheless, OpInf typically assumes that the reduced operators have a polynomial structure, which can be restrictive in cases where the dynamics involve more complex nonlinear interactions.

Latent Dynamics modeling frameworks [29, 30] take yet another route. The central idea is that the solution trajectories of high-dimensional problems evolve on a low-dimensional manifold, which can be identified either with linear projections (e.g. POD) or with nonlinear encoders (e.g., autoencoders [31, 32]). Once the manifold is fixed, the latent dynamics can be learned in different ways. Neural ODEs [33] model continuous-time dynamics by learning the time derivative as a neural network. Sparse Identification of Nonlinear Dynamics (SINDy) [34] also parametrizes the derivative in continuous time, but restricts it to a sparse linear combination of predefined candidate functions, leading to interpretable governing equations. More recent developments further combine physics-based structure and data-driven learning for parametric reduced-order modeling [35], and explore advanced nonlinear manifold-learning strategies such as probabilistic manifold decomposition [36], as well as autoencoder-based ROMs enhanced through self-attention mechanisms [37].

In the HTS field as well, machine learning-based approaches have recently been applied to accelerate quench detection and AC loss estimation [38, 39]. However, their application to reduced-order modeling of nonlinear electromagnetic simulations remains largely unexplored.

In this work, we build on the latent dynamics paradigm and propose a novel neural architecture for nonlinear electromagnetic modeling that operates directly in the reduced-order space. The starting point




Figure 1: Geometrical model of HTS tapes under analysis (units in mm), and 2D surface mesh.

is a discretized integral equation model of HTS conductors, which is compressed through POD-based dimensionality reduction. Instead of relying on intrusive hyper-reduction techniques such as DEIM, we design a Structured Neural ODE that learns only the nonlinear component of the reduced dynamics, while the linear part is inherited directly from the full-order model. This idea of embedding problem structure into Neural ODEs architectures has also been explored in recent works on Structured Neural ODEs [40, 41, 42], albeit motivated primarily by stiffness and time-integration considerations. In this way, physical knowledge is embedded into the reduced-order evolution, ensuring parsimony, interpretability, and efficiency. The approach is validated on a HTS CORC[®] cable case study and benchmarked against the standard POD–DEIM method, demonstrating the potential of physics-informed latent dynamics modeling for challenging electromagnetic simulation. The results demonstrate improved accuracy and significantly lower computational cost, highlighting the potential of the proposed strategy for real-time, industrially relevant superconducting applications.

The main contributions of this work are twofold:

- We demonstrate, for the first time, the application of the POD-DEIM methodology to IEM-based models of HTS systems;
- We introduce and validate a Structured Neural ODE approach that directly learns the nonlinear dynamics in the reduced-order space, offering a less-intrusive alternative to traditional hyperreduction techniques.

The remainder of this paper is organized as follows. Section 2 describes the physical system under study, including the mesh and material properties. Section 3 introduces the IEM formulation. Section 4 discusses the application of the POD-DEIM method, while Section 5 presents the proposed Neural ODE approach. Section 6 reports the numerical results and Section 7 discusses them. Finally, Section 8 summarizes the conclusions and outlines directions for future research.

2 Model Description

The CORC[®] cables are realized by winding HTS tapes made of REBCO, arranged in a helical geometry, around a thick core made of Copper [43]. The cable is subjected to various AC loss mechanisms, including magnetization losses and coupling losses. From the EM modeling perspective, a common practice is to focus only on the HTS tapes and determine the AC losses due to the current loops therein [44, 45, 46, 47]. Moreover, the REBCO tapes with a thickness of approximately 1 μm , represent a small portion of the whole thickness of the HTS tapes ($\approx 100 \mu\text{m}$) [44]. Owing to the extremely high aspect ratio of the superconducting layer, the tapes are modeled as two-dimensional surfaces embedded in three-dimensional space, using surface formulations, e.g. exploiting the $T - A$ approach [13]. Following [44], the HTS tapes analyzed in this work are extracted from a CORC[®] conductor made of 3 tapes in a single layer, arranged in a helical geometry with a winding angle of 40° , representing an 18 mm-long segment of the cable. The diameter of the Copper former is 5.6 mm. The model representation is shown in Figure 1.

Within the superconducting domain $\Omega_{sc} \subseteq \Omega$, the resistivity depends on the local current density. In

this work, we adopt a power-law constitutive relation [48]:

$$\mathbf{E}(\mathbf{J}) = \frac{E_c}{J_c} \left(\frac{\|\mathbf{J}\|}{J_c} \right)^{n-1} \mathbf{J}, \quad (1)$$

where the electric field is assumed to be parallel to the current density. The corresponding resistivity can thus be expressed as

$$\rho_{sc}(\mathbf{J}) = \frac{E_c}{J_c} \left(\frac{\|\mathbf{J}\|}{J_c} \right)^{n-1}. \quad (2)$$

Here $E_c = 10^{-4}$ V/m is the critical electric field, $J_c = 236$ MA/m² is the critical current density, and $n = 25$ is the power-law exponent characterizing the superconducting material. In general, both n and J_c depend on temperature and magnetic field; here, they are assumed constant for simplicity. Being this constitutive equation highly nonlinear, the resulting problem is also highly nonlinear, which represents a significant challenge when solving it numerically.

We investigate the current density distribution induced by an external time-varying magnetic field. Specifically, the magnetic field is assumed to be oriented in the y -direction. This external field is modeled as:

$$\mathbf{B}_{\text{ext}}(t) = B_y(t)\hat{\mathbf{y}} = B_0 \sin(2\pi ft)\hat{\mathbf{y}}, \quad (3)$$

where B_0 is the magnetic field amplitude, and $f = 50$ Hz is the frequency. According to Faraday's law of induction, this time-varying magnetic field induces an electric field:

$$\nabla \times \mathbf{E}_{\text{ext}} = -\frac{\partial \mathbf{B}}{\partial t}. \quad (4)$$

For the given magnetic field orientation, the induced electric field can be expressed as

$$\mathbf{E}_{\text{ext}}(x, y, z) = \begin{bmatrix} -\frac{\partial B_y}{\partial t} z \\ 0 \\ \frac{\partial B_y}{\partial t} x \end{bmatrix}. \quad (5)$$

This induced electric field acts as the external excitation driving the current distribution in the superconducting tapes.

3 $J - \varphi_e$ Formulation

The $J - \varphi_e$ formulation [12] for solving electromagnetic problems is based on the Electric Field Integral Equation (EFIE), which serves as its starting point. We consider an electromagnetic analysis over a bounded domain $\Omega \subset \mathbb{R}^3$, where the EFIE is coupled with the current continuity equation:

$$\begin{aligned} \mathbf{E}(\mathbf{r}) &= -\frac{\partial \mathbf{A}(\mathbf{r})}{\partial t} - \nabla \varphi_e(\mathbf{r}) + \mathbf{E}_{\text{ext}}(\mathbf{r}), \\ \nabla \cdot \mathbf{J}(\mathbf{r}) &= 0, \end{aligned} \quad (6)$$

where $\mathbf{E}(\mathbf{r})$ is the electric field, $\mathbf{A}(\mathbf{r})$ is the magnetic vector potential, $\varphi_e(\mathbf{r})$ is the scalar electric potential, $\mathbf{J}(\mathbf{r})$ is the current density, and $\mathbf{E}_{\text{ext}}(\mathbf{r})$ is the external applied electric field. The constitutive relation between electric field and current density is given by Ohm's law:

$$\mathbf{E} = \rho(\mathbf{J})\mathbf{J}, \quad (7)$$

with $\rho(\mathbf{J})$ denoting the electrical resistivity. The magnetic vector potential \mathbf{A} can be expressed in terms of \mathbf{J} via the integral relation:

$$\mathbf{A}(\mathbf{r}) = \mu_0 \int_{\Omega} g(\mathbf{r}, \mathbf{r}') \mathbf{J}(\mathbf{r}') d\Omega, \quad (8)$$

where

$$g(\mathbf{r}, \mathbf{r}') = \frac{1}{4\pi \|\mathbf{r} - \mathbf{r}'\|} \quad (9)$$

is the static Green's function.

To solve (6) numerically, we adopt the Partial Equivalent Element Circuit (PEEC) methodology. The computational domain Ω is discretized into triangular elements. The current density and electric potential are expanded using vector basis functions $\mathbf{w}_h(\mathbf{r})$ (associated with edges), and piecewise constant basis functions $p_h(\mathbf{r})$ (associated with faces):

$$\mathbf{J}(\mathbf{r}) = \sum_{h=1}^{n_e} i_h \mathbf{w}_h(\mathbf{r}), \quad \varphi_e(\mathbf{r}) = \sum_{h=1}^{n_f} \varphi_j p_h(\mathbf{r}), \quad (10)$$

where n_e is the number of edges and n_f the number of faces. Applying Galerkin-testing to (6) yields a system of nonlinear Differential Algebraic Equation (DAE):

$$\begin{bmatrix} \mathbf{R}(\mathbf{i}) & \mathbf{G}^T \\ \mathbf{G} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i} \\ \varphi \end{bmatrix} + \begin{bmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \frac{d}{dt} \begin{bmatrix} \mathbf{i} \\ \varphi \end{bmatrix} = \begin{bmatrix} \mathbf{e}_s(t) \\ \mathbf{0} \end{bmatrix}. \quad (11)$$

Here $\mathbf{i} \in \mathbb{R}^{n_e}$ collects the unknown edge currents, and $\varphi \in \mathbb{R}^{n_f}$ contains the face-associated scalar potentials. The matrix $\mathbf{L} \in \mathbb{R}^{n_e \times n_e}$ is the inductance matrix, $\mathbf{R}(\mathbf{i}) : \mathbb{R}^{n_e} \rightarrow \mathbb{R}^{n_e \times n_e}$ is the current-dependent resistance operator, and \mathbf{G} is the faces-to-edges incidence matrix, acting as discrete counterpart of the divergence operator. The source term $\mathbf{e}_s(t)$ arises from the discretization of the external electric field.

Although this formulation leads to non-sparse matrices, it is advantageous because only the regions carrying current—namely, the superconducting domains—need to be meshed. This avoids the need to mesh the surrounding air, unlike the H and $H - \varphi_e$ formulations, which require solving for air elements as well. As a result, the overall model size is significantly reduced.

4 POD-DEIM

4.1 Structure-Preserving POD

The POD method [49] is used to derive a reduced-order basis for both current and potential fields while preserving the structure of the underlying equations [50]. The procedure is as follows:

1. Using FOM simulations, snapshots of the current and potential are collected:

$$\mathbf{I} = \begin{bmatrix} | & & | \\ \mathbf{i}^{(1)} & \dots & \mathbf{i}^{(n_s)} \\ | & & | \end{bmatrix} \in \mathbb{R}^{n_e \times n_s}, \quad \mathbf{\Phi} = \begin{bmatrix} | & & | \\ \varphi^{(1)} & \dots & \varphi^{(n_s)} \\ | & & | \end{bmatrix} \in \mathbb{R}^{n_f \times n_s}, \quad (12)$$

where n_e , n_f , and n_s are the number of edges, faces, and snapshot instances, respectively. Each column of the snapshot matrices corresponds to the system state at a given time instant, and the variability across snapshots arises from both temporal evolution and changes in the simulation setup, such as different inputs or initial conditions, depending on the specific application.

2. Apply Singular Value Decomposition (SVD) separately to the current and potential snapshot matrices:

$$\mathbf{V}_i \mathbf{\Sigma}_i \mathbf{U}_i^T = \mathbf{I}, \quad \mathbf{V}_\varphi \mathbf{\Sigma}_\varphi \mathbf{U}_\varphi^T = \mathbf{\Phi}. \quad (13)$$

The columns of \mathbf{V}_i and \mathbf{V}_φ define orthonormal basis vectors for the current and potential, respectively;

3. Retain r_i and r_φ dominant modes, forming the reduced basis

$$\mathbf{V}_{i,r} \in \mathbb{R}^{n_e \times r_i}, \quad \mathbf{V}_{\varphi,r} \in \mathbb{R}^{n_f \times r_\varphi}. \quad (14)$$

The state variables are projected onto the reduced space as:

$$\mathbf{i}_r = \mathbf{V}_{i,r}^T \mathbf{i}, \quad \varphi_r = \mathbf{V}_{\varphi,r}^T \varphi, \quad (15)$$

and the system matrices are modified accordingly:

$$\mathbf{R}_r(\mathbf{i}) = \mathbf{V}_{i,r}^T \mathbf{R}(\mathbf{i}) \mathbf{V}_{i,r}, \quad \mathbf{L}_r = \mathbf{V}_{i,r} \mathbf{L} \mathbf{V}_{i,r}^T, \quad \mathbf{G}_r = \mathbf{V}_{\varphi,r} \mathbf{G} \mathbf{V}_{\varphi,r}^T, \quad \mathbf{e}_{s,r}(t) = \mathbf{V}_{i,r}^T \mathbf{e}_s(t). \quad (16)$$

This structure-preserving POD approach ensures that the ROM captures the dominant dynamics of both the current and potential while maintaining the original physical structure of the equations.

An a priori estimate of the error introduced by truncating the POD modes is given by the retained energy ratio:

$$\epsilon = \frac{\sum_{k=1}^r \sigma_k}{\sum_{k=1}^{n_s} \sigma_k}, \quad (17)$$

where σ_k are the singular values of \mathbf{I} or Φ , stored in descending order in the diagonal of Σ_i or Σ_φ , respectively. A higher value of ϵ indicates a more accurate ROM relative to the full model.

4.2 DEIM

The DEIM [18] efficiently approximates nonlinear functions in ROMs by evaluating them at a small number of carefully selected spatial points. We define the nonlinear function:

$$\mathbf{f}(\mathbf{i}) := \mathbf{R}(\mathbf{i})\mathbf{i}. \quad (18)$$

DEIM proceeds by constructing a low-rank approximation of \mathbf{f} through the following steps:

1. *snapshots* of \mathbf{f} are generated under various input conditions by evaluating (18), using the FOM simulations performed for POD. These snapshots are assembled into a matrix:

$$\mathbf{F} = \begin{bmatrix} | & & | \\ \mathbf{f}^{(1)} & \dots & \mathbf{f}^{(n_s)} \\ | & & | \end{bmatrix} \in \mathbb{R}^{n_e \times n_s}. \quad (19)$$

2. as in POD, a SVD is applied to the snapshot matrix:

$$\mathbf{F} = \mathbf{V}_f \Sigma_f \mathbf{U}_f^T, \quad (20)$$

and the reduced basis $\mathbf{V}_{f,r} \in \mathbb{R}^{n_e \times r_{deim}}$ is formed from the first $r_{deim} \ll n$ columns of \mathbf{V}_f ;

3. a set of n_p optimal interpolation points is determined using a greedy algorithm [17]. These points are encoded in a selection matrix $\mathbf{P} \in \mathbb{R}^{n_e \times n_p}$.

The number of modes retained, r_{DEIM} , can be chosen based on the retained energy criterion, similar to POD.

Using DEIM, the nonlinear function is approximated as:

$$\mathbf{f}(\mathbf{i}) \approx \mathbf{V}_{f,r} (\mathbf{P}^T \mathbf{V}_{f,r})^{-1} \mathbf{P}^T \mathbf{f}(\mathbf{i}). \quad (21)$$

This formulation allows for evaluating the nonlinear term $\mathbf{f}(\mathbf{i})$ only at the selected spatial points determined by \mathbf{P} , significantly reducing computational effort.

4.3 POD-DEIM ROM Simulation

We have described how the system is assembled. To simulate its dynamics, we must discretize the system of ODEs in time and solve it step by step. We consider the time step interval $[t_n, t_{n+1}] \subset \mathbb{R}_{\geq 0}$ of length $\Delta t = t_{n+1} - t_n$. Given the initial values of the current and potential, $\mathbf{i}_r^{(n)} = \mathbf{i}_r(t_n)$ and $\varphi_r^{(n)} = \varphi_r(t_n)$, our goal is to compute the next-step values $\mathbf{i}_r^{(n+1)} = \mathbf{i}_r(t_{n+1})$ and $\varphi_r^{(n+1)} = \varphi_r(t_{n+1})$.

To this end, we approximate the time derivatives over $[t_n, t_{n+1}]$ using a first-order backward finite difference, namely

$$\frac{d\mathbf{i}}{dt} \approx \frac{\mathbf{i}_r^{(n+1)} - \mathbf{i}_r^{(n)}}{\Delta t}, \quad \frac{d\varphi}{dt} \approx \frac{\varphi_r^{(n+1)} - \varphi_r^{(n)}}{\Delta t}. \quad (22)$$

Using this approximation, we discretize the system with a backward Euler scheme:

$$\begin{bmatrix} \Delta t \mathbf{R}_{deim}(\mathbf{i}_r^{(n+1)}) + \mathbf{L}_r & \Delta t \mathbf{G}_r^T \\ \Delta t \mathbf{G}_r & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r^{(n+1)} \\ \varphi_r^{(n+1)} \end{bmatrix} = \begin{bmatrix} \mathbf{L}_r \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r^{(n)} \\ \varphi_r^{(n)} \end{bmatrix} + \begin{bmatrix} \Delta t \mathbf{e}_{s,r}^{(n+1)} \\ \mathbf{0} \end{bmatrix}, \quad (23)$$

with \mathbf{R}_{deim} is assembled using only the rows selected by the DEIM procedure. This discretization leads to a nonlinear system, since \mathbf{R}_{deim} depends nonlinearly on $\mathbf{i}_r^{(n+1)}$. Solving it requires a Newton solver, which in turn involves repeated Jacobian evaluations. This process is computationally expensive. To reduce the cost, we can, in principle, introduce an approximation: we evaluate \mathbf{R}_{deim} at the current from the previous time step $\mathbf{i}_r^{(n)}$, rather than at the unknown current $\mathbf{i}_r^{(n+1)}$:

$$\begin{bmatrix} \Delta t \mathbf{R}_{deim}(\mathbf{i}_r^{(n)}) + \mathbf{L}_r & \Delta t \mathbf{G}_r^T \\ \Delta t \mathbf{G}_r & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r^{(n+1)} \\ \varphi_r^{(n+1)} \end{bmatrix} = \begin{bmatrix} \mathbf{L}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r^{(n)} \\ \varphi_r^{(n)} \end{bmatrix} + \begin{bmatrix} \Delta t \mathbf{e}_{s,r}^{(n+1)} \\ \mathbf{0} \end{bmatrix}, \quad (24)$$

where $\mathbf{e}_{s,r}^{(n+1)} = \mathbf{e}_{s,r}(t_{n+1})$, and \mathbf{R}_{deim} again assembled from the DEIM-selected rows. This approximation linearizes the system in the unknowns, so advancing the state reduces to solving a linear system—no Newton iterations or Jacobian evaluations are required. However, care must be taken: if the problem is strongly nonlinear, this approximation may introduce large errors, since even a small error in the current can significantly affect \mathbf{R}_{deim} , and hence the computed next-state current. For the highly nonlinear problem addressed in this paper, this approach proves infeasible, as demonstrated in Section 6.7.

5 Structured Neural ODE

5.1 Model Definition

We aim at constructing a non-intrusive ROM that avoids hyperreduction of the nonlinear operator. Starting from the POD-reduced formulation, the reduced dynamics can be written in the form

$$\begin{bmatrix} \mathbf{R}_r(\mathbf{i}_r) & \mathbf{G}_r^T \\ \mathbf{G}_r & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r \\ \varphi_r \end{bmatrix} + \begin{bmatrix} \mathbf{L}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \frac{d}{dt} \begin{bmatrix} \mathbf{i}_r \\ \varphi_r \end{bmatrix} = \begin{bmatrix} \mathbf{e}_{s,r}(t) \\ \mathbf{0} \end{bmatrix}. \quad (25)$$

where the linear operators \mathbf{L}_r and \mathbf{G}_r are inherited from POD projection of the FOM, while the only remaining nonlinear contribution is the reduced resistance operator $\mathbf{R}_r(\mathbf{i}_r)$:

$$\begin{aligned} \mathbf{R}_r : \mathbb{R}^{r_i} &\longrightarrow \mathbb{R}^{r_i \times r_i}, \\ \mathbf{i}_r &\longmapsto \mathbf{R}_r(\mathbf{i}_r). \end{aligned} \quad (26)$$

Instead of approximating $\mathbf{R}_r(\mathbf{i}_r)$ via POD-DEIM, we introduce a neural-network surrogate that operates entirely in the reduced current space. Specifically, we approximate the reduced resistance operator with a feed-forward neural network

$$\widehat{\mathbf{R}}_r(\mathbf{i}_r; \psi) = \text{NN}_\psi(\mathbf{i}_r), \quad \text{NN}_\psi : \mathbb{R}^{r_i} \rightarrow \mathbb{R}^{r_i \times r_i}, \quad (27)$$

where ψ denotes the trainable parameters (weights and biases). For ease of notation, in the remainder of the manuscript, we denote this neural surrogate simply by $\mathbf{R}_\psi(\mathbf{i}_r)$. The network takes as input the reduced current degrees of freedom \mathbf{i}_r and outputs the entries of the reduced resistance operator. Substituting $\mathbf{R}_\psi(\mathbf{i}_r)$ into (25) yields a reduced dynamical system parametrized by ψ ; in this sense, the neural network defines a learnable vector field and the resulting model falls within the Neural ODE paradigm.

5.2 Training

Once the approximated operator \mathbf{R}_ψ has been defined, the next step is to train it. A naive approach would be to treat the problem as a standard regression task: one would simulate the FOM, project the resulting current snapshots onto the POD basis to obtain reduced states $\{\mathbf{i}_r^{(n)}\}$, evaluate (and, if needed, project) the corresponding nonlinear operators to form the reduced resistance matrices $\{\mathbf{R}_r(\mathbf{i}_r^{(n)})\}$, and then train the network by pointwise regression of the mapping $\mathbf{i}_r \mapsto \mathbf{R}_r(\mathbf{i}_r)$. However, this strategy is (i) intrusive, since it requires access to snapshots of the nonlinearity, which is often infeasible when dealing with commercial solvers, and (ii) less effective than training on entire trajectories. In fact, as demonstrated in [51], trajectory-based training is superior to local derivative fitting in terms of robustness and generalization capabilities.

For this reason, we adopt a *discretize-then-optimize* approach, where the time discretization is embedded into the learning process, and the parameters are updated by backpropagation through the solver. The procedure to advance the reduced system from t_n to t_{n+1} is as follows:

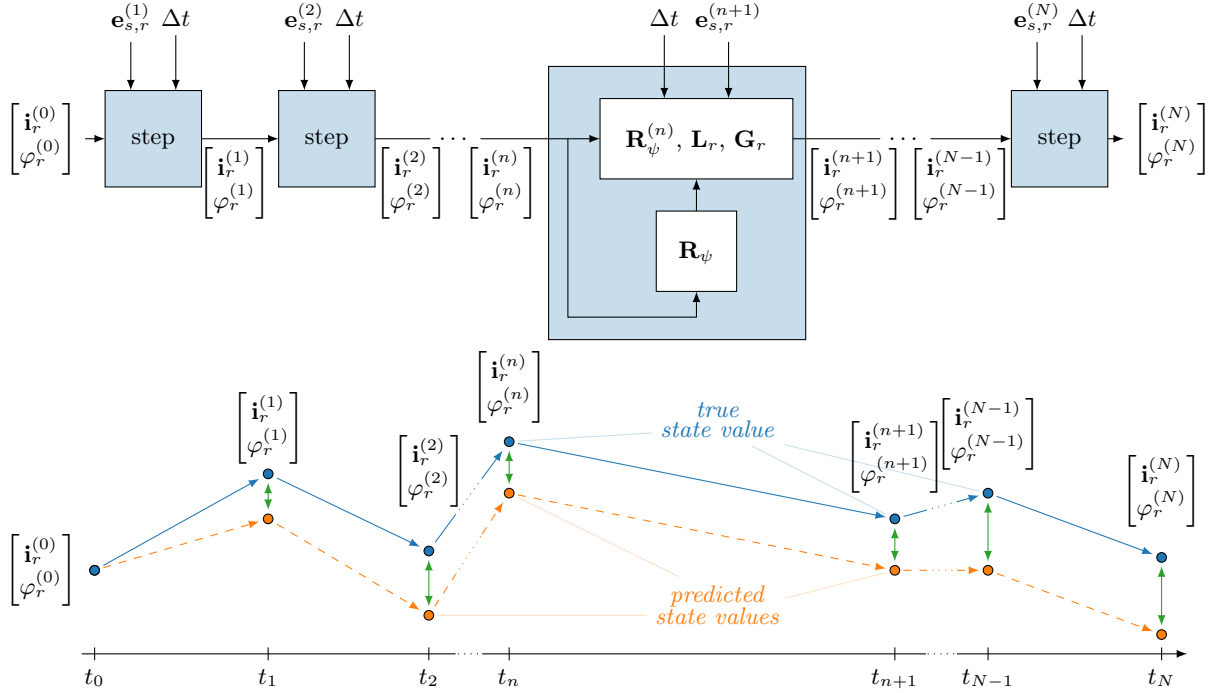


Figure 2: Graphical representation of the training process. Green arrows illustrate the discrepancy between the true state values and the predicted values, which the optimization procedure seeks to minimize.

1. *Construction of the approximated operator*: evaluate the reduced-order operator

$$\mathbf{R}_\psi^{(n)} = \mathbf{R}_\psi(\mathbf{i}_r^{(n)}), \quad (28)$$

using the reduced current at the previous time step.

2. *Time discretization*: apply the Backward Euler scheme to the dynamical system in reduced coordinates.
3. *Solution of the linear system*: solve the resulting linear system:

$$\begin{bmatrix} \Delta t \mathbf{R}_\psi^{(n)} + \mathbf{L}_r & \Delta t \mathbf{G}_r^T \\ \mathbf{G}_r & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r^{(n+1)} \\ \varphi_r^{(n+1)} \end{bmatrix} = \begin{bmatrix} \mathbf{L}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{i}_r^{(n)} \\ \varphi_r^{(n)} \end{bmatrix} + \begin{bmatrix} \mathbf{e}_{s,r}^{(n+1)} \\ \mathbf{0} \end{bmatrix}. \quad (29)$$

Since \mathbf{R}_ψ is evaluated at the previous time step, the resulting system is linear in $\mathbf{i}_r^{(n+1)}$ and can be solved without Newton iterations. In this case, considering the previous time current during training allows the network to implicitly account for the resulting lag. Therefore, unlike DEIM, this approximation error is explicitly embedded into the training process: the network learns to compensate for it, thereby reducing the mismatch and improving accuracy.

Repeating this procedure for all time steps $n = 0, \dots, N-1$, with $N\Delta t = T$, produced the time evolution of the reduced system. The predicted trajectories $\{\mathbf{i}_r^{(n)}\}_{n=0}^{N-1}$, $\{\varphi_r^{(n)}\}_{n=0}^{N-1}$ are then compared with the reference solutions obtained from the FOM, and the parameters ψ are optimized by minimizing the deviation between the reduced and full-order trajectories over the entire time horizon.

This training procedure follows the Neural ODE [33] paradigm, where the dynamics are unrolled across the temporal domain and the parameters are optimized end-to-end. The loss is defined as the Mean Squared Error (MSE) by the predicted states and the true states obtained by simulating the FOM. Figure 2 provides a graphical interpretation of this training process.

Importantly, inference must employ the same integration scheme used in training, since the learned operator is biased toward that discretization.

6 Results

This section presents the numerical results obtained for the HTS CORC[®]. We evaluate the performance of both the classical POD-DEIM approach and the proposed Structured Neural ODE model in terms of accuracy and computational efficiency.

6.1 Data Generation

The resulting model features $n_e = 4518$ edges and $n_f = 3140$ faces, yielding a total of 7658 Degrees of Freedom (DOFs). To generate the dataset used for POD-DEIM reduction and for training the Structured Neural ODE, five transient simulations were performed. Each transient corresponds to a sinusoidal excitation at 50 Hz with different amplitudes. This setup reflects the physical scenario under consideration, where the external magnetic field induced by surrounding cables is sinusoidal at 50 Hz, and its amplitude varies with both current magnitude and cable distance. To evaluate the generalization capability of the proposed approach, additional transient simulations were conducted with varying amplitudes—both within and outside the training distribution—as well as with different excitation frequencies. The corresponding excitation parameters defining the training and testing datasets are reported in Table 1.

Freq.	Ampl.	Training	Within-the-distr. val.	Outside-the-distr. val.	Different freq. val.
50 Hz	13 mT	✓			
50 Hz	15 mT	✓			
50 Hz	18 mT	✓			
50 Hz	22 mT	✓			
50 Hz	24 mT	✓			
50 Hz	20 mT		✓		
50 Hz	10 mT			✓	
50 Hz	30 mT			✓	
40 Hz	20 mT				✓
60 Hz	20 mT				✓
100 Hz	20 mT				✓

Table 1: Overview of the excitation parameters adopted for dataset generation.

Due to the pronounced nonlinearities of the model, time integration was carried out using a fully implicit Backward Euler scheme, with a Newton–Raphson solver applied at each time step. The simulations span a time window from 0 to 35 ms, discretized into 560 steps of duration 0.0625 ms. These runs are computationally expensive: the system is both high-dimensional and nonlinear, requiring many Newton iterations, each involving the solution of a large, dense linear system. All computations were performed on a 13th Gen Intel[®] Core[™] i7-13700 CPU @2.1 GHz, with 24 cores, equipped with 64 GB of RAM. On this hardware, a single transient required roughly three hours to complete, resulting in about 15 hours for the training set and an additional 18 hours for validation. These times should be understood as indicative rather than absolute, as they depend on the hardware characteristics and implementation details. Importantly, this computational effort belongs to the offline stage of the methodology and must be carried out only once. In practice, the relevant cost is that of repeatedly simulating the ROM, which is several orders of magnitude lower and therefore well suited to design and control tasks.

6.2 POD-DEIM Practical Implementation

In the implementation of the proposed reduction framework, the training snapshots were used for constructing the snapshot matrices. This procedure yielded a total of $n_s = 2800$ snapshots for the current, potential, and nonlinear term. For the current field, the dominant modes were selected based on the energy retained, resulting in $r_i = 48$ retained modes. Also for the potential field, the energy retention criterion was applied, leading to the selection of the first five modes, i.e., $r_\varphi = 5$. The singular values decay for the two variables are shown in Figure 3.

For the DEIM approximation of the nonlinear term, the number of interpolation points n_p was chosen to be equal to the number of retained modes. In this case, 150 modes were retained. The relatively large number of DEIM modes is motivated by the strong nonlinearities present in the system: accurately capturing localized variations in the current density requires a correspondingly dense set of interpolation

points. The singular values decay for the nonlinear term, with the details of retained modes, are shown in Figure 3.




Figure 3: Singular value decay of the current, potential, and nonlinear term $\mathbf{R}(\mathbf{i})\mathbf{i}$ snapshot matrices. Retained modes are shown in orange, discarded modes in blue.

6.3 Structured Neural ODE Practical Implementation

The neural network \mathbf{R}_ψ is implemented as a fully connected feed-forward architecture with $L_h = 4$ hidden layers, each consisting of 140 neurons. As a preprocessing step, the absolute value of the reduced current $|\mathbf{i}_r|$ is fed to the network, reflecting the fact that resistivity—and therefore the \mathbf{R}_ψ —depends on the magnitude of the current rather than its sign. All the layers employ *selu* [52] activation functions. The training dataset considers the very same transients used for constructing the POD-DEIM snapshot matrix. This dataset is split into fixed-length sequences that are organized into mini-batches, and the model is trained using the Adam optimizer [53].

6.4 Accuracy

The accuracy of each reduced-order model was assessed by comparing the predicted current density with the reference solution obtained from the FOM. Specifically, the instantaneous absolute error was computed as the absolute difference between the ROM-predicted current and the corresponding FOM value. In the time-dependent plots, the mean, 95th-percentile, and maximum errors were evaluated across all current states at each time instant, thereby illustrating the temporal evolution of the error. Conversely, in the summary tables, the reported mean, 95th-percentile, and maximum values were computed over all current states and all time instants of the transient simulations. Moreover, the last row of the accuracy tables reports the relative discrepancy between the predicted and reference current trajectories, measured through the Frobenius norm:

$$\text{err}_{\|\cdot\|_F} = \frac{\|\mathbf{I} - \mathbf{I}_{\text{true}}\|_F}{\|\mathbf{I}_{\text{true}}\|_F}, \quad (30)$$

where \mathbf{I} denotes the matrix collecting the full-order current states at all time instants reconstructed from the ROM solution (i.e., obtained by projecting the reduced currents back to the full space), and \mathbf{I}_{true} is defined analogously using the FOM solution. This quantity provides a global measure of relative error over the entire transient. We note that a pointwise relative error is not well defined in this context, since the current may vanish at certain spatial locations and time instants.

6.4.1 Within-the-distribution Validation

To assess the accuracy of the proposed approach within the training distribution, the models were evaluated using sinusoidal inputs whose amplitude and frequency fall inside the range of the training data ($f = 50$ Hz, $B_0 = 20$ mT). The POD-DEIM ROM was simulated with the same discretization strategy employed for the FOM, whereas the Structured Neural ODE ROM was advanced according to (29). As illustrated qualitatively at the top of Figure 4, both reduced-order models reproduce the state evolution with good accuracy.

A quantitative comparison with the FOM simulation further supports this result: the bottom part of Figure 4 shows the mean, 95th-percentile, and maximum error among the states as functions of time, while the corresponding numerical values are summarized in Table 2. These results indicate that the

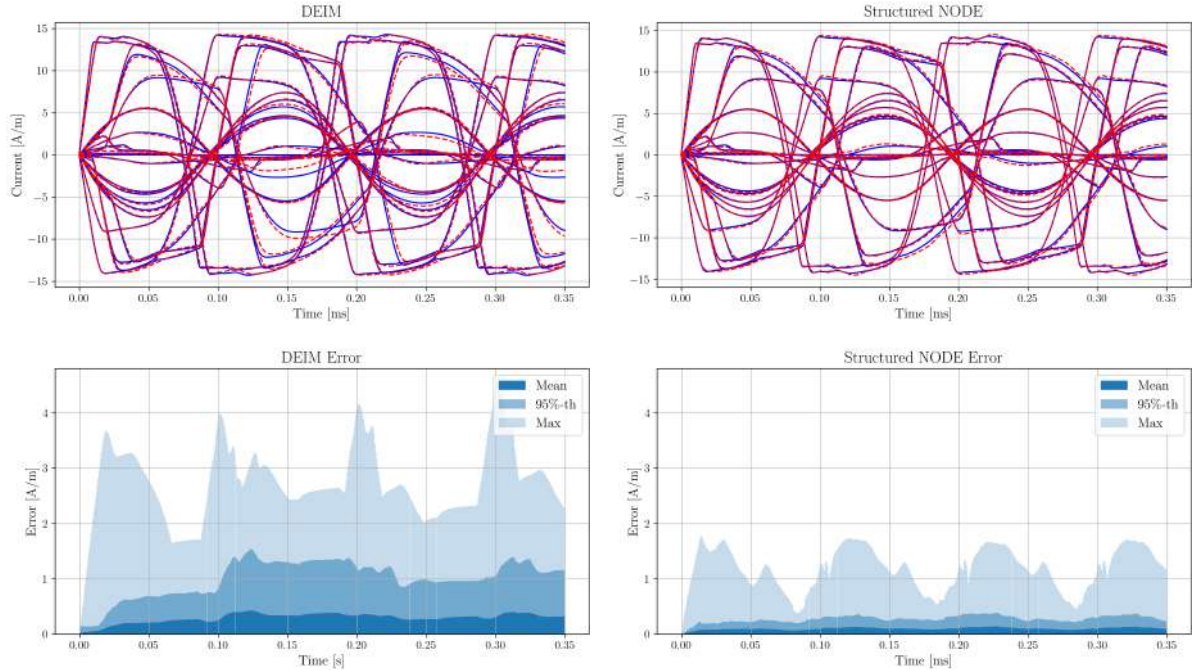


Figure 4: Top: Evolution of selected current states in the *within-the-distribution* validation transient for the ROM simulations (dashed red) compared to the FOM (solid blue). Bottom: Corresponding mean, 95th-percentile, and maximum errors between the ROMs and the FOM.

	DEIM	Neural ODE
Mean Error	0.1709 A/m	0.1164 A/m
95th-percentile	0.7921 A/m	0.3398 A/m
Max Error	4.5680 A/m	1.7775 A/m
$\text{err}_{\ \cdot\ _F}$	$7.5194 \cdot 10^{-2}$	$2.3042 \cdot 10^{-2}$

Table 2: Mean, 95-th percentile and max absolute error for the DEIM ROM and the Neural ODE ROM in the *within-the-distribution* validation transient.

errors are of the same order of magnitude, with the Structured Neural ODE yielding slightly lower values. Moreover, the current density map obtained with the proposed Neural ODE approach is consistent with that from the FOM simulation, as shown in Figure 5, which depicts the current distribution and the associated error at selected time instants.

6.4.2 *Outside-the-distribution* Validation

To evaluate the performance of the ROMs outside the training distribution, two additional transient simulations were performed at 50 Hz with excitation amplitudes of 10 mT and 30 mT—values lying below and above the training range (13–24 mT), respectively. The corresponding quantitative results, summarized in Table 3, indicate that the accuracy of the Structured Neural ODE decreases noticeably, whereas the POD-DEIM ROM maintains performance comparable to the within-distribution case.

	DEIM	Neural ODE
Mean Error	0.1021 A/m	0.4566 A/m
95th-percentile	0.3860 A/m	2.0990 A/m
Max Error	6.0196 A/m	6.7252 A/m
$\text{err}_{\ \cdot\ _F}$	$7.6155 \cdot 10^{-2}$	$1.8361 \cdot 10^{-1}$

Table 3: Mean, 95-th percentile and max absolute error for the DEIM ROM and the Neural ODE ROM in the *outside-the-distribution* validation transients.




Figure 5: Comparison between the current density map of the FOM simulation and the Neural ODE-based ROM on the *within-the-distribution* validation transient. The results demonstrate high accuracy.

This outcome is consistent with expectations: data-driven models such as the Structured Neural ODE are inherently constrained by the range of the training data. When the excitation amplitude extends beyond this range, the resulting state trajectories explore regions of the state space not represented during training, leading to a loss in predictive accuracy.

6.4.3 Different Frequency Validation

To further assess the generalization capability of the proposed approach, additional simulations were carried out at an excitation amplitude of 20 mT—within the training range—but with frequencies of 40, 60, and 100 Hz, differing from the 50 Hz training frequency. The quantitative results, summarized in Table 4, demonstrate that the Structured Neural ODE achieves very good accuracy across all tested frequencies, with errors of the same order of magnitude as in the 50 Hz validation and consistently lower than those of the POD-DEIM ROM.

	DEIM	Neural ODE
Mean Error	0.1729 A/m	0.1484 A/m
95th-percentile	0.7854 A/m	0.4595 A/m
Max Error	4.6427 A/m	2.7591 A/m
$\text{err}_{\ \cdot\ _F}$	$7.5948 \cdot 10^{-2}$	$3.6895 \cdot 10^{-2}$

Table 4: Mean, 95-th percentile and max absolute error for the DEIM ROM and the Neural ODE ROM in the *different frequency* validation transients.

This strong performance can be attributed to the intrinsic structure of the proposed model: since the Neural ODE component depends solely on the state variables, the learned dynamics remain accurate as long as the system evolution explores regions of the state space already represented in the training data. Although the model implicitly captures some temporal dependencies through the state update formulation, these do not significantly affect the accuracy, explaining the robust behavior observed across different excitation frequencies.

6.5 Computational Cost

To compare the computational cost of the two ROMs, we evaluate the number of Floating-Point Operations (FLOPs) required to advance the system by one time step. The time-stepping scheme used to simulate the POD-DEIM ROM employs Newton iterations at each step, with residual and Jacobian evaluations at each iteration—an approach necessitated by the strong nonlinearity of the problem. In both cases, an optimal implementation scenario is assumed.

6.5.1 POD-DEIM ROM Computational Cost

The computational cost of the POD-DEIM reduced-order model involves mapping the reduced solution to the full current density \mathbf{J} , evaluating the resistivity ρ and its derivative ρ' , assembling the DEIM-selected rows of \mathbf{R} and its Jacobian, performing matrix-vector products and projections, and solving the linear system within each Newton-Raphson iteration. Contributions from the right-hand side, which are independent of the solution, are computed once per time step.

6.5.2 Structure Neural ODE ROM Computational Cost

The computational cost for advancing the state in the Structured Neural ODE ROM includes evaluating the neural network, assembling the left- and right-hand sides through matrix-vector operations, and solving the resulting linear system.

6.5.3 Results of the Evaluations

For a single time step, the POD-DEIM ROM requires approximately 4.5×10^7 operations, whereas the Structured Neural ODE ROM requires only 2.6×10^5 , corresponding to a nominal FLOPs-based speed-up of $\times 176$. This estimate reflects the algorithmic cost of advancing the reduced state under the fully implicit scheme employed in the POD-DEIM ROM, which requires repeated residual and Jacobian evaluations within Newton-Raphson iterations. In our experiments, the POD-DEIM ROM required an average of

Approach	# FLOPs	Eval. Time
DEIM	45,330,701 ($\times 176$)	334 s ($\times 1500$)
Neural ODE	257,395 ($\times 1$)	0.2 s ($\times 1$)

Table 5: Number of FLOPs per time step and total evaluation time for reconstructing a full transient.

19 Newton iterations per time-step, and the dominant cost arises from repeated assembly of the DEIM-reconstructed nonlinear term and its Jacobian. By contrast, the Structured Neural ODE ROM advances the reduced state by evaluating the neural-network once per step and solving a single reduced linear system, thereby avoiding Newton iterations and the associated repeated nonlinear assembly.

If, instead of considering operation counts, we measure computational cost in terms of wall-clock time for the full transient simulation, the observed speed is significantly larger. While the Structured Neural ODE can be implemented in a highly optimized manner using only efficient neural-network evaluation and basic algebraic operations, the POD-DEIM ROM requires repeated operator assembly, and these assembly routines are more difficult to optimize in practice. As a result, the overhead associated with repeated assembly and Newton iterations dominates the execution time of the POD-DEIM ROM. In our implementation, this leads to a practical speed-up of approximately $\times 1500$, reducing the runtime for a full transient simulation from 5 minutes to just 0.2 seconds. These results are summarized in Table 5.

6.6 Time-step sensitivity analysis

We additionally investigate the effect of the time-step size Δt on accuracy and computational cost. The Structured Neural ODE, trained at Δt_{train} , remains accurate when deployed with smaller time steps, whereas accuracy degrades for $\Delta t > \Delta t_{\text{train}}$, consistently with [30]. The fully implicit POD-DEIM scheme is, in principle, less sensitive to Δt , but in our experiments, Newton iterations fail to converge beyond a certain Δt , limiting the admissible time-step range unless additional solver enhancements are introduced (e.g., problem-specific preconditioning strategies and improved initial guesses), which bring non-trivial algorithmic complexity. These trends are summarized in Fig. 6.

In terms of computational cost, the Structured Neural ODE exhibits an approximately constant per-step complexity (dominated by a single neural-network evaluation and reduced linear algebra), so the total transient evaluation time scales linearly with the number of time steps, i.e., proportional to $1/\Delta t$. The computational cost of the fully implicit POD-DEIM approach instead depends both on the number of time steps and on the number of Newton iterations per step. Moreover, as discussed in Section 6.7, the per-step cost of POD-DEIM with an explicit time integrator is already comparable to that of the Structured Neural ODE. Therefore, even in the most favorable implicit scenario (i.e., a single Newton iteration), one should expect a higher per-step cost due to the additional assembly of the Jacobian on top of the residual evaluation. Consequently, across the tested range of Δt , the proposed method retains a substantial computational advantage while maintaining competitive accuracy for $\Delta t \leq \Delta t_{\text{train}}$.

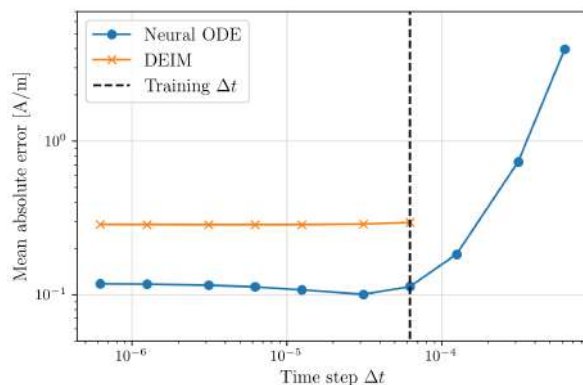


Figure 6: Time-step sensitivity study. Each marker reports the mean absolute current error averaged over space and time with respect to the FOM reference solution, so that each point summarizes the discrepancy over the full transient. The dashed line denotes the training time step Δt_{train} . The DEIM curve terminates at the largest time step for which the Newton solver converges.

6.7 Note on the Time-stepping Strategy

As explained in Section 4.3, the DEIM approach could also be implemented by evaluating \mathbf{R}_r at the previously computed state, in which case only a linear system would need to be solved at each time step. This would reduce the computational effort to a level comparable with that of our Neural ODE. In particular, if we consider this approximation, the DEIM requires approximately 1.99×10^5 FLOPs to advance the state, which is about 1.29 times fewer operations than the Structured Neural ODE. However, a crucial difference lies in the fact that the neural ODE has been explicitly trained to perform this task, and the training process inherently accounts for this simplification, thereby ensuring accurate results. By contrast, due to the strong nonlinearity of the problem, applying the same type of approximation within the POD-DEIM ROM leads to unreliable and ultimately uninformative simulations, as illustrated in Figure 7.




Figure 7: Evolution of selected current states in the validation transient for the POD-DEIM ROM (dashed red) compared with the FOM (solid blue). The corresponding mean, 95th-percentile, and maximum errors between the ROM and the FOM are also reported. In these simulations, the nonlinear operator is evaluated using the current state from the previous iteration. The results clearly show that this approach does not yield accurate or practically useful solutions.

6.8 Post-processing: Power losses

As a post-processing step, we evaluate the AC losses in the CORC[®] cable, which serve as input for a thermal model aimed at predicting the cable’s temperature distribution—a key quantity for quench prediction in both design and control. The AC losses are computed as:

$$p(t) = \int_{\Omega_{sc}} \mathbf{J} \cdot \mathbf{E} d\Omega = \int_{\Omega_{sc}} \rho \|\mathbf{J}\|^2 d\Omega. \quad (31)$$

These losses are related to the current loops in the HTS tapes, also recognized as magnetization losses [54]. The integral is performed numerically, assuming the electric field and current density are constant within each element. Power losses are evaluated using current densities obtained from both the POD-DEIM ROM and the Neural ODE ROM. Figure 8 shows the results of these evaluations, along with the error relative to the FOM values. The comparison indicates that both ROMs achieve comparable accuracy in predicting power loss.

7 Discussion

The comparison between the Structured Neural ODE and the POD-DEIM approaches highlights a clear trade-off between accuracy, computational efficiency, intrusiveness, and training requirements.

Both ROMs achieve comparable overall accuracy, with the Structured Neural ODE generally providing slightly higher precision. In particular, the Structured Neural ODE yields the most accurate results for inputs within the training distribution, while the POD-DEIM ROM proves more reliable when the excitation amplitude lies outside the training range. For excitations at different frequencies, the Structured Neural ODE performs remarkably well, maintaining accuracy comparable to the training frequency case and outperforming the POD-DEIM model. This behavior confirms that the proposed data-driven structure generalizes effectively across variations in frequency, provided that the system trajectories remain within regions of the state space represented during training.

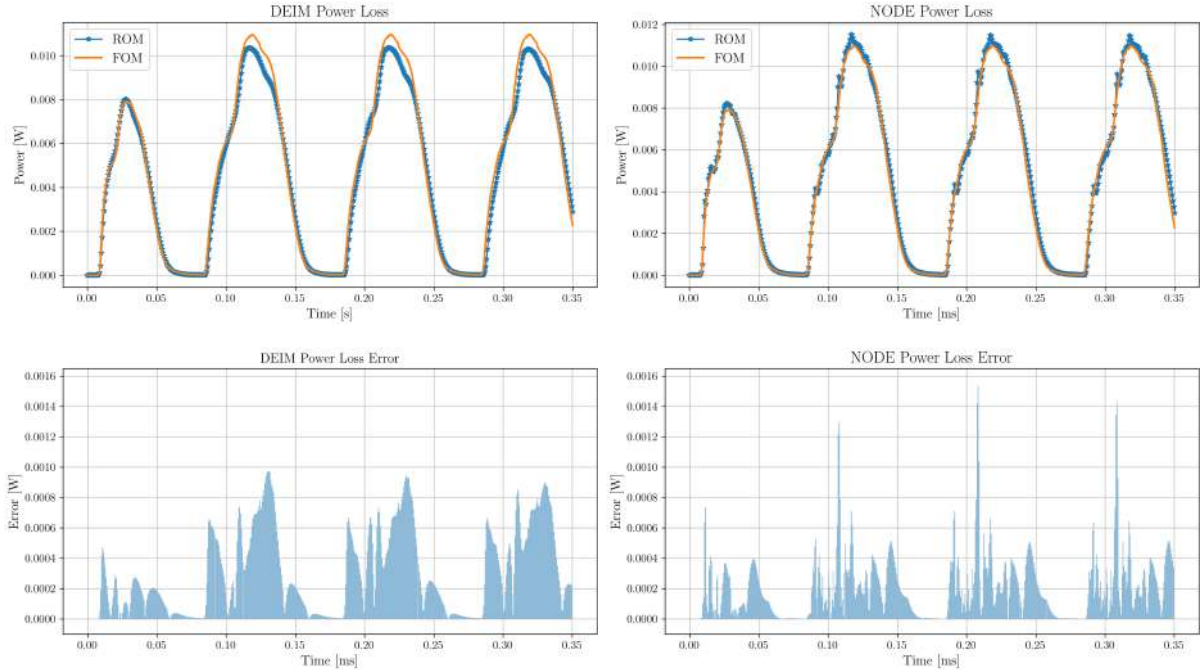


Figure 8: Comparison of power loss of the *within-the-distribution* validation transient computed from the current density obtained using the FOM and the two ROMs, including the associated error.

In terms of computational cost, the Structured Neural ODE is significantly faster because it advances the state by solving only a linear system at each time step, whereas the POD-DEIM ROM requires a high number of Newton-Raphson iterations due to the strong nonlinearity of the model, with repeated evaluations of the Jacobian at each step.

The Structured Neural ODE is also considerably less intrusive, as it only requires access to the assembled linear operators and to evaluations of the nonlinear operator at selected working points. Such information can typically be extracted without modifying the underlying simulation code, making the approach compatible with commercial solvers and straightforward to integrate into existing workflows. In contrast, the POD-DEIM method is highly intrusive, since it requires direct access to and modification of the FOM assembly routines to incorporate the DEIM interpolation points. This need for code-level intervention limits its applicability in commercial or legacy software environments and makes the optimization of the DEIM approximation more demanding in practice.

When it comes to training, POD-DEIM is easier to set up and tune due to its linear algebra-based structure, whereas Structured Neural ODEs demand careful tuning of hyperparameters and more complex training procedures.

Finally, data requirements differ: POD-DEIM typically needs less data and generalizes well across operating conditions, while Structured Neural ODEs require larger, representative datasets to ensure reliable performance during operation. A concise summary of these pros and cons is provided in Table 6.

Taken together, these results indicate that the Structured Neural ODE offers clear advantages in terms of intrusiveness and online efficiency, particularly in strongly nonlinear regimes where Newton iterations dominate the runtime. POD-DEIM remains preferable primarily in data-scarce or extrapolative scenarios, whereas the Structured Neural ODE is the method of choice whenever sufficient training data are available.

Aspect	Neural ODE ROM	POD-DEIM ROM
Accuracy	High (overall, slightly higher)	High
Cost	Low ($\sim 176\times$ speedup)	High (Newton Iterations; Jacobian Evaluations)
Data requirements	High (Dense Dataset)	Low (Few snapshots are enough)
Training procedure	Empirical and delicate	Straightforward (Linear Algebra-based)
Intrusiveness	Non-intrusive	Intrusive

Table 6: Detailed comparison between Neural ODE-based ROM and POD-DEIM ROM.

8 Conclusion

In this work, we developed and assessed reduced-order modeling strategies for nonlinear EM simulations of HTS conductors within the IEM framework. Specifically, we presented the first application of POD-DEIM to IEM-based HTS models, and proposed a Structured Neural ODE approach that learns the nonlinear dynamics directly in the reduced-order space.

These methods were pursued to enable real-time prediction and monitoring of quench phenomena in superconducting devices. While conventional FEM or IEM remain too costly for online use, the proposed reduced-order strategies significantly improve computational efficiency without sacrificing accuracy. In particular, the Structured Neural ODE demonstrated superior predictive capability compared to the standard POD-DEIM approach, achieving faster evaluations while maintaining high fidelity to the FOM.

Future work will extend this framework toward multiphysics simulations, where the computed AC losses from the electromagnetic model will serve as inputs to thermal solvers. This coupling will allow for a more comprehensive description of quench initiation and propagation, thereby bringing reduced-order modeling closer to practical deployment in the design and operation of next-generation superconducting systems.

References

- [1] Y. Zhai, T. Brown, J. E. Menard, D. C. van der Laan, J. D. Weiss, and Z. Johnson. HTS Cable Conductor for Compact Fusion Tokamak Solenoids. *IEEE Transactions on Applied Superconductivity*, 32(6):1–5, 2022. doi: 10.1109/TASC.2022.3167343.
- [2] Lucio Rossi. Particle Accelerators and Cuprate Superconductors. *Physica C: Superconductivity and its Applications*, 614:1354360, 2023. ISSN 0921-4534. doi: doi.org/10.1016/j.physc.2023.1354360.
- [3] Makoto Takayasu, Luisa Chiesa, Nathaniel C. Allen, and Joseph V. Minervini. Present Status and Recent Developments of the Twisted Stacked-Tape Cable Conductor. *IEEE Transactions on Applied Superconductivity*, 26(2):25–34, 2016. doi: 10.1109/TASC.2016.2521827.
- [4] Wilfried Goldacker, Francesco Grilli, Enric Pardo, Anna Kario, Sonja I Schlachter, and Michal Vojenčiak. Roebel cables from REBCO coated conductors: a one-century-old concept for the superconductivity of the future. *Superconductor Science and Technology*, 27(9):093001, aug 2014. doi: 10.1088/0953-2048/27/9/093001.
- [5] D C van der Laan, J D Weiss, and D M McRae. Status of CORC® cables and wires for use in high-field magnets and power systems a decade after their introduction. *Superconductor Science and Technology*, 32(3):033001, feb 2019. doi: 10.1088/1361-6668/aafc82.
- [6] Rui Kang, Juan Wang, and Qingjin Xu. Detecting Quench in HTS Magnets With LTS Wires—A Theoretical and Numerical Analysis. *IEEE Transactions on Applied Superconductivity*, 32(6):1–5, 2022. doi: 10.1109/TASC.2022.3171185.
- [7] Maxim Marchevsky. Quench Detection and Protection for High-Temperature Superconductor Accelerator Magnets. *Instruments*, 5(3), 2021. ISSN 2410-390X. doi: 10.3390/instruments5030027.
- [8] Brad J. Merrill. Modeling an unmitigated thermal quench event in a large field magnet in a DEMO reactor. *Fusion Engineering and Design*, 98-99:2196–2200, 2015. ISSN 0920-3796. doi: doi.org/10.1016/j.fusengdes.2015.03.007. Proceedings of the 28th Symposium On Fusion Technology (SOFT-28).
- [9] A Zappatore. Full 3D thermal-hydraulic and electric modelling of quench propagation in HTS conductors. *Superconductor Science and Technology*, 37(12):125012, nov 2024. doi: 10.1088/1361-6668/ad8e8a.
- [10] Boyang Shen, Francesco Grilli, and Tim Coombs. Overview of H-Formulation: A Versatile Tool for Modeling Electromagnetics in High-Temperature Superconductor Applications. *IEEE Access*, 8: 100403–100414, 2020. doi: 10.1109/ACCESS.2020.2996177.
- [11] Felix Huber, Wenjuan Song, Min Zhang, and Francesco Grilli. The T-A formulation: an efficient approach to model the macroscopic electromagnetic behaviour of HTS coated conductor applications. *Superconductor Science and Technology*, 35(4):043003, mar 2022. doi: 10.1088/1361-6668/ac5163.

- [12] Francesco Lucchini, Riccardo Torchio, Antonio Morandi, and Fabrizio Dughiero. A fast integral equation $j - \varphi$ formulation for superconducting structures. *IEEE Transactions on Applied Superconductivity*, 34(4):1–8, 2024. doi: 10.1109/TASC.2024.3366189.
- [13] Francesco Lucchini. Evaluating Magnetization Losses in 3-D CORC Tapes With Integral and Finite-Element Methods. *IEEE Transactions on Applied Superconductivity*, 35(3):1–8, 2025. doi: 10.1109/TASC.2025.3544512.
- [14] Yizhao Huang, Mohammad Yazdani-Asrami, and Wenjuan Song. A Deep Learning-Based Ultra-Fast Surrogate Model for AC Loss Estimation in Superconductors Using COMSOL. *IEEE Access*, 13: 153845–153854, 2025. doi: 10.1109/ACCESS.2025.3603953.
- [15] Linjie Zhou, Yihan Wang, Qi Yuan, Xiaowei Song, Liang Li, and Qiuliang Wang. AC Loss Calculation of High Temperature Superconducting Coils Based on a Surrogate Model. *IEEE Transactions on Applied Superconductivity*, 35(5):1–5, 2025. doi: 10.1109/TASC.2025.3561097.
- [16] Stefano Sorti, Lorenzo Balconi, Lucio Rossi, Carlo Santini, and Marco Statera. Toward Real-Time Electromagnetic Simulations of HTS Non-Insulated Coils Through Proper Orthogonal Decomposition. *IEEE Transactions on Applied Superconductivity*, 35(5):1–5, 2025. doi: 10.1109/TASC.2025.3526741.
- [17] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. An ‘Empirical Interpolation’ Method: Application to Efficient Reduced-Basis Discretization of Partial Differential Equations. *Comptes Rendus. Mathématique*, 339(9):667–672, Nov. 2004. doi: 10.1016/j.crma.2004.08.006.
- [18] Saifon Chaturantabut and Danny C. Sorensen. Discrete Empirical Interpolation for nonlinear model reduction. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 4316–4321, 2009. doi: 10.1109/CDC.2009.5400045.
- [19] Ze Guo, Zuqi Tang, and Zhuoxiang Ren. Tensor decomposition-based DEIM for model order reduction applied to nonlinear parametric electromagnetic problems. *Journal of Computational Physics*, 542:114352, 2025. ISSN 0021-9991. doi: doi.org/10.1016/j.jcp.2025.114352.
- [20] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural Operator: Learning Maps Between Function Spaces With Applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- [21] Peter Benner, Serkan Gugercin, and Karen Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review*, 57(4):483–531, 2015. doi: 10.1137/130932715.
- [22] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2 edition, 2022.
- [23] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, March 2021. ISSN 2522-5839. doi: 10.1038/s42256-021-00302-5.
- [24] Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier Neural Operator for Parametric Partial Differential Equations. In *International Conference on Learning Representations*, 2021.
- [25] Anima Anandkumar, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Nikola Kovachki, Zongyi Li, Burigede Liu, and Andrew Stuart. Neural operator: Graph kernel network for partial differential equations. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2019.
- [26] 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*, 378:686–707, 2019. ISSN 0021-9991. doi: doi.org/10.1016/j.jcp.2018.10.045.

- [27] B. Peherstorfer and K. Willcox. Data-driven Operator Inference for Nonintrusive Projection-based Model Reduction. *Computer Methods in Applied Mechanics and Engineering*, 306:196–215, Jul. 2016. doi: 10.1016/j.cma.2016.03.025.
- [28] Peter Benner, Pawan Goyal, Boris Kramer, Benjamin Peherstorfer, and Karen Willcox. Operator inference for non-intrusive model reduction of systems with non-polynomial nonlinear terms. *Computer Methods in Applied Mechanics and Engineering*, 372:113433, 2020. ISSN 0045-7825. doi: doi.org/10.1016/j.cma.2020.113433.
- [29] Christophe Bonneville, Xiaolong He, April Tran, Jun Sur Park, William Fries, Daniel A. Messenger, Siu Wun Cheung, Yeonjong Shin, David M. Bortz, Debojyoti Ghosh, Jiun-Shyan Chen, Jonathan Belof, and Youngsoo Choi. A comprehensive review of latent space dynamics identification algorithms for intrusive and non-intrusive reduced-order-modeling. *arXiv preprint arXiv:2403.10748*, 2024. URL <https://arxiv.org/abs/2403.10748>.
- [30] Nicola Farenga, Stefania Fresca, Simone Brivio, and Andrea Manzoni. On latent dynamics learning in nonlinear reduced order modeling. *Neural Networks*, 185:107146, 2025. ISSN 0893-6080. doi: doi.org/10.1016/j.neunet.2025.107146.
- [31] B. Lusch, J.N. Kutz, and S.L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, Apr. 2018. doi: 0.1038/s41467-018-07210-0.
- [32] Boris Kramer and Karen E. Willcox. Nonlinear model order reduction via lifting transformations and proper orthogonal decomposition. *AIAA Journal*, 57(6):2297–2307, Jun. 2019. doi: 10.2514/1.J057791.
- [33] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [34] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.
- [35] Xinyu Pan, Dunhui Xiao, Lihua Wang, and Xiaoquan Yang. A physics-data combined neural network-based finite volume parametric reduced order model. *Physics of Fluids*, 37(12):123602, 12 2025. ISSN 1070-6631. doi: 10.1063/5.0308144.
- [36] Jiaming Guo and Dunhui Xiao. Nonlinear model reduction by probabilistic manifold decomposition. *SIAM Journal on Scientific Computing*, 48(1):A209–A235, 2026. doi: 10.1137/25M1738863.
- [37] R. Fu, D. Xiao, I.M. Navon, F. Fang, L. Yang, C. Wang, and S. Cheng. A non-linear non-intrusive reduced order model of fluid flow by auto-encoder and self-attention deep learning methods. *International Journal for Numerical Methods in Engineering*, 124(13):3087–3111, 2023. doi: <https://doi.org/10.1002/nme.7240>.
- [38] Kexing Li, Ke Li, Liang Zou, Yutong Fu, Longhao Yang, Weihang Peng, and Yawei Wang. Over-current Quench Detection of Parallel-Wound No-Insulation High Temperature Superconductor Coil Based on Digital Twin Method. *IEEE Transactions on Applied Superconductivity*, 34(5):1–6, 2024. doi: 10.1109/TASC.2024.3357445.
- [39] Pai Peng, Yutong Fu, Weihang Peng, and Yawei Wang. A Quench Behavior Predictive Model for High Temperature Superconducting Magnet Based on Deep-Learning Neural Network. *IEEE Transactions on Applied Superconductivity*, 35(5):1–6, 2025. doi: 10.1109/TASC.2025.3543793.
- [40] Alec J. Linot, Joshua W. Burby, Qi Tang, Prasanna Balaprakash, Michael D. Graham, and Romit Maulik. Stabilized neural ordinary differential equations for long-time forecasting of dynamical systems. *Journal of Computational Physics*, 474:111838, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2022.111838>.
- [41] Daniel A. Serino, Allen Alvarez Loya, J.W. Burby, Ioannis G. Kevrekidis, and Qi Tang. Fast-slow neural networks for learning singularly perturbed dynamical systems. *Journal of Computational Physics*, 537:114090, 2025. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2025.114090>.

- [42] Allen Alvarez Loya, Daniel A. Serino, J. W. Burby, and Qi Tang. Structure-preserving neural ordinary differential equations for stiff systems. *arXiv preprint arXiv:2503.01775*, 2025. URL <https://arxiv.org/abs/2503.01775>.
- [43] Boyang Shen, Xiaoyuan Chen, Lin Fu, Shan Jiang, Huayu Gou, Jie Sheng, Zhen Huang, Wei Wang, Yujia Zhai, Yupeng Yuan, Shuo Gao, Sheng Wang, Chuanyue Li, Wei Chen, Shuai Mu, Qian Zhou, Wenjuan Song, Xiaoze Pei, Yavuz Öztürk, James Gawith, Yi Liu, Ismail Patel, Mengyuan Tian, Jiabin Yang, and Tim Coombs. Superconducting Conductor on Round Core (CORC) Cables: 2D or 3D Modeling? *IEEE Transactions on Applied Superconductivity*, 31(8):1–5, 2021. doi: 10.1109/TASC.2021.3091091.
- [44] Yawei Wang, Min Zhang, Francesco Grilli, Zixuan Zhu, and Weijia Yuan. Study of the magnetization loss of CORC® cables using a 3D T-A formulation. *Superconductor Science and Technology*, 32(2):025003, jan 2019. doi: 10.1088/1361-6668/aaf011.
- [45] M. U. Fareed, M. Kapolka, B. C. Robert, M. Clegg, and H. S. Ruiz. 3D FEM Modeling of CORC Commercial Cables With Bean’s Like Magnetization Currents and Its AC-Losses Behavior. *IEEE Transactions on Applied Superconductivity*, 32(4):1–5, 2022. doi: 10.1109/TASC.2022.3145309.
- [46] B. H. Wu, S. Y. Gao, X.S. Yang, J. Jiang, H. Zhang, and Y. Zhao. Magnetization Losses in HTS Coil With Different Winding Method. *IEEE Transactions on Applied Superconductivity*, 34(8):1–4, 2024. doi: 10.1109/TASC.2024.3425323.
- [47] Hasan N. H. Al-Ssalih, Matthew Clegg, Antonio Badía-Majós, and Harold S. Ruiz. Magnetization Losses and Non-Reduced 3D Modeling of Hybrid CORC-TSTC Composite Conductors. *IEEE Access*, 13:186952–186964, 2025. doi: 10.1109/ACCESS.2025.3627080.
- [48] Frédéric Sirois, Francesco Grilli, and Antonio Morandi. Comparison of constitutive laws for modeling high-temperature superconductors. *IEEE Transactions on Applied Superconductivity*, 29(1):1–10, 2019. doi: 10.1109/TASC.2018.2848219.
- [49] Anindya Chatterjee. An introduction to the proper orthogonal decomposition. *Current Science*, 78(7):808–817, 2000.
- [50] Riccardo Torchio, Sebastian Schops, and Francesco Lucchini. Block Structure Preserving Model Order Reduction for A-EFIE Integral Equation Method. In *2024 IEEE International Symposium on Antennas and Propagation and INC/USNC-URSI Radio Science Meeting (AP-S/INC-USNC-URSI)*, pages 977–978, 2024. doi: 10.1109/AP-S/INC-USNC-URSI52054.2024.10686274.
- [51] Hugo Melchers, Daan Crommelin, Barry Koren, Vlado Menkovski, and Benjamin Sanderse. Comparison of neural closure models for discretised PDEs. *Computers & Mathematics with Applications*, 143:94–107, 2023. ISSN 0898-1221. doi: doi.org/10.1016/j.camwa.2023.04.030.
- [52] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-Normalizing Neural Networks. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [53] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2017. URL <https://arxiv.org/abs/1412.6980>.
- [54] Mohammad Yazdani-Asrami, S Asghar Gholamian, Seyyed Mehdi Mirimani, and Jafar Adabi. Calculation of AC magnetizing loss of ReBCO superconducting tapes subjected to applied distorted magnetic fields. *Journal of Superconductivity and Novel Magnetism*, 31(12):3875–3888, 2018.