

Gradient-Informed Machine Learning in Electromagnetics

Matteo Zorzetto¹, Merle Backmeyer², Michael Wiesheu²,
Riccardo Torchio^{1,3}, Fabrizio Dughiero¹, and Sebastian Schöps²

¹Department of Industrial Engineering, University of Padova, 35131 Padova, Italy

²Computational Electromagnetics Group, Technische Universität Darmstadt, 64289 Darmstadt, Germany

³Department of Information Engineering, University of Padova, 35131 Padova, Italy

Simulation techniques such as the finite element method are essential for designing electrical devices, but their computational cost can be prohibitive for repeated or real-time computations. Projection-based model order reduction techniques mitigate this by reducing the model size and complexity, yet face challenges when extended to nonlinear or non-affine parametric models. In this work, Isogeometric Analysis (IGA) is combined with proper orthogonal decomposition and Gaussian process regression to construct a non-intrusive surrogate model of a parametric nonlinear model of a permanent magnet synchronous machine. The differentiable nature of IGA allows for computationally efficient extraction of parametric sensitivities, which are leveraged for gradient-enhanced surrogate modeling.

***Index Terms*—Gaussian Process Regression, Proper Orthogonal Decomposition, Isogeometric Analysis, Electromagnetic Modeling.**

I. INTRODUCTION

SIMULATION techniques offer powerful tools for analyzing, designing, and monitoring electrical devices. Methods such as the finite element method convert physical equations into large systems of algebraic equations, which can then be solved using numerical linear algebra. However, the computational cost associated with solving these systems can become a limiting factor when repeated computations are required, such as in system-level simulations, design optimization, or real-time control. Projection-based Model Order Reduction (MOR) techniques address this problem by projecting the system matrices onto a lower-dimensional subspace, obtaining a computationally efficient Reduced-Order Model (ROM). Approaches such as balanced truncation, moment matching, and Proper Orthogonal Decomposition (POD) [1], [2] are particularly effective for linear systems and differ by how they find the reduced basis. However, their extension to nonlinear and parametric models is challenging and intrusive [1]. Algorithms such as the discrete empirical interpolation method can treat nonlinear terms [3], [4], [5], but their application requires access to the nonlinear representation of the system, which is not guaranteed, especially when commercial software is involved [6]. Moreover, although projection-based methods can be extended to handle affine parametric dependencies, producing a parametric ROM that preserves the structure of the original model, certain types of dependencies remain difficult to accommodate [1]. In particular, geometric parameter variations often violate the assumptions required for an affine decomposition, making them poorly suited for standard projection-based MOR frameworks. In recent years, machine learning techniques have been increasingly employed to overcome these challenges by learning the relationship between system parameters and quantities of interest directly from data. In this framework, the high-dimensional field variables can be decomposed using methods such as POD [7], [8] or autoen-

coders [9], [10], resulting in a compact set of latent variables that capture the dominant dynamics of the system. Machine learning models such as neural networks [11], radial basis function interpolation [7], and Gaussian process regression (GPR) [6], [9], [8], can then be trained to map the system parameters to this reduced space, enabling the reconstruction of the full-order field with significantly lower computational cost. In this work, building upon [12], we propose a Gradient-Enhanced (GE) surrogate modeling framework that combines Isogeometric Analysis (IGA) with POD and GPR (aka Kriging), extending the previously proposed Gradient-Free (GF) approach [11]. Leveraging the exact geometric parametrization of IGA, the proposed approach enables the non-intrusive and computationally efficient extraction of the state vector sensitivities with respect to geometric parameters, quantities not readily accessible in finite element method formulations. This gradient information can be incorporated into the reduced basis to construct a more expressive subspace for representing the state vector. Moreover, it can be used within gradient-enhanced interpolation schemes, such as GPR, to improve the predictive accuracy of the surrogate in the latent space with respect to gradient-free approaches [12]. The resulting method is capable of generating accurate and efficient surrogate models of electrical devices with nonlinear material behavior and geometric parameter variations in a non-intrusive manner. The effectiveness of the approach is demonstrated on a Permanent-Magnet Synchronous Machine (PMSM) through two studies. First, POD is combined with gradient-enhanced GPR to construct a surrogate model of the electromagnetic field distribution. Moreover, the derivative information is directly leveraged to build surrogate models for Key Performance Indicators (KPIs) such as torque.

II. MOTOR MODEL

The method is applied to a PMSM from [13], [14], shown in Figure 1, function of the geometric design variables

Corresponding author: M. Zorzetto




Fig. 1: Parametrization of the PMSM geometry, based on [14].

$\mathbf{p} = [\text{MH}, \text{MW}, \text{MAG}, \theta_1]$, and solved for a fixed number of rotation angles α . The iron cores are shown in gray, the (homogenized) copper slots in red, the rotor magnet in green, and the air gap and air pockets in blue. The physical model is described in detail in [14], while the parameterization is discussed in [11]. For brevity, only the key aspects are outlined here.

A. Governing Equations

For the laminated PMSM, the parameter-dependent magnetic vector potential in the rotor and stator, $A_{z,R}$ and $A_{z,S}$ respectively, is obtained from the two-dimensional magneto-static Poisson problem [15], [16]

$$\nabla \cdot (\nu \nabla A_{z,R}^{(\mathbf{p})}) = \nu \nabla \cdot \mathbf{B}_{\text{rem}}^{\perp}(\mathbf{p}) \quad \text{in } \Omega_R^{(\mathbf{p})}, \quad (1)$$

$$\nabla \cdot (\nu \nabla A_{z,S}^{(\mathbf{p})}) = -J_{z,\text{src}} \quad \text{in } \Omega_S^{(\mathbf{p})}, \quad (2)$$

driven by the source current density $J_{z,\text{src}}$. The 2D remanent flux density of the permanent magnets $\mathbf{B}_{\text{rem}}^{\perp}(\mathbf{p}) = B_{\text{rem}} [-\sin(\alpha), \cos(\alpha)]^T$ is computed from the remanence magnitude B_{rem} and the rotation angle α . It may also depend on the parameter vector. The reluctivity ν is set linear for air, copper, and the magnet. For the iron parts, we have $\nu = \nu(B)$ given through a nonlinear B–H curve. The magnetic flux density is given by $\mathbf{B}^{(\mathbf{p})} = \nabla \times (A_z^{(\mathbf{p})} \mathbf{e}_z)$. Problem (1–2) is complemented with standard boundary conditions, and the coupling across $\Gamma_{\text{ag}} = \Omega_R^{(\mathbf{p})} \cap \Omega_S^{(\mathbf{p})}$ is enforced by a Lagrange multiplier and additional coupling conditions (see [17]). The electromagnetic torque is computed from the field solution; thus it also depends on the parametric realization of the problem [15].

B. IGA Discretization

The discretization of the geometry as well as the solution is performed using Non-Uniform Rational B-Spline basis functions, following the principles of IGA [18]. The physical domain is described by a multi-patch parametrization, decomposing it into a collection of subdomains, each with a corresponding projection map to the same reference domain [19]. Despite different geometrical realizations, the reference domain stays the same, which is essential if the parametric problem is learned in terms of its coefficients.

For reasonable parametric changes, no remeshing is required, and the learned coefficients belong to the same nodes in the reference domain which are then mapped to the corresponding physical domain [11]. Problem (1–2) is discretized with B-splines following the standard Ritz-Galerkin approach. The magnetic vector potential in the reference domain is approximated by

$$A_z^{(\mathbf{p})}(\mathbf{x}) \approx \sum_{i=1}^N B_i^q(\mathbf{p}, \mathbf{x}) u_i \quad (3)$$

where u_i are the unknown coefficients and $B_i^q(\mathbf{p}, \mathbf{x})$ are the spline basis functions of degree q . The physical problem results in a nonlinear system of equations dependent on the geometrical parameters $\mathbf{p}$:

$$\mathbf{A}(\mathbf{p}, \mathbf{u}) \mathbf{u} = \mathbf{b}, \quad (4)$$

where $\mathbf{u}$ is the state vector, including variables for the rotor and stator, as well as Lagrange multipliers for rotor-stator coupling. $\mathbf{A}(\mathbf{p}, \mathbf{u})$ is assembled with stiffness matrices and coupling terms that depend on geometrical parameters and the state vector due to the presence of nonlinear materials. The vector $\mathbf{b}$ describes the excitations of rotor and stator that are given by the permanent magnet and the coils. For more details on constructing the matrix system (4) the reader is referred to [16]. As outlined in [16], analytical computation of the gradients is feasible through the adjoint method. Our implementation of IGA [20] not only provides the solution vector $\mathbf{u}(\mathbf{p}_*)$ but also enables efficient evaluation of its sensitivities with respect to each component of $\mathbf{p}$, as briefly discussed in the following.

C. Calculation of Derivatives

The sensitivities with respect to the parameters are obtained by applying the adjoint method. After solving (4) iteratively with the Newton-Raphson scheme, the total derivative of the torque T is obtained by

$$\frac{dT}{d\mathbf{p}} = \mathbf{z}^T \left(\frac{\partial(\mathbf{A}\mathbf{u})}{\partial\mathbf{p}} - \frac{\partial\mathbf{b}}{\partial\mathbf{p}} \right) \quad (5)$$

introducing the adjoint variable $\mathbf{z}$ with

$$\mathbf{J}^T \mathbf{z} = -\frac{\partial T}{\partial \mathbf{u}}, \quad (6)$$

where $\mathbf{J}$ is the Jacobian matrix of (4) after the Newton solve. This makes the evaluation of gradients very efficient, as even if the system is nonlinear, the adjoint approach just requires one linear system solve to obtain the gradient.

This leaves only the calculation of partial derivatives needed to evaluate (5) and (6). The partial derivative with respect to the torque needed in (6) is simple, because the torque is a closed-form expression directly depending on $\mathbf{u}$, see [16, Eq. 30].

The partial derivatives in (5) are more involved, as the changes in system stiffness matrix and right-hand side vector with respect to parameters need to be derived. For standard Finite Elements, this might prove difficult, since a new geometry usually involves remeshing, making it impossible to track the contributions in the system matrix and vector. This is different in the IGA case, where the geometry mapping is predefined

by patches and their corresponding control points. There are efficient analytical expressions for calculating the derivatives of $\mathbf{A}$ and $\mathbf{b}$ with respect to the control points [16]. The derivatives with respect to the parameters are then obtained in a postprocessing step by realizing that changing one parameter directly corresponds to changing all control points that are related to the parameter. Mathematically, this corresponds to applying the chain rule. This makes the evaluation of sensitivities computationally inexpensive, even for a large parameter set. If an additional parameter should be optimized, it only needs to be determined how the control points depend on it.

III. SURROGATE MODELING

The surrogate modeling approach employed in this work consists of two main components: POD, used to obtain a low-dimensional representation of the solution fields, and GPR, used to model the dependence of the reduced coefficients on the parameters. Following [12], both POD and GPR are enriched with gradient information from the IGA model to improve accuracy.

A. Proper Orthogonal Decomposition

POD finds a reduced representation of the solution starting from the snapshot matrix, formed by concatenating a set of solutions of the problem:

$$\mathbf{U} = [\mathbf{u}(\mathbf{p}_1), \dots, \mathbf{u}(\mathbf{p}_{n_s})] \in \mathbb{R}^{n \times n_s}, \quad (7)$$

where n is the number of degrees of freedom, while n_s is the number of samples in the set. The main assumption behind POD is that the solution vector can be approximated by a linear combination of $r \ll n$ basis vectors ϕ , weighted by parameter-dependent coefficients $\xi(\mathbf{p})$. The reduced basis $\mathbf{Q} = [\mathbf{q}_1, \dots, \mathbf{q}_r] \in \mathbb{R}^{n \times r}$ is commonly obtained by taking the first r left singular vectors from the Singular Value Decomposition (SVD) of $\mathbf{U}$. Following [11], POD is performed with a weighted inner product obtained via a weighting matrix $\bar{\mathbf{A}} = \mathbf{A}(\bar{\mathbf{p}}, \mathbf{0})$, chosen as the system matrix from Equation 4, evaluated with homogeneous unit material at the expected parameter combination $\bar{\mathbf{p}}$. Although $\mathbf{A}$ depends on the parameters, it is kept constant to simplify the basis computation. With this weighted inner product, the POD modes are obtained by solving the generalized eigenvalue problem

$$\mathbf{U}^\top \bar{\mathbf{A}} \mathbf{U} \phi_i = \lambda_i \phi_i, \quad i = 1, \dots, r, \quad (8)$$

and defining the reduced basis vectors as

$$\mathbf{q}_i = \frac{1}{\sqrt{\lambda_i}} \mathbf{U} \phi_i, \quad i = 1, \dots, r. \quad (9)$$

The reduced coefficients can then be computed for a generic parameter combination $\mathbf{p}^*$ by

$$\xi(\mathbf{p}^*) = \mathbf{Q}^\top \bar{\mathbf{A}} \mathbf{u}(\mathbf{p}^*). \quad (10)$$

Truncating the basis at a fixed value r introduces a reconstruction error, which can be monitored by:

$$\epsilon_{\text{rel}}^2(\mathbf{p}^*) = \frac{(\mathbf{u}(\mathbf{p}^*) - \mathbf{Q}\tilde{\xi}(\mathbf{p}^*))^\top \bar{\mathbf{A}} (\mathbf{u}(\mathbf{p}^*) - \mathbf{Q}\tilde{\xi}(\mathbf{p}^*))}{\mathbf{u}(\mathbf{p}^*)^\top \bar{\mathbf{A}} \mathbf{u}(\mathbf{p}^*)}. \quad (11)$$

After dimensionality reduction, machine learning approaches such as GPR can be used to find a function $\mathbf{f}$, mapping $\mathbf{p}$ to ξ , from which the field distribution can be reconstructed by:

$$\mathbf{u}(\mathbf{p}^*) \approx \tilde{\mathbf{u}}(\mathbf{p}^*) = \mathbf{Q}\xi, \quad \xi \approx \mathbf{f}(\mathbf{p}^*). \quad (12)$$

The accuracy of the POD reconstruction is directly influenced by the size of the basis. While increasing r improves accuracy, it also raises computational costs, requiring a compromise between efficiency and error. A typical approach is to choose r so that the relative cumulative energy:

$$\frac{\sum_{k=1}^r \lambda_k}{\sum_{k=1}^{n_s} \lambda_k} > \text{tol}, \quad (13)$$

is greater than a user-defined tolerance, typically set to be greater than 90% [11][21]. In this work, the rotation angle is treated explicitly by solving each rotor geometry for a set of discrete rotation angles. To achieve faster inference, rather than introducing it as an additional parameter, POD is applied to an expanded snapshot matrix obtained by concatenating the solutions corresponding to the discrete rotor positions. The resulting snapshot vectors therefore have dimension $n_\alpha n$. The proposed treatment of the rotor position does not increase the size of the eigenvalue problem (8), since its dimension is determined by the number of sampled parameter combinations n_s rather than by the size of the snapshot vectors. Furthermore, it does not affect the complexity of the subsequent GPR fit. Since the regression model is trained only on the reduced coefficients ξ , its computational cost is independent of the number of discrete rotor positions n_α .

When gradient information is available, a more general POD basis can be derived by augmenting the snapshot matrix $\mathbf{U}$ [12]:

$$\mathbf{U}_{\text{aug}} = [\mathbf{u}(\mathbf{p}_1), \mathbf{u}_1(\mathbf{p}_1), \dots, \mathbf{u}_{n_p}(\mathbf{p}_1), \dots, \mathbf{u}_{n_p}(\mathbf{p}_{n_s})], \quad (14)$$

where n_p is the number of parameters and $\mathbf{u}_i(\mathbf{p}_j)$ stands for [12]:

$$\mathbf{u}_i(\mathbf{p}_j) = \mathbf{u}(\mathbf{p}_j) + h \cdot \frac{\partial \mathbf{u}}{\partial p_i}(\mathbf{p}_j). \quad (15)$$

where h is a small value, chosen here to be 10^{-3} . Since the magnetic vector potential is learned as coefficients with respect to spline basis functions, the solution inherits the properties of the basis. Therefore, the proposed approach preserves several relevant structural and physical properties by construction. For instance, the resulting magnetic flux densities are solenoidal and exhibit the appropriate normal continuity across material interfaces, while the magnetic vector potentials and electric fields exhibit the corresponding tangential continuity. Moreover, all field energies are non-negative by construction. In general, however, the resulting fields do not satisfy the governing equations exactly, neither in strong nor in weak form, and other derived quantities may lack certain physical

properties. This distinguishes the proposed non-intrusive POD-regression framework from many classical projection-based MOR techniques [22], [23], which may preserve additional equation-based properties but are typically intrusive and require the solution of nonlinear systems for each evaluation in parameter space.

B. Gaussian Process Regression

This section briefly outlines GPR and its extension to the gradient-enhanced case. A more detailed presentation can be found in [24]. A Gaussian process $\mathcal{GP}$ is a distribution over functions and is defined by a mean function $\mu(\mathbf{x})$ and a covariance function $k(\mathbf{x}, \mathbf{x}')$ [24] as

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (16)$$

where the mean function

$$\mu(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})] \quad (17)$$

represents the expected function value for input $\mathbf{x}$, while the covariance function

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{x}') - \mu(\mathbf{x}'))] \quad (18)$$

models the relationship between function values at different input points $\mathbf{x}$ and $\mathbf{x}'$ [24], [25]. A set of training data $(\mathbf{X}_t, \mathbf{y}_t)$ is considered, where $\mathbf{X}_t$ corresponds to the parameter vectors $\mathbf{p}$ and $\mathbf{y}_t$ to the associated reduced coefficients ξ . A kernel is then selected, such as the radial basis function:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\lambda^2}\right), \quad (19)$$

where σ_f is the signal variance and λ is the length scale. Assuming zero mean, the function value for an unseen input $\mathbf{x}_*$ can then be obtained by computing the posterior mean of the distribution:

$$\mu(\mathbf{x}_*) = \mathbf{K}(\mathbf{x}_*, \mathbf{X}_t) [\mathbf{K}(\mathbf{X}_t, \mathbf{X}_t) + \sigma_\epsilon^2 \mathbf{I}]^{-1} \mathbf{y}_t, \quad (20)$$

where σ_ϵ^2 is the assumed noise level of the observations, $\mathbf{I}$ is the identity matrix and entries in the covariance matrix $\mathbf{K}$ are defined by evaluating the kernel function

$$\mathbf{K}(\mathbf{X}_a, \mathbf{X}_b) = \{k(\mathbf{x}_i^a, \mathbf{x}_j^b)\}_{ij}, \quad (21)$$

where $\mathbf{X}_a$ and $\mathbf{X}_b$ represent arbitrary sets of input points. One of the main advantages of GPR compared to other interpolation approaches is the availability of a confidence interval, which can be estimated by the variance function

$$\begin{aligned} \sigma^2(\mathbf{x}_*) &= k(\mathbf{x}_*, \mathbf{x}_*) \\ &\quad - \mathbf{K}(\mathbf{x}_*, \mathbf{X}_t) [\mathbf{K}(\mathbf{X}_t, \mathbf{X}_t) + \sigma_\epsilon^2 \mathbf{I}]^{-1} \mathbf{K}(\mathbf{X}_t, \mathbf{x}_*). \end{aligned} \quad (22)$$

The extension of GPR to the gradient-enhanced case (GEGPR) is performed by modeling the solution and its partial derivatives jointly. For each training point $\mathbf{x}$, the solution is extended from $y = f(\mathbf{x})$ to:

$$\mathbf{y} = \left[f(\mathbf{x}), \frac{\partial f}{\partial x_i}(\mathbf{x}), \dots, \frac{\partial f}{\partial x_{n_p}}(\mathbf{x}) \right]. \quad (23)$$




Fig. 2: Comparison between GFGPR and GEGPR. Black markers indicate the samples of the true function (represented by the dashed black line) used to fit the models. Red lines correspond to the prediction, while the shaded area is the 95% confidence interval.

Matrix $\mathbf{K}$ is then modified to include entries corresponding to the covariance between function values $k(\mathbf{x}, \mathbf{x}')$, covariance between function values and each partial derivative:

$$\frac{\partial k(\mathbf{x}, \mathbf{x}')}{\partial x_i}, \quad (24)$$

and covariance between partial derivatives [26]:

$$\frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_j}. \quad (25)$$

The primary computational bottleneck of GPR is the inversion of the covariance matrix, which scales asymptotically as $\mathcal{O}(n_s^3)$. In GEGPR, the covariance matrix is expanded to incorporate both function values and their gradients, resulting in a size of $n_s(1 + n_p)$. Consequently, the computational cost increases to $\mathcal{O}(n_s^3(1 + n_p)^3)$ [27]. Although this higher scaling may initially appear disadvantageous, the inclusion of gradient information effectively augments the dataset, increasing the number of observations from n_s to $n_s(1 + n_p)$. The main advantage of combining GEGPR with IGA is that sensitivity information can be obtained at a significantly lower cost, i.e., one linear system for all derivatives, than generating additional high-fidelity samples, requiring multiple linear systems to be solved. As a result, gradient information can enrich the training dataset with only a modest increase in the data-generation cost. Nevertheless, for large datasets and high-dimensional parameter spaces, the computational complexity of GPR can become prohibitive, in which case sparse approximations may be employed to improve scalability [28]. In Fig. 2, the standard GPR is compared to the gradient-enhanced version for modeling a simple one-dimensional function. The inclusion of derivative information improves the function estimate and tightens the confidence bounds.

IV. RESULTS

A parametrized IGA model of the geometry in Fig. 1 was developed, discretizing the physical equations into a nonlinear system of 6177 variables, which is a function of the geometrical parameters presented in Table I. To monitor the accuracy

TABLE I: Parameter bounds

Name	unit	lower bound	upper bound
MH	mm	2	12
MW	mm	8	22
MAG	mm	5	15
θ_1	°	15	23

TABLE II: Times for dataset generation and model fitting for different dataset partitions. Each sample is associated with 14 simulations with increasing rotation angle.

n_s	15	31	61	119
Simulation time (min)	30	62	123	239
Field GFGPR (s)	8.5	8.7	7.8	11.2
Field GEGPR (s)	16.8	25.2	53.4	124.5
Torque GFGPR (s)	13.2	9.6	10.9	13.3
Torque GEGPR (s)	23.7	55.0	60.7	103.0

of the approach in data-scarce conditions, Sobol sampling was employed to generate training datasets of increasing size. Four nested sets of 16, 32, 64, and 128 samples were initially generated within the bounds in Table I. All combinations leading to a magnet exceeding the geometric bounds were excluded from the dataset. Additionally, combinations in which the angle between the magnet edge and the adjacent rotor iron boundary was less than 10° were removed to eliminate geometries where the flux barrier becomes excessively skewed. All geometries were solved for a set of $n_\alpha = 14$ equidistant rotation angles in $[0, 20]^\circ$. The model was simulated on a workstation equipped with a 3.10 GHz Intel® Xeon® processor, with an average time of 9 seconds per solve, resulting in approximately 120 seconds per data point due to the 14 rotation angles. The resulting training datasets are summarized in Table II, while all models were tested on the same collection of 30 unseen combinations. The torque was subsequently obtained through postprocessing of the solution, yielding a distribution with a mean of 0.86 Nm and a standard deviation of 0.35 Nm.

A. Field Decomposition

For each dataset partition, the standard and augmented snapshot matrices were realized according to (7) and (14). To compare the two approaches, their accuracy was evaluated by monitoring the average relative L_2 error of the magnetic flux density through (11) over the test dataset, computed using $\hat{\mathbf{A}}$. In Fig. 3, the relative error is presented for the POD basis trained on the smallest dataset of 15 samples. In this scenario, the augmented snapshot matrix allows for an increased basis size without requiring additional physical samples, thus leading to a smaller reconstruction error. It should be noted, though, that a rank of 14 constitutes an edge case where the training data is perfectly recovered, while generalization to new samples is not guaranteed. In contrast, as presented in Fig. 4, when the number of available samples grows, an improvement in the reconstruction error begins to appear for truncations close to n_s . This behavior is attributed to information redundancy in the non-augmented dataset, which limits generalization when r approaches n_s . Since the goal of MOR is to obtain accurate but compact models, large truncation levels are undesirable. Consequently, for this application, snapshot augmentation provides only




Fig. 3: Average relative L_2 reconstruction error of the magnetic flux density for the gradient-free and gradient-enhanced basis for $n_s = 15$.




Fig. 4: Average relative L_2 reconstruction error of the magnetic flux density for the gradient-free and gradient-enhanced basis for $n_s = 119$.

marginal benefits, limited to the smallest dataset. Therefore, the standard POD basis is adopted in all scenarios, fixing the number of modes to 14, ensuring a fair comparison for the following latent-space modeling.

B. Field reconstruction

Using the POD basis, the reduced coefficients $\xi(\mathbf{p}^*)$ and their sensitivities are obtained for all parameter combinations in the training dataset. For each component of ξ , two separate GPR models (GF and GE) are trained to map the parameter space to the corresponding reduced coefficient. Both input and output data are standardized, and the corresponding gradients are consistently rescaled to match the transformed variables. Models are then trained using the L-BFGS optimizer with 80 iterations and 3 restarts using GPyTorch [29]. The two surrogate models are then compared by predicting the reduced coefficients $\hat{\xi}(\mathbf{p})$ starting from the parameter combinations in the test dataset. These predictions are subsequently used to reconstruct the full field solution $\hat{\mathbf{u}}(\mathbf{p})$ via (12), while the

relative error of the magnetic flux density is computed through (11), updating $\mathbf{A}$ for each parameter combination. As shown in Fig. 6 and Fig. 7, increasing the number of training samples improves the accuracy of both the GF and GE models, with their errors approaching the POD reconstruction error, which constitutes a baseline for the model accuracy. This indicates that the reduced coefficients ξ are predicted with high accuracy and that the remaining error is primarily due to the POD truncation. The GE approach consistently outperforms the GF method, particularly in data-scarce regimes, motivating the use of gradient information for latent-space modeling. However, as reported in Table II, this improved accuracy comes at the cost of increased training time due to the higher complexity of the model. A visual example of the error distribution of the POD-GEGPR model is presented Fig. 5 for a representative case with error close to the average.

C. KPI prediction

For the prediction of KPIs such as torque, two surrogate modeling strategies can be considered. The first consists of computing KPIs from the full-field surrogate model obtained in subsection IV-B. Alternatively, KPIs can be learned directly as functions of the input parameters using GPR, potentially leveraging the availability of torque sensitivities from the IGA model. For each rotor position, GF and GE GPR models are trained for torque prediction using the same standardization and training setup described in subsection IV-B. Their performance is assessed by comparing the Mean Absolute Percentage Error (MAPE) of the torque prediction, as shown in Fig. 8. Results indicate that, while both methods yield accurate predictions, directly learning the parameter-to-torque map leads to superior accuracy, especially in data-rich regimes where the field reconstruction error constitutes the limiting factor of POD-GPR. In addition, this direct approach eliminates the need to store or construct the full IGA model, resulting in a more compact surrogate representation. Moreover, since only one output is required, as presented in Table II, the training time is also decreased. Ultimately, this comparison is application-dependent, as it is influenced by the smoothness of the KPI and the reduced coefficients with respect to the parameter space, the truncation error, and the number of desired KPIs.

V. CONCLUSION

In this work, the differentiable nature of Isogeometric Analysis (IGA), enabling the efficient calculation of fields, Key Performance Indicators (KPIs), and their parametric sensitivities, is leveraged for efficient surrogate modeling of electric devices. Proper orthogonal decomposition is employed for dimensionality reduction, while gradient-enhanced Gaussian Process Regression (GPR) learns a mapping between parameters and field coefficients to generate surrogate models for the magnetic flux density distribution. Moreover, GPR is also used for parametric surrogate modeling of KPIs. The approach is assessed on a parametric model of a permanent magnet synchronous motor, considering both data-scarce and data-rich regimes. Results demonstrate that, for both the magnetic flux density and the motor torque, including gradient

information in the surrogate models significantly improves the model quality. While including gradient information increases the computational cost, it is justified in data-scarce regimes, where the available samples alone are insufficient to capture the underlying behavior accurately. Moreover, the surrogate models reduce evaluation times from tens of seconds for the full IGA simulations to milliseconds, while maintaining high predictive accuracy. Given the short training times and the very fast prediction of new solutions, this approach represents a promising tool for parametric surrogate modeling of electric devices.

ACKNOWLEDGMENT

This work was made possible by a scholarship from the German Academic Exchange Service (DAAD), through the program Research Grants in Germany, 2025 (57742125). The work is also supported by the joint DFG/WWF Collaborative Research Centre CREATOR (DFG: Project-ID 492661287/TRR 361; WWF: 10.55776/F90) at TU Darmstadt, TU Graz, and JKU Linz.

REFERENCES

- [1] P. Benner, S. Gugercin, and K. Willcox, "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems," *SIAM Review*, vol. 57, pp. 483–531, Jan. 2015.
- [2] A. Chatterjee, "An introduction to the proper orthogonal decomposition," *Current Science*, vol. 78, no. 7, pp. 808–817, 2000.
- [3] S. Chaturantabut and D. C. Sorensen, "Nonlinear Model Reduction via Discrete Empirical Interpolation," *SIAM Journal on Scientific Computing*, vol. 32, pp. 2737–2764, Jan. 2010.
- [4] T. Henneron and S. Cl  net, "Model-Order Reduction of Multiple-Input Non-Linear Systems Based on POD and DEI Methods," *IEEE Transactions on Magnetics*, vol. 51, pp. 1–4, Mar. 2015.
- [5] M. Zorzetto, R. Torchio, F. Lucchini, S. Massei, L. Robol, and F. Dughiero, "Reduced Order Modeling for Thermal Simulations of Electric Components With Surface-to-Surface Radiation," *IEEE Access*, vol. 12, pp. 178117–178126, 2024.
- [6] G. Ortali, N. Demo, G. Rozza, Mathematics Area, mathLab, SISSA, via Bonomea 265, I-34136 Trieste, Italy, and Department of Applied Physics, Eindhoven University of Technology, The Netherlands, "A Gaussian Process Regression approach within a data-driven POD framework for engineering problems in fluid dynamics," *Mathematics in Engineering*, vol. 4, no. 3, pp. 1–16, 2021.
- [7] T. Henneron, A. Pierquin, and S. Cl  net, "Surrogate Model Based on the POD Combined With the RBF Interpolation of Nonlinear Magnetostatic FE Model," *IEEE Transactions on Magnetics*, vol. 56, pp. 1–4, Jan. 2020.
- [8] M. Zorzetto, R. Torchio, F. Lucchini, M. Forzan, and F. Dughiero, "Proper Orthogonal Decomposition for Parameterized Macromodeling of a Longitudinal Electromagnetic Levitator," *IEEE Transactions on Magnetics*, vol. 61, pp. 1–7, Apr. 2025.
- [9] M. Zorzetto, R. Torchio, F. Lucchini, P. D. Barba, M. E. Mognaschi, M. Forzan, and F. Dughiero, "Machine learning-based reduced order modeling of nonlinear and multiphysics magnetic devices," *IEEE Transactions on Magnetics*, pp. 1–1, 2025.
- [10] F. Pichi, B. Moya, and J. S. Hesthaven, "A graph convolutional autoencoder approach to model order reduction for parametrized PDEs," *Journal of Computational Physics*, vol. 501, p. 112762, Mar. 2024.
- [11] M. Backmeyer, M. Wiesheu, and S. Sch  ps, "Learning electromagnetic fields based on finite element basis functions," *IEEE Transactions on Magnetics*, pp. 1–1, 2025.
- [12] R. Zimmermann, "Gradient-enhanced surrogate modeling based on proper orthogonal decomposition," *Journal of Computational and Applied Mathematics*, vol. 237, pp. 403–418, Jan. 2013.
- [13] U. Pahner, *A General Design Tool for the Numerical Optimisation of Electromagnetic Energy Transducers*. Phd thesis, KU Leuven, Leuven, 05 1998.




Fig. 5: Comparison between true solution and the POD-GEGPR ($n_s = 31$) model's prediction for $\mathbf{p} = [\text{MH} = 8.47 \text{ mm}, \text{MW} = 20.84 \text{ mm}, \text{MAG} = 10.17 \text{ mm}, \theta_1 = 21.68^\circ]$ and $\alpha = 6^\circ$, corresponding to a relative error of 2.84%.




Fig. 6: Average relative L_2 error of the magnetic flux density for gradient-free and gradient-enhanced GPR as a function of the number of samples in the training dataset, compared to the reconstruction error of POD.




Fig. 8: MAPE of the torque prediction from the gradient-free and gradient-enhanced POD-GPR models compared with GPR performed directly on the KPI (1% corresponds to 0.01 Nm at 1 Nm).




Fig. 7: Average relative L_2 error of the magnetic flux density in the air region, for gradient-free and gradient-enhanced GPR as a function of the number of samples in the training dataset, compared to the reconstruction error of POD.

- [14] T. Komann, M. Wiesheu, S. Ulbrich, and S. Schöps, "Robust design optimization of electric machines with isogeometric analysis," *Math.*, vol. 12, 04 2024.
- [15] S. J. Salon, *Finite Element Analysis of Electrical Machines*. Kluwer, 1995.
- [16] M. Wiesheu, T. Komann, M. Merkel, S. Schöps, S. Ulbrich, and I. Cortes Garcia, "Combined parameter and shape optimization of electric machines with isogeometric analysis," *Optim. Eng.*, 08 2024. arxiv:2311.06046.
- [17] H. Egger, M. Harutyunyan, R. Löscher, M. Merkel, and S. Schöps, "On torque computation in electric machine simulation by harmonic mortar methods," *J. Math. Ind.*, vol. 12, no. 6, 2022. arxiv:2112.05572.
- [18] A. Buffa, G. Sangalli, and R. Vázquez, "Isogeometric analysis in electromagnetics: B-splines approximation," *Comput. Meth. Appl. Mech. Eng.*, vol. 199, pp. 1143–1152, 2010.
- [19] A. Buffa, R. Vázquez Hernández, G. Sangalli, and L. Beirão da Veiga, "Approximation estimates for isogeometric spaces in multipatch geometries," *Numer. Meth. Part. Differ. Equat.*, vol. 31, no. 2, pp. 422–438, 2015.
- [20] M. Backmeyer, M. Wiesheu, and S. Schöps, "MotorLearning: Version paper with training data." Zenodo, 11 2025.
- [21] R. Swischuk, L. Mainini, B. Peherstorfer, and K. Willcox, "Projection-based model reduction: Formulations for physics-based machine learning," *Computers & Fluids*, vol. 179, pp. 704–717, Jan. 2019.
- [22] S. Volkwein, *Proper Orthogonal Decomposition: Theory and Reduced-Order Modelling*. Konstanz: University of Konstanz, 2013.

- [23] T. Henneron and S. Clénet, "Model order reduction of non-linear magnetostatic problems based on POD and DEI methods," *IEEE Trans. Magn.*, vol. 50, pp. 33–36, 02 2014.
- [24] E. Schulz, M. Speekenbrink, and A. Krause, "A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions," *Journal of Mathematical Psychology*, vol. 85, pp. 1–16, Aug. 2018.
- [25] T. J. Santner, B. J. Williams, and W. I. Notz, *The Design and Analysis of Computer Experiments*. Springer Series in Statistics, New York, NY: Springer New York, 2018.
- [26] C. E. Rasmussen and C. K. I. Williams, "Gaussian Processes for Machine Learning," in *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning, pp. 189–198, Cambridge, Mass.: MIT Press, 3. print ed., 2008.
- [27] V. Bone, C. van der Heide, K. Mackle, I. Jahn, P. M. Dower, and C. Manzie, "Gradient-Enhanced Deep Gaussian Processes for Multifidelity Modeling," *Journal of Computational Physics*, vol. 520, p. 113474, Jan. 2025.
- [28] J. Quinero-Candela and C. E. Rasmussen, "A Unifying View of Sparse Approximate Gaussian Process Regression," *Journal of machine learning research*, vol. 6, no. Dec, pp. 1939–1959, 2005.
- [29] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, "Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration," in *Advances in Neural Information Processing Systems*, 2018.