Engineering with Computers (2025) 41:2619–2639
https://doi.org/10.1007/s00366-025-02123-1

ORIGINAL ARTICLE

Fourier-enhanced reduced-order surrogate modeling for uncertainty quantification in electric machine design

Aylar Partovizadeh[1,2] · **Sebastian Schöps**[1,2,3] · **Dimitrios Loukrezis**[2,3]

Received: 27 November 2024 / Accepted: 16 February 2025 / Published online: 18 March 2025

Abstract

This work proposes a data-driven surrogate modeling framework for cost-effectively inferring the torque of a permanent magnet synchronous machine under geometric design variations. The framework is separated into a reduced-order modeling and an inference part. Given a dataset of torque signals, each corresponding to a different set of design parameters, torque dimension is first reduced by post-processing a discrete Fourier transform and keeping a reduced number of frequency components. This allows to take advantage of torque periodicity and preserve physical information contained in the frequency components. Next, a response surface model is computed by means of machine learning regression, which maps the design parameters to the reduced frequency components. The response surface models of choice are polynomial chaos expansions, feedforward neural networks, and Gaussian processes. Torque inference is performed by evaluating the response surface model for new design parameters and then inverting the dimension reduction. Numerical results show that the resulting surrogate models lead to sufficiently accurate torque predictions for previously unseen design configurations. The framework is found to be significantly advantageous compared to approximating the original (not reduced) torque signal directly, as well as slightly advantageous compared to using principal component analysis for dimension reduction. The combination of discrete Fourier transform-based dimension reduction with Gaussian process-based response surfaces yields the best-in-class surrogate model for this use case. The surrogate models replace the original, high-fidelity model in Monte Carlo-based uncertainty quantification studies, where they provide accurate torque statistics estimates at significantly reduced computational cost.

Keywords Dimension reduction · Discrete Fourier transform · Electric machine design · Machine learning regression · Permanent magnet synchronous machine · Reduced-order model · Response surface · Surrogate model · Uncertainty quantification

1 Introduction

Several sectors of major societal and economic interest, such as industry, infrastructure, and transportation, currently undergo transformative changes towards their electrification [1]. These changes drive an ever increasing usage of electric machines, which in turn imposes high demands on design quality and reliability. Considering computer-aided design (CAD) in particular, discrepancies between computer models and manufactured machines are all but inevitable, often affecting crucial quantities of interest (QoIs) [2]. To a large extent, these discrepancies can be attributed to *epistemic* uncertainty, referring to systematic modeling errors due to inexactly known physical mechanisms, or *aleatory* uncertainty, referring to inherently random variations in the design parameters [3]. This work focuses on aleatory uncertainty

✉ Aylar Partovizadeh
aylar.partovizadeh@tu-darmstadt.de

Sebastian Schöps
sebastian.schoeps@tu-darmstadt.de

Dimitrios Loukrezis
dimitrios.loukrezis@tu-darmstadt.de

[1] Computational Electromagnetics Group, Technische Universität Darmstadt, Schlossgartenstr. 8, 64289 Darmstadt, Germany

[2] Institute for Accelerator Science and Electromagnetic Fields, Technische Universität Darmstadt, Schlossgartenstr. 8, 64289 Darmstadt, Germany

[3] Graduate School of Computational Engineering, Technische Universität Darmstadt, Dolivostr. 15, 64293 Darmstadt, Germany

only. In this context, uncertainty quantification (UQ) aims at assessing the impact of random design parameters upon QoIs [4], such that more robust and reliable designs may be developed. However, workhorse UQ methods such as Monte Carlo sampling [5] can incur severe computational costs, due to their need for numerous evaluations of possibly time and resource-demanding numerical models and simulations, as is also the case in electric machine design.

A common way to mitigate the computational cost of UQ—or other studies that require repetitive model evaluations for varying parameter values, such as optimization or design exploration—is to replace the computationally demanding, high-fidelity model with an inexpensive albeit sufficiently accurate *surrogate* model [6]. A surrogate modeling approach that dates back to the work of Box and Wilson in the early 1950s, is to use a dataset of model parameters and corresponding QoI values to establish a so-called response surface model (RSM) that mimics the functional relation between parameters and QoI [7]. Among other options, polynomial chaos expansions (PCEs), feedforward neural networks (FNNs), and Gaussian processes (GPs) have been routinely used as RSMs [8–13]. An alternative approach is to decrease the cost of a computational model by reducing the dimensionality of the algebraic objects encountered in its numerical solution, resulting in a so-called reduced-order model (ROM). Depending on application and ROM method of choice, dimension reduction can be applied to the input parameters, the QoI, the system matrices of the numerical model, or combinations thereof. This typically entails computing an approximate, low-dimensional representation of the original, high-dimensional algebraic object, using methods such as balanced truncation, proper orthogonal decomposition, Krylov subspaces, tensor decompositions, or dynamic mode decomposition, to name but a few relevant options [14–17]. As a natural extension, surrogate modeling methods that combine ROMs with RSMs have also been developed [18–20]. Note that the literature does not always clearly distinguish between ROMs and surrogate models, even using the terms interchangeably on occasion. In the context of this work, these terms will be used as described above.

The present work focuses on the design of a permanent magnet synchronous machine (PMSM) under geometric variations, possibly random ones. Different PMSM geometries are obtained by varying 20 parameters of a CAD model [21], which take values within predefined ranges. The considered QoI is the electromagnetic torque developed during the machine's rotation, which is computed using a numerical model based on isogeometric analysis (IGA) [22]. For the torque signal to have sufficient resolution, e.g., to account for so-called torque ripples, the torque must be evaluated for a large number of rotation angles, thus resulting in a high-dimensional QoI. The high dimensionality of the QoI can hinder the application of direct surrogate modeling approaches, such as RSMs mapping the design parameters to the QoI, by increasing computation time and failing to reach sufficient approximation accuracy. An obvious solution is to use a dimension reduction method upon the torque signal. However, common methods do not take into account important properties of the torque signal, such as periodicity (for steady-state operation) and frequency content.

To compute a sufficiently accurate surrogate model that predicts the torque of the PMSM given different geometric designs and thus enable UQ studies, this work suggests a data-driven surrogate modeling framework utilizing discrete Fourier transform (DFT) for dimension reduction in combination with RSMs. The framework is separated into a reduced-order modeling part and an inference part. For reduced-order modeling, DFT is first applied to a dataset of torque signals, each obtained for different design parameters. Next, critical frequency components are retained while the rest are omitted, resulting in a reduced QoI that preserves the most important frequency content of the signal. Then, machine learning (ML) regression is used to compute an RSM that maps the design parameters to the reduced frequency components. To infer the torque given a new geometric design of the PMSM, the RSM is first evaluated, followed by inverting the DFT, thus obtaining the full torque signal.

From a methodological perspective, the use of DFT for the purpose of dimension reduction and its utilization in a surrogate modeling context is the main novelty of the suggested framework. The review of Hou and Behdinan [23] lists several prior works which have pursued similar ideas, wherein DFT-based dimension reduction is not included. For this particular use case, DFT ensures that important physical information contained in the frequency components is retained, thus leading to superior results in terms of surrogate modeling accuracy, as also supported by the numerical studies available in this work. The use of a reduced-order surrogate model, i.e., one that integrates dimension reduction of the QoI in combination with an RSM, is also a novel contribution concerning surrogate-assisted electric machine design, see for example the review of Cheng et al. [24] and the references therein. Last, contrary to most works encountered in the literature, the suggested framework is not restricted to a single RSM. Instead, different RSM options are evaluated, namely, PCE, FNN, and GP, in an attempt to identify the best-in-class surrogate model for this use case.

The remaining of this paper is organized as follows. Section 2 presents the PMSM model, its numerical approximation, and the considered geometric variations. Next, the data-driven reduced-order surrogate modeling

framework proposed in this work is detailed in Sect. 3. Numerical studies regarding the performance of surrogate models obtained with the suggested framework are presented in Sect. 4. Last, Sect. 5 contains concluding remarks.

2 Problem setting: PMSM with geometry variations

2.1 Geometry representation in CAD

We consider a PMSM similar to the one depicted in Fig. 1, where only one-fourth of a two-dimensional cross-section of its geometry is shown, taking into account rotational symmetry and translational invariance. The PMSM consists of a stator and a rotor, separated by an air gap. The rotor is composed of an iron part, permanent magnets, and air slots. The stator comprises an iron part and six slots where the copper windings carrying the source current are placed. For computational purposes, the air gap is divided into a rotor part and a stator part, separated by the interface Γ_{ag}. This is a benchmark PMSM geometry obtained from the commercial software `JMAG` [25]. For the purposes of the present study, the geometry is rebuilt using the `GeoPDEs` software [26, 27].

As is common in CAD, the PMSM's geometry is parametrized using free-form curves, such that the geometry is defined through a projection map from a reference domain $[0,1]^d$ to the physical domain $\Omega \subset \mathbb{R}^r$, $r \geq d$. The standard tools for this process are B-splines and non-uniform rational B-splines (NURBS) [28]. Given a knot vector $\Xi = \{\xi_1, \ldots, \xi_{n+p+1}\} \subset [0,1]$ such that $\xi_1 \leq \xi_2 \leq \cdots \leq \xi_{n+p+1}$, a basis of n univariate B-splines of degree p can be defined using the Cox-de-Boor recursion formula

$$B_i^p(\xi) = \frac{\xi - \xi_i}{\xi_{i+p} - \xi_i} B_i^{p-1}(\xi) + \frac{\xi_{i+p+1} - \xi}{\xi_{i+p+1} - \xi_{i+1}} B_{i+1}^{p-1}(\xi), \quad i = 1, \ldots, n, \tag{1a}$$

$$B_i^0(\xi) = \begin{cases} 1, & \xi_1 \leq \xi < \xi_{i+1}, \\ 0, & \text{otherwise.} \end{cases} \tag{1b}$$

Univariate NURBS basis functions are defined as

$$N_i^p(\xi) = \frac{w_i B_i^p(\xi)}{\sum_{j=1}^n w_j B_j^p(\xi)}, \tag{2}$$

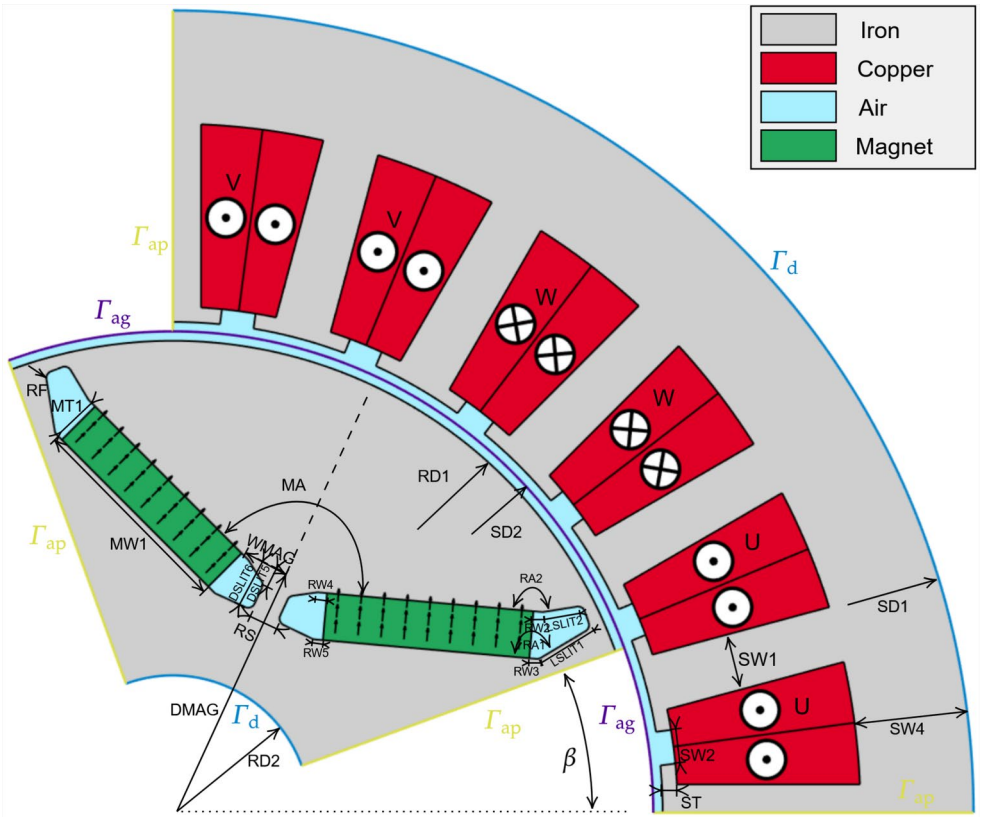

Fig. 1 One-fourth of the PMSM's geometry in two dimensions, including material distribution, boundary conditions, and geometric design parameters (see Table 1). The phase and current direction in the copper coils is also shown. Figure adapted from [21]

where w_i are positive weights. Using n NURBS basis functions along with n control points $\{\mathbf{P}_i\}_{i=1}^{n} \subset \Omega$, a NURBS curve can be defined as

$$\mathbf{C}(\xi) = \sum_{i=1}^{n} \mathbf{P}_i N_i^p(\xi). \tag{3}$$

NURBS surfaces or higher-dimensional objects can be created using tensor products of NURBS curves. Using these constructions, it is possible to define a map $F : [0,1]^d \to \Omega$ that parametrizes the physical domain such that $\mathbf{x} = F(\boldsymbol{\xi})$, $\mathbf{x} \in \Omega$, $\boldsymbol{\xi} \in [0,1]^d$.

Not that, in most practical applications, the geometry cannot be parametrized using a single projection map from the reference to the physical domain. This is also the case for the considered PMSM, which features a complex geometry with different material subdomains. In such cases, a multi-patch parametrization is used, where the physical domain is decomposed into a collection of subdomains, each with a corresponding projection map, to be appropriately combined [29].

2.2 Magnetostatic model

The magnetic field distribution on the PMSM is obtained by the magnetostatic formulation, which for a magnetic vector potential $\mathbf{A}$ reads

$$\nabla \times (\nu \nabla \times \mathbf{A}) = \mathbf{J}_{\text{src}} + \nabla \times (\nu \mathbf{B}_{\text{r}}), \tag{4}$$

where $\mathbf{B} = \nabla \times \mathbf{A}$ denotes the magnetic flux density, $\nu = \nu(|\mathbf{B}|)$ the (nonlinear) reluctivity, $\mathbf{J}_{\text{src}}$ the source current density, and $\mathbf{B}_{\text{r}}$ the remanent flux density of the permanent magnets. Note that this formulation neglects displacement and eddy currents, which is common practice for laminated PMSMs. Moreover, assuming invariance along the z-axis and neglecting three-dimensional effects, e.g., due to end windings, the two-dimensional problem is considered, simplifying (4) to the Poisson problem

$$\nabla \cdot (\nu \nabla A_{z,\text{rt}}) = \nu \nabla \cdot \mathbf{B}_{\text{r}}, \quad \text{in } \Omega_{\text{rt}}, \tag{5a}$$

$$\nabla \cdot (\nu \nabla A_{z,\text{st}}) = -J_{z,\text{src}}, \quad \text{in } \Omega_{\text{st}}, \tag{5b}$$

where Ω_{rt} and Ω_{st} denote the rotor and stator domains, respectively. Note that (5) requires only the z-components of the magnetic vector potential and the source current density, i.e., $\mathbf{A} = \left(0, 0, A_z\right)$ and $\mathbf{J}_{\text{src}} = \left(0, 0, J_{z,\text{src}}\right)$. Problem (5) is complemented with homogeneous Dirichlet and anti-periodic boundary conditions, respectively applied to the boundaries denoted with Γ_{d} and Γ_{ap} in Fig. 1. Additionally, field continuity at the air gap interface Γ_{ag} is ensured by the coupling conditions

$$A_{z,\text{st}}(\theta) = A_{z,\text{rt}}(\theta - \beta), \quad \text{on } \Gamma_{\text{ag}}, \tag{6a}$$

$$H_{\theta,\text{st}}(\theta) = H_{\theta,\text{rt}}(\theta - \beta), \quad \text{on } \Gamma_{\text{ag}}, \tag{6b}$$

where $\mathbf{H} = \nu \mathbf{B}$ is the magnetic field strength and A_z, H_θ are evaluated in local, rotor- or stator-specific coordinate systems that depend on the rotation angle β [30]. The remanent flux density $\mathbf{B}_{\text{r}}$ is given as $\mathbf{B}_{\text{r}} = B_{\text{r}}(-\sin(\alpha), \cos(\alpha))$, where α is the magnets' direction angle. Last, the source current density is given as $J_{z,\text{src}} = \sum_k J_{z,\text{src}}^{(k)}$, $k \in \{1,2,3\}$, and the k-th phase current is given as

$$J_{z,\text{src}}^{(k)} = J_0 \sin\left(n_{\text{pp}}\beta + \phi_0 + \frac{2\pi}{3}k\right), \tag{7}$$

where the magnitude J_0 depends on the application current and the coils' characteristics, ϕ_0 is the electric phase offset, and n_{pp} the number of pole pairs. A uniform distribution of $J_{z,\text{src}}$ within the coils is assumed.

2.3 Numerical approximation

The IGA method is used to solve problem (5) numerically, such that the basis functions of the CAD geometry representation are also used as basis functions for the numerical approximation of the solution [22]. Additionally, using harmonic mortaring for rotor-stator coupling [31] results in the matrix system

$$\underbrace{\begin{pmatrix} \mathbf{K}_{\text{rt}} & \mathbf{0} & -\mathbf{G}_{\text{rt}} \\ \mathbf{0} & \mathbf{K}_{\text{st}} & \mathbf{G}_{\text{st}}\mathbf{R}_\beta \\ -\mathbf{G}_{\text{rt}}^\top & \mathbf{R}_\beta^\top \mathbf{G}_{\text{st}}^\top & \mathbf{0} \end{pmatrix}}_{=:\mathbf{K}} \underbrace{\begin{pmatrix} \mathbf{u}_{\text{rt}} \\ \mathbf{u}_{\text{st}} \\ \lambda \end{pmatrix}}_{=:\mathbf{u}} = \underbrace{\begin{pmatrix} \mathbf{b}_{\text{rt}} \\ \mathbf{b}_{\text{st}} \\ \mathbf{0} \end{pmatrix}}_{=:\mathbf{b}}, \tag{8}$$

where $\mathbf{K}_{\text{rt}}$, $\mathbf{K}_{\text{st}}$ are stiffness matrices and $\mathbf{G}_{\text{rt}}$, $\mathbf{G}_{\text{st}}$ coupling matrices for the rotor and stator, respectively, $\mathbf{R}_\beta$ is the rotation matrix for a given rotation angle β, λ is the vector of Lagrange multipliers, and $\mathbf{b}_{\text{rt}}$, $\mathbf{b}_{\text{st}}$ are the rotor- and stator-specific right-hand side vectors. For more details on constructing the matrix system (8) the reader is referred to [21]. Note that the system matrix $\mathbf{K}$ actually depends on the solution vector $\mathbf{u}$ due to the nonlinear constitutive law. Based on energy conservation principles [30], the electromagnetic torque, denoted as τ_β for a given rotation angle β, is computed as

$$\tau_\beta(\mathbf{u}) = -\mathbf{u}_{\text{st}}^\top \mathbf{G}_{\text{st}} \frac{\mathrm{d}\mathbf{R}_\beta}{\mathrm{d}\beta} \lambda L, \tag{9}$$

where L is the axial length of the PMSM. The latter must be explicitly taken into account since the numerical solution corresponds to the two-dimensional problem.

Table 1 Geometric design parameters of the PMSM and their value ranges

#	Parameter	Lower bound	Upper bound
1	LSLIT1	6.1	6.7
2	LSLIT2	4.1	4.5
3	DSLIT5	0.9	1.1
4	DSLIT6	1.9	2.1
5	MA	142.5	157.5
6	MT1	3.8	4.2
7	MW1	20.9	23.1
8	RA1	136.8	151.2
9	RA2	157.7	174.3
10	RS	0.9	1.1
11	RW2	0.9	1.1
12	RW3	0.9	1.1
13	RW4	0.9	1.1
14	RW5	0.9	1.1
15	WMAG	3.8	4.2
16	DMAG	28.5	31.5
17	ST	1.5	1.7
18	SW1	5.7	6.3
19	SW2	3.4	3.8
20	SW4	24.2	26.8

The lower and upper bounds of the parameter values are given in mm

2.4 PMSM model parameters and design variations

Following Wiesheu et al. [21], a PMSM model with the standard M330-50A material is utilised. The equivalent M27 data from the FEMM software [32] are used to define the nonlinear relative permeability, i.e., the inverse of the magnetic reluctivity (see Sect. 2.2) of the iron parts. Nd-Fe-B (linear) magnets with remanence $B_r = 1.0$ T and relative permeability $\mu_r = 1.05$ are used. The PMSM operates with an applied current $I_{app} = 3$ A and its windings have $n_{wind} = 35$ turns. The IGA discretization of the PMSM's geometry utilises 32 patches with 444 control points for the rotor domain and 144 patches with 365 control points for the stator domain. The geometry of the PMSM depends on $P = 20$ design parameters, which are shown in Fig. 1 and listed in Table 1. These parameters take values within predefined ranges, also listed in Table 1. The different geometric design configurations arising from the possible parameter combinations have a significant impact on the computational domain of the PMSM model, its numerical solution, and finally on QoIs estimated by post-processing the solutions, such as the electromagnetic torque. The latter is the main QoI in this work. Collecting all geometric design parameters in a vector $\mathbf{p} \in \mathbb{R}^P$, system (8) becomes parameter-dependent with system matrix $\mathbf{K} = \mathbf{K}(\mathbf{p})$, right-hand side $\mathbf{b} = \mathbf{b}(\mathbf{p})$, and solution $\mathbf{u} = \mathbf{u}(\mathbf{p})$. Accordingly, the electromagnetic torque is parameter-dependent as well. We write, by abuse of notation, $\tau_\beta(\mathbf{u}) = \tau_\beta(\mathbf{u}(\mathbf{p})) = \tau_\beta(\mathbf{p})$.

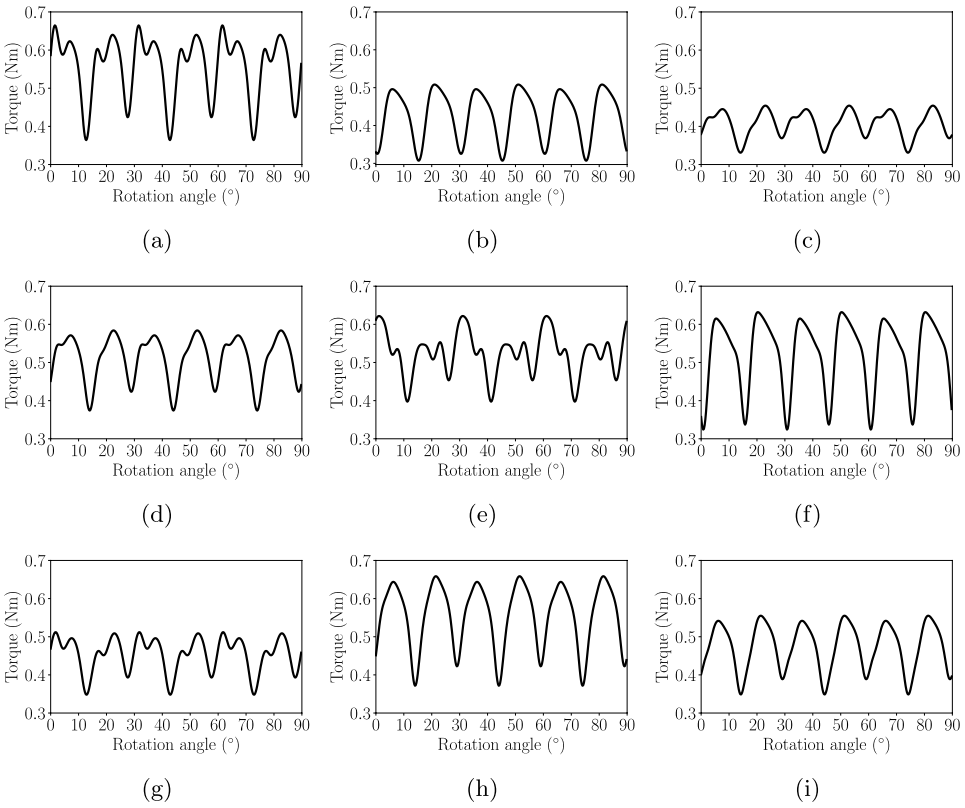

Fig. 2 Torque signals for different geometric design configurations of the PMSM

Table 2 Geometric design parameters values (in mm) corresponding to the torque signals shown in Fig. 2

#	Parameter	2(a)	2(b)	2(c)	2(d)	2(e)	2(f)	2(g)	2(h)	2(i)
1	LSLIT1	6.27	6.26	6.51	6.25	6.54	6.18	6.53	6.29	6.39
2	LSLIT2	4.26	4.09	4.26	4.20	4.10	4.17	4.38	4.40	4.45
3	DSLIT5	1.02	0.96	0.96	0.98	0.97	1.04	0.98	0.99	1.02
4	DSLIT6	1.96	2.00	1.93	1.94	2.07	2.01	1.99	1.98	2.06
5	MA	156.89	147.96	155.08	155.12	153.99	149.09	156.30	149.13	151.70
6	MT1	3.80	3.92	3.82	4.09	3.84	3.86	4.16	3.91	4.03
7	MW1	22.99	20.93	20.98	21.64	22.94	20.98	21.68	22.11	21.18
8	RA1	138.02	136.84	145.08	150.34	148.50	138.26	139.43	146.07	143.15
9	RA2	158.11	172.73	157.91	163.86	173.99	162.99	169.16	159.35	173.63
10	RS	0.97	1.01	1.00	1.02	0.99	1.05	0.99	0.96	0.99
11	RW2	1.01	0.99	0.98	0.96	0.96	0.97	0.96	1.04	0.98
12	RW3	0.99	0.97	0.99	1.00	0.97	1.02	0.97	0.96	0.95
13	RW4	1.02	0.96	0.99	1.02	1.02	1.04	1.02	0.96	0.96
14	RW5	1.00	1.05	1.00	1.00	1.01	0.98	1.02	1.01	0.98
15	WMAG	4.00	3.99	3.86	4.02	3.87	4.05	3.98	4.19	3.85
16	DMAG	30.01	28.97	28.52	30.25	29.06	30.59	29.30	29.27	29.84
17	ST	1.65	1.67	1.63	1.63	1.62	1.56	1.65	1.54	1.63
18	SW1	6.01	5.85	6.13	6.09	5.90	5.99	5.90	5.90	5.85
19	SW2	3.74	3.61	3.52	3.56	3.48	3.48	3.61	3.68	3.55
20	SW4	25.48	25.41	26.66	24.56	26.13	25.85	26.71	24.41	26.07

To illustrate the impact of geometric design variations upon the torque of the PMSM, Fig. 2 shows samples of torque signals, each corresponding to a different set of design parameter values given in Table 2. Note that the torque repeats every 30° of rotation due to the symmetry in the stator's winding pattern. As can be observed, geometric design variations lead to significant changes in the torque signal. Also note that generating the torque signal for each geometric design variation entails running a computationally expensive numerical simulation for the solution of problem (5). A computationally inexpensive estimation of the torque would be preferable, especially considering UQ [33–36] or related studies like robust or reliability-based optimization [37–39], that necessitate repetitive and possibly numerous model evaluations for varying PMSM designs.

3 Reduced-order surrogate modeling framework

To address the need for computationally inexpensive torque estimates given different design configurations of the PMSM, we propose a data-driven framework for computing a surrogate model that can reliably replace the original, high-fidelity model in parameter studies such as UQ. The framework consists of a reduced-order modeling part and an inference part.

The reduced-order modeling part of the framework consists of the following steps:

0. *Data acquisition:* We assume that a dataset $\mathcal{D} = \left\{\mathbf{p}^{(m)}, \boldsymbol{\tau}^{(m)} = \boldsymbol{\tau}\left(\mathbf{p}^{(m)}\right)\right\}_{m=1}^{M}$ is made available, containing different parameter realizations $\mathbf{p}^{(m)} \in \mathbb{R}^P$ and the corresponding torque signals $\boldsymbol{\tau}^{(m)} \in \mathbb{R}^N$, where N is the signal's dimension, defined by the number of rotation angles β_n for which the corresponding torque values τ_{β_n} are obtained, such that $\boldsymbol{\tau}^{(m)} = \left\{\tau_{\beta_n}^{(m)}\right\}_{n=0}^{N-1}$. Continuous torque signals, e.g., as in Fig. 2, are obtained by means of linear interpolation. In this work, such a dataset is obtained by evaluating the PMSM's numerical model for different design configurations. However, it is possible to use experimental data as well, e.g., torque signals measured for different machine configurations, possibly even mixing simulation-generated and measured data.
1. *Dimension reduction of the QoI:* The torque signals $\boldsymbol{\tau}^{(m)}$, $m = 1, \ldots, M$, are processed such that the original dimension of the QoI is significantly reduced. In this work, dimension reduction is performed by applying DFT upon the torque signals and retaining the most important frequency components. Essentially, the dimension reduction step creates a map $\mathcal{R} : \boldsymbol{\tau} \rightarrow \mathbf{r}$, where $\mathbf{r}$ denotes the reduced QoI, equivalently, the representation of the original QoI in a reduced frequency space.

2. *ML regression:* RSMs are trained by means of ML regression, as to map the design parameters to the reduced frequency components of the torque signals. We denote this map as $\mathcal{S} : \mathbf{p} \rightarrow \mathbf{r}$. Three RSMs are employed for that purpose, namely, PCE, FNN, and GP.

The inference part of the framework consists of the following steps:

1. *Response surface evaluation:* Given a new design configuration and the corresponding design parameter vector $\mathbf{p}^*$, the RSM is evaluated as to yield an estimate of the reduced QoI, such that $\mathbf{r}^* = \mathcal{S}(\mathbf{p}^*)$.
2. *Dimension reduction inversion:* The torque estimate for the new design configuration is obtained by inverting the dimension reduction procedure used in the reduced-order modeling part, such that $\boldsymbol{\tau}^* = \mathcal{R}^{-1}(\mathbf{r}^*)$.

A visualization of the complete framework is provided in Fig. 3. For UQ purposes, the resulting surrogate model can replace the high-fidelity model within a Monte Carlo sampling algorithm to obtain estimates of the torque's statistics (e.g., expected value, variance) cost-effectively. Details regarding the DFT-based dimension reduction approach are given in Sect. 3.1. The computation of RSMs by means of ML regression is presented in Sect. 3.2.

Remark on computational cost The computational cost of the suggested surrogate modeling framework concerns mainly the reduced-order modeling part, in particular data acquisition and ML regression, collectively called "offline" costs. In this work, data acquisition is the main computational cost to be considered, since datasets are generated by running computationally expensive numerical model evaluations. In other cases, e.g., for datasets that are readily available, this cost can become insignificant. ML regression costs may range from negligible to quite significant, depending on choices regarding the type of ML model, training/optimization algorithm, and size of training dataset. In this work, ML regression costs can be neglected next to data acquisition costs. Inference (also called "online") costs can similarly be neglected, even if one considers numerous surrogate model evaluations, for example, in the context of sampling-based UQ studies. The cost remains negligible due to the computationally inexpensive surrogate model. A detailed discussion on the computational costs and gains in the context of the present work is provided in Sect. 4.3.3.

3.1 Dimension reduction via discrete Fourier transform

In the following, we assume that M torque signals $\boldsymbol{\tau}^{(m)}$, $m = 1, \dots, M$, are available, consisting of N torque values within one rotation period. Longer torque signals can always be reduced to a single period without loss of information. Due to the fact that the dimensionality of the torque signal can be prohibitively large for the subsequent steps of the surrogate modeling framework, we suggest an approach based on the DFT to reduce the torque's dimension to its minimum necessary frequency components.

The DFT is first used to convert a finite sequence of equally-spaced torque samples (in time, equivalently, with respect to rotation angle) into a sequence of the same length, where the latter contains information about the underlying frequency components that contribute to the given torque signal. Given the torque signal as a sequence of torque values per rotation angle, i.e., $\boldsymbol{\tau}^{(m)} = \left\{ \tau_{\beta_n}^{(m)} \right\}_{n=0}^{N-1}$, the DFT yields the sequence of complex values $\mathbf{c}^{(m)} = \left\{ c_k^{(m)} \right\}_{k=0}^{N-1}$, such that

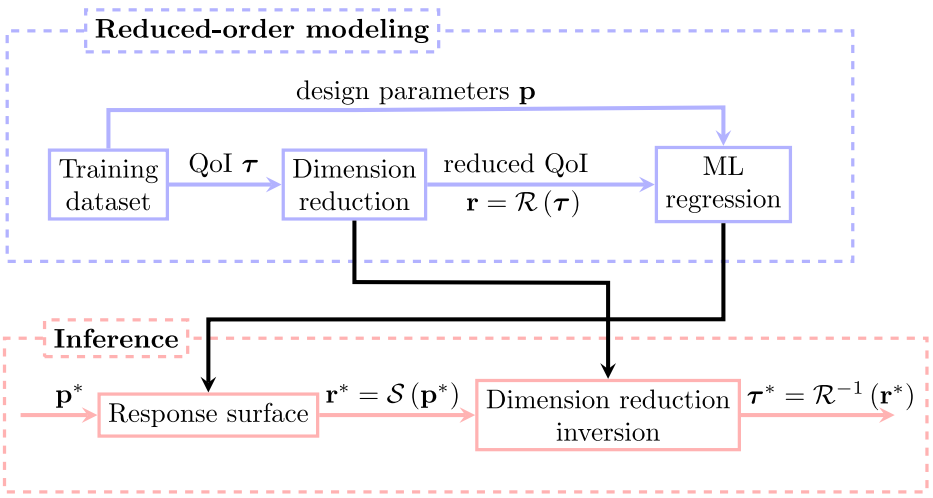

Fig. 3 Schematic illustration of the surrogate modeling framework, separated into reduced-order modeling and inference parts

$$c_k^{(m)} = \sum_{n=0}^{N-1} \tau_{\beta_n}^{(m)} \exp\left(-\frac{\imath 2\pi kn}{N}\right) = \sum_{n=0}^{N-1} \tau_{\beta_n}^{(m)} \left(\cos\left(\frac{2\pi kn}{N}\right) - \imath \sin\left(\frac{2\pi kn}{N}\right)\right), \tag{10}$$

where $c_0^{(m)} \in \mathbb{R}$ corresponds to the average value of the original torque signal $\boldsymbol{\tau}^{(m)}$ across all rotation-angle samples. Since the torque signal is real-valued, its average value is real-valued as well. The DFT can be inverted to reconstruct the original signal, such that

$$\tau_{\beta_n}^{(m)} = \frac{1}{N} \sum_{k=0}^{N-1} c_k^{(m)} \exp\left(-\frac{\imath 2\pi kn}{N}\right). \tag{11}$$

For the purpose of dimension reduction, we take advantage of the fact that the power spectrum of a given torque signal $\boldsymbol{\tau}^{(m)}$ can be given in terms of its frequency components as

$$\mathrm{PS}_k^{(m)} = \frac{1}{N}\left|c_k^{(m)}\right|^2, \tag{12}$$

meaning that each frequency component has a distinct contribution to the power spectrum, quantified by the value $\eta_k^{(m)} = \left|c_k^{(m)}\right|^2$. Omitting components with negligible contributions to the power spectrum allows to reduce the dimensionality of the output of the DFT without significantly affecting the value of the power spectrum, but also without compromising the accuracy of the torque's reconstruction given in (11) due to the small values of the omitted constants $c_k^{(m)}$.

However, for varying design configurations, each corresponding to a different torque signal, it is to be expected that the DFT output will differ from one signal to another. In turn, the classification of frequency components to negligible or retainable shall differ among the torque signals. To address this issue, we suggest to retain the frequency components that remain important over the possible design configurations in the sense of expected contribution to the power spectrum. To that end, the contribution of each frequency component to the power spectrum over all given torque signals $\boldsymbol{\tau}^{(m)}$, $m = 1, \ldots, M$, is averaged as to obtain a single contribution per component, computed as

$$\overline{\eta}_k = \frac{1}{M} \sum_{m=1}^{M} \eta_k^{(m)} = \frac{1}{M} \sum_{m=1}^{M} \left|c_k^{(m)}\right|^2. \tag{13}$$

Subsequently, the frequency components are sorted based on their averaged power spectrum contributions. The dimension reduction of the DFT's output proceeds by keeping the frequency components corresponding to significant averaged power spectrum contributions $\overline{\eta}_k$ and omitting the rest. One way of distinguishing between frequency components being negligible or not is presented in Sect. 4.1, based on the accuracy of the reconstructed torque signal.

This procedure results in a reduced QoI $\mathbf{r} \in \mathbb{C}^R$, $R < N$, with $\mathbf{r} = \left\{c_{k'}\right\}_{k' \in \mathcal{I}}$, where $\mathcal{I} \subset \{k\}_{k=0}^{N-1}$ is an index set with cardinality equal to R that contains the indices of the retained frequency components. Put in a ROM context, $\mathbf{r} \in \mathbb{C}^R$ is the reduced representation (here, in frequency space) of $\boldsymbol{\tau} \in \mathbb{R}^N$ and $\left\{c_{k'}\right\}_{k' \in \mathcal{I}}$ is the corresponding reduced basis. Accordingly, the realizations of the reduced QoI are given as $\mathbf{r}^{(m)} = \left\{c_{k'}^{(m)}\right\}_{k' \in \mathcal{I}}$, $m = 1. \ldots, M$. Note that the reduced QoI can be given equivalently as a real vector $\mathbf{r} \in \mathbb{R}^{2R-1}$, by considering separately the real and imaginary parts of the complex elements $c_{k'}$, $k' \neq 0$, where $c_0 \in \mathbb{R}$. We opt for the more concise complex notation.

3.2 Regression-based response surface models

Within the suggested surrogate modeling framework, data-driven, regression-based RSMs are employed to approximate the functional relationship between the geometric parameters of the PMSM and the reduced frequency components of the torque signal, thus providing a map $\mathcal{S} : \mathbf{p} \to \mathbf{r}$. For comparison purposes, in the numerical experiments presented in Sect. 4, the RSMs are employed as approximations to full torque signals and to their reduced principal components. Therefore, in the following, the functional relation is denoted as $\mathbf{y} = f(\mathbf{p})$, where $\mathbf{y}$ may refer to either the full or reduced QoI. Without loss of generality, we assume that $\mathbf{y} \in \mathbb{R}^N$. For the regression-based computation of the RSMs, we assume the existence of a dataset $\mathcal{D}_\mathrm{t} = \left\{\mathbf{p}^{(m)}, \mathbf{y}^{(m)}\right\}_{m=1}^{M}$ containing input parameter realizations along with the corresponding values of the QoI. Three RSMs are examined, namely, PCE, FNN, and GP, which are briefly presented next.

3.2.1 Polynomial chaos expansion (PCE)

For response surface modeling by means of the PCE, the parameter vector $\mathbf{p}$ is assumed to be a random vector defined on the probability space $(\Theta, \Sigma, \mathcal{P})$, where Θ denotes the sample space, Σ the sigma algebra of events, and $\mathcal{P} : \Sigma \to [0, 1]$ a probability measure, and with the probability density function (PDF) $\varrho(\mathbf{p}) : \Xi \to \mathbb{R}_{\geq 0}$, where $\Xi \subset \mathbb{R}^P$ denotes the random vector's support. Under these assumptions, a PCE is a global polynomial approximation of the form

$$f(\mathbf{p}) \approx f_{\mathrm{PCE}}(\mathbf{p}) = \sum_{k=1}^{K} \mathbf{a}_k \Psi_k(\mathbf{p}), \tag{14}$$

where $\mathbf{a}_k \in \mathbb{R}^N$ are expansion coefficients and Ψ_k are multivariate polynomials that satisfy the orthogonality property

$$\mathbb{E}\left[\Psi_k \Psi_l\right] = \int_{\Xi} \Psi_k(\mathbf{p}) \Psi_l(\mathbf{p}) \varrho(\mathbf{p}) \mathrm{d}\mathbf{p} = \mathbb{E}\left[\Psi_k^2\right] \delta_{kl}, \tag{15}$$

where δ_{kl} is the Kronecker delta. Depending on the PDF, the polynomials are chosen either according to the Wiener-Askey scheme [40] or are constructed numerically [41]. To evenly cover the design space, we assume here that all design parameters are uniformly distributed within their value ranges, also see Sect. 4.1. Therefore, Legendre polynomials are used. The polynomial basis of the PCE can be formed in several ways, which also define the truncation limit K. Common options include total degree, hyperbolic truncation, sparse, and adaptive bases, and combinations thereof [42].

Training the PCE means fitting its coefficients based on the data $\mathcal{D}_{\mathrm{t}}$. In particular, the coefficients are computed by solving the least squares regression problem

$$\arg\min_{\mathbf{a}_1,\ldots,\mathbf{a}_K} \left\{ \frac{1}{M} \sum_{m=1}^{M} \left(\mathbf{y}^{(m)} - \sum_{k=1}^{K} \mathbf{a}_k \Psi_k\left(\mathbf{p}^{(m)}\right) \right)^2 \right\}. \tag{16}$$

For sparse PCEs, the regression problem (16) is complemented with an additional penalty term that induces sparsity in its solution. In this work, we opt for a sparse and adaptive PCE method based on least angle regression (LAR) [43], which is implemented with the `UQLab` software [44].

3.2.2 Feedforward neural network (FNN)

To define an FNN, we must first introduce the concepts of neuron and layer. A neuron is the smallest unit of an FNN and is mathematically described by the function

$$\nu(\mathbf{z}) = \sigma\left(\mathbf{z}^\top \mathbf{w} + b\right), \tag{17}$$

where $\mathbf{z} \in \mathbb{R}^D$ is the input to the neuron, $\mathbf{w} \in \mathbb{R}^D$ a weight vector, $b \in \mathbb{R}$ a bias term, and $\sigma : \mathbb{R} \to \mathbb{R}$ a nonlinear function commonly referred to as the activation. A layer is formed by a set of K neurons that receive the same input $\mathbf{z} \in \mathbb{R}^D$ and is mathematically described by the function

$$\ell(\mathbf{z}) = \sigma\left(\mathbf{z}^\top \mathbf{W} + \mathbf{b}\right), \tag{18}$$

where $\mathbf{W} \in \mathbb{R}^{D \times K}$ is a matrix that contains the weights of the K neurons in the layer, $\mathbf{b} \in \mathbb{R}^K$ is the corresponding bias vector, and the activation function is applied elementwise per neuron. The FNN is formed as a sequence of L layers and is mathematically described by the layer composition formula

$$f_{\mathrm{FNN}}(\mathbf{z}) = \ell^L\left(\ell^{L-1}\left(\cdots\left(\ell^1(\mathbf{z})\right)\right)\right). \tag{19}$$

To be used as an RSM $f_{\mathrm{FNN}}(\mathbf{p}) \approx f(\mathbf{p})$, the FNN's trainable parameters, i.e., the collection of its weights and biases $\boldsymbol{\theta} = \left\{\mathbf{W}^{(l)}, \mathbf{b}^{(l)}\right\}_{l=1}^{L}$ must be fitted to the training data $\mathcal{D}_{\mathrm{t}}$. Here, this is accomplished by minimizing the mean square error (MSE) loss function

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{M} \sum_{m=1}^{M} \left(\mathbf{y}^{(m)} - f_{\mathrm{FNN}}\left(\boldsymbol{\theta}; \mathbf{p}^{(m)}\right)\right)^2, \tag{20}$$

resulting in a nonlinear optimization problem, the solution of which is typically computed by means of stochastic gradient descent algorithms [45] enabled by automatic differentiation [46].

In this work we opt for an FNN consisting of four fully-connected hidden layers with 45, 60, 80, and 25 neurons, respectively. The number of layers and neurons per layer has been determined by considering an initial FNN with three layers and 20 neurons per layer, and then incrementally increasing the number of layers and neurons per layer until no significant improvement in the approximation error was observed. Naturally, the number of neurons used in the output layer coincide with the dimensions of the QoI (original or reduced). Rectified linear unit (ReLU) activation functions are used for all neurons. The FNN is trained for 600 epochs using a standard stochastic gradient descent optimizer with scheduled learning rate. The `Tensorflow` software [47] is used for implementation.

3.2.3 Gaussian process (GP)

A GP is formally defined as a collection of random variables, any finite number of which have a joint Gaussian distribution [48]. Considering temporarily a scalar QoI $y = f(\mathbf{p})$, an approximation by a GP interprets the functional relation as a probability distribution in function space, such that

$$f(\mathbf{p}) \sim \mathcal{GP}\big(m(\mathbf{p}), k\big(\mathbf{p}, \mathbf{p}'\big)\big), \tag{21}$$

where $m(\mathbf{p})$ is a mean function and $k(\mathbf{p}, \mathbf{p}')$ a covariance function. The mean and covariance functions encode prior assumptions regarding the target function, e.g., concerning its expected behavior and regularity. The prior assumptions are then combined with the data available in the training dataset $\mathcal{D}_\mathrm{t}$ as to obtain a new GP, called the posterior, with updated mean and covariance functions. Defining the matrix $\mathbf{K} \in \mathbb{R}^{M\times M}$ with elements $k_{ij} = k\big(\mathbf{p}^{(i)}, \mathbf{p}^{(j)}\big)$, and the vector $\mathbf{m} = \big(m\big(\mathbf{p}^{(1)}\big), \ldots, m\big(\mathbf{p}^{(M)}\big)\big) \in \mathbb{R}^M$, the updated (posterior) distribution of the target function, i.e., conditioned on the training dataset $\mathcal{D}_\mathrm{t}$, is given as

$$f(\mathbf{p})|\mathcal{D}_\mathrm{t} \sim \mathcal{N}(\mathbf{m}, \mathbf{K}), \tag{22}$$

where $\mathcal{N}(\cdot,\cdot)$ denotes a Gaussian distribution. For a new data point $\mathbf{p}^* \notin \mathcal{D}_\mathrm{t}$, the predictive distribution is given as

$$f(\mathbf{p})|\mathbf{p}^*, \mathcal{D}_\mathrm{t} \sim \mathcal{N}\Bigg(\underbrace{m(\mathbf{p}^*) + \mathbf{k}(\mathbf{p}^*, \mathbf{P})\mathbf{K}^{-1}(\mathbf{y} - \mathbf{m})^\top}_{\text{mean}}, \underbrace{k(\mathbf{p}^*, \mathbf{p}^*) - \mathbf{k}(\mathbf{p}^*, \mathbf{P})\mathbf{K}^{-1}\mathbf{k}(\mathbf{P}, \mathbf{p}^*)^\top}_{\text{variance}}\Bigg), \tag{23}$$

where $\mathbf{P} = \big(\mathbf{p}^{(m)}\big)_{m=1}^M$, $\mathbf{y} = \big(y^{(m)}\big)_{m=1}^M$, $\mathbf{k}(\mathbf{p}^*, \mathbf{P}) = \big(k\big(\mathbf{p}^*, \mathbf{p}^{(m)}\big)\big)_{m=1}^M$, and $\mathbf{k}(\mathbf{P}, \mathbf{p}^*) = \big(k\big(\mathbf{p}^{(m)}, \mathbf{p}^*\big)\big)_{m=1}^M$. The predictions of the GP-based RSM are obtained by evaluating the mean function of the predictive distribution upon new data points, equivalently,

$$f(\mathbf{p}) \approx f_\mathrm{GP}(\mathbf{p}) = m(\mathbf{p}) + \mathbf{k}(\mathbf{p}, \mathbf{P})\mathbf{K}^{-1}(\mathbf{y} - \mathbf{m})^\top. \tag{24}$$

In this work, the assumed prior has zero mean and the ellipsoidal Gaussian covariance function

$$k(\mathbf{p}, \mathbf{p}'|\boldsymbol{\theta}) = \exp\left(-\frac{1}{2}\sum_{i=1}^{P}\left(\frac{p_i - p_i'}{\theta_i}\right)^2\right), \tag{25}$$

where $\boldsymbol{\theta} = \big(\theta_i\big)_{i=1}^P$ is a vector of hyperparameters. Note that, for ease of notation, the hyperparameter vector $\boldsymbol{\theta}$ has been omitted prior to (25). Training the GP-based RSM entails optimizing the hyperparameters such that the likelihood of observing the data contained within the training dataset $\mathcal{D}_\mathrm{t}$, denoted as $\mathcal{P}\big(\mathcal{D}_\mathrm{t}, \boldsymbol{\theta}\big)$, is maximized [48]. Here, this is accomplished with a covariance matrix adaptation evolution strategy (CMA-ES) optimizer [49]. Similar to the PCE, the implementation of the GP-based RSM is based on the `UQLab` software [44].

Note that the presentation above concentrated on a scalar QoI and its approximation by a GP. Multi-output GPs can be obtained in various ways, for example, using a so-called co-Kriging model or assuming a multi-output covariance [50]. In this work, we opt for the simplest solution of considering the elements of a vector-valued QoI to be independent and treating them separately, which was found to be sufficient.

4 Numerical studies and results

In the following, we assess several aspects of suggested reduced-order surrogate modeling framework, in particular concerning the dimension reduction achieved for the QoI, the performance of the resulting surrogates in terms

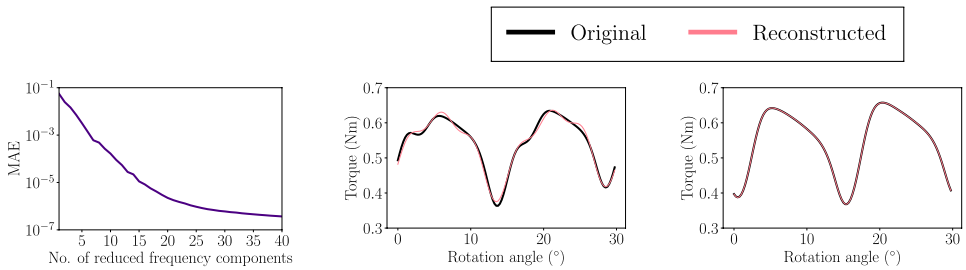

(a) Reconstruction error versus reduced frequency components.

(b) Worst-case reconstruction with $R = 5$ frequency components.

(c) Worst-case reconstruction with $R = 11$ frequency components.

Fig. 4 Reconstruction error versus dimension reduction and worst-case torque signal reconstructions using DFT-based dimension reduction with $R = 5$ and $R = 11$ reduced frequency components

of prediction accuracy, and the accuracy of statistics estimates when the surrogates are used for UQ. The corresponding numerical studies are presented in Sects. 4.1, 4.2, and 4.3, respectively.

4.1 Dimension reduction of the QoI

As previously explained in Sect. 3.1, by keeping a reduced number of the frequency components obtained via DFT, the dimensions of the QoI can be reduced significantly. However, the number of reduced frequency components must be determined considering the trade-off between dimension reduction and the error between the reconstructed QoI, obtained by inverting the DFT, and the original. Therefore, in this first numerical study, we examine the optimal number of reduced QoI dimensions, such that torque signal reconstruction is sufficiently accurate compared to the original signals.

To identify the optimal level of dimension reduction, we first evaluate the mean absolute error (MAE)

$$\mathrm{MAE} = \frac{1}{M}\sum_{m=1}^{M}\left\|\boldsymbol{\tau}^{(m)} - \mathcal{R}^{-1}\left(\mathbf{r}^{(m)}\right)\right\|_1, \tag{26}$$

where $\mathcal{R} : \boldsymbol{\tau} \in \mathbb{R}^N \rightarrow \mathbf{r} \in \mathbb{C}^R$, $R < N$, as discussed in Sect. 3.1. To that end, we use a sample of $M = 2000$ torque signals, each containing $N = 120$ torque values τ_{β_n} evaluated at equidistantly distributed rotation angles $\beta_n \in [0°, 30°)$, $n = 0, \ldots, N-1$. The torque signals correspond to uniform random samples of the design space defined in Table 1, i.e., for sampling purposes, each parameter is assumed to be a random variable that is uniformly distributed within its value range, additionally assuming that the parameters are independent from one another. It should be noted that if $R = N$, the original torque signal can be fully recovered without information loss, equivalently, the reconstruction error should be equal to zero (bearing machine accuracy). Figure 4a shows the relationship between reduced frequency components and the associated reconstruction error, which may provide an indication as to how many frequency components should be retained to achieve an average reconstruction accuracy.

However, the MAE is not informative concerning worst-case reconstruction errors. This can be verified by the results shown in Fig. 4b, which shows the worst-case signal reconstruction for $R = 5$ frequency components. In this case, the MAE is equal to $3.2 \cdot 10^{-3}$, which should be sufficient for a reconstruction that is visually identical to the original signal, based on the torque values. Still, the reconstructed signal deviates significantly from the original, thus indicating that the retained frequency components are insufficient. The minimum number of frequency components to obtain a worst-case signal reconstruction that is visually identical to the original torque signal is $R = 11$, as can be observed in Fig. 4c. In this case, the MAE equals $9.0 \cdot 10^{-5}$. Based on these observations and numerical results, in the following, $R = 11$ frequency components will be retained after the DFT.

Remark on the selection of frequency components In principle, it is possible to determine the set of influential frequency components based on the electric machine's geometry and configuration, e.g., based on the number of phases, poles, and slots. However, this requires expert knowledge. Moreover, geometric and material uncertainties can cause the number of influential frequency components to rise significantly [51, 52], in which case the geometrical characteristics of the machine alone do not suffice. The suggested approach allows to select the important frequency components in an automatic manner, without resorting to expert knowledge.

4.2 Surrogate model accuracy

In this second numerical study, we assess the performance of surrogates obtained with the suggested framework in terms of prediction accuracy. Moreover, comparisons against alternative surrogates are performed, which use the same RSMs but either employ principal component analysis (PCA) [53] instead of DFT for dimension reduction or omit dimension reduction altogether. In the former case, 21 principal components are kept, resulting in an MAE equal to $2.8 \cdot 10^{-5}$ with respect to torque signal reconstruction. Note that the chosen number of principal components coincides with the number of reduced components obtained if the DFT-based, complex-valued, reduced QoI, $\mathbf{r} \in \mathbb{C}^R$, is transformed to an equivalent real-valued QoI, $\mathbf{r} \in \mathbb{R}^{2R-1}$, by considering separately the real and imaginary parts of the frequency components (see Sect. 3.1), for $R = 11$ (see Sect. 4.1).

For training and validation, an initial dataset $\mathcal{D} = \left\{\mathbf{p}^{(m)}, \boldsymbol{\tau}^{(m)} = \boldsymbol{\tau}\left(\mathbf{p}^{(m)}\right)\right\}_{m=1}^{M}$, $M = 2000$, is partitioned into a training dataset $\mathcal{D}_\mathrm{t}$ and a validation dataset $\mathcal{D}_\mathrm{v}$, where $\mathcal{D}_\mathrm{t} \cap \mathcal{D}_\mathrm{v} = \emptyset$. As previously discussed in Sect. 4.1, the dataset $\mathcal{D}$ is generated by means of uniform random sampling of the design parameters. The training dataset size, M_t, is progressively increased such that $M_\mathrm{t} \in \{600, 1200, 1800\}$, which allows to examine the effect of training data availability on surrogate modeling accuracy. Contrarily, the validation dataset remains fixed with size $M_\mathrm{v} = 200$. Given a validation sample $(\mathbf{p}, \boldsymbol{\tau}) \in \mathcal{D}_\mathrm{v}$, the absolute percentage error (APE) for a specific rotation angle β with corresponding torque value $\tau_\beta \in \boldsymbol{\tau}$ is computed as

$$\mathrm{APE} = \frac{\left|\tau_\beta - \hat{\tau}_\beta\right|}{\tau_\beta}, \tag{27}$$

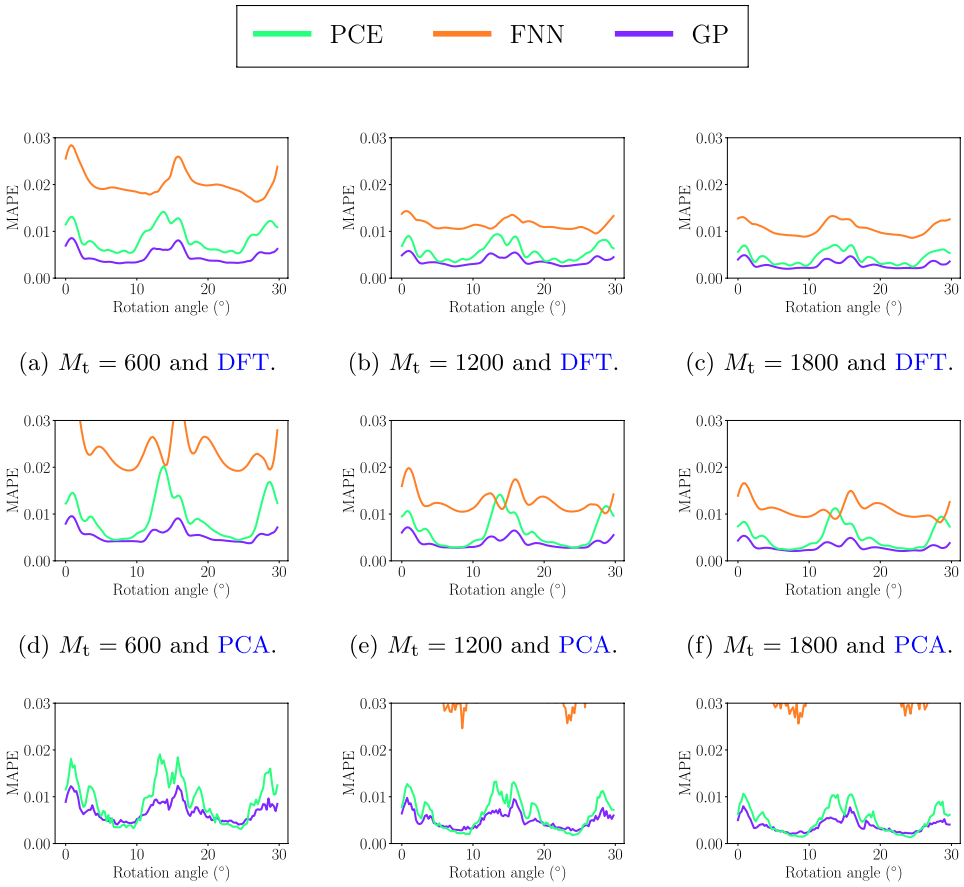

(a) $M_t = 600$ and DFT. (b) $M_t = 1200$ and DFT. (c) $M_t = 1800$ and DFT.

(d) $M_t = 600$ and PCA. (e) $M_t = 1200$ and PCA. (f) $M_t = 1800$ and PCA.

(g) $M_t = 600$ and no reduction. (h) $M_t = 1200$ and no reduction. (i) $M_t = 1800$ and no reduction.

Fig. 5 Mean absolute percentage error (MAPE) for surrogates computed with different combinations of training dataset size, RSM, and dimension reduction approach

where $\hat{\tau}_\beta$ denotes the surrogate-based torque estimate. Aggregating the APE over the validation dataset, the error mean (MAPE) and standard deviation (SDAPE) can be computed, where the latter indicates the robustness of the surrogate with respect to its predictive accuracy. Doing the same for all rotation angles β_n, $n = 0, \ldots, N-1$, the MAPE and SDAPE are computed over a full signal period.

The MAPE and SDAPE results are shown in Figs. 5 and 6, respectively, for all combinations of training dataset size, RSMs, and dimension reduction approach (or lack thereof). As would be expected, the accuracy and robustness of all surrogates is improved with more training data. The improvement is significant when increasing the training dataset size from $M_t = 600$ to $M_t = 1200$, but only minor gains are obtained for a further increase to $M_t = 1800$. Irrespective of training dataset size or dimension reduction approach, GPs are generally the best RSM option in terms of surrogate modeling accuracy and robustness, especially if combined with dimension reduction. On the opposite side, FNN is consistently the worst-performing RSM option. Using DFT for dimension reduction yields overall more accurate and robust surrogates compared to PCA or no dimension reduction, especially given smaller training datasets. The combination of DFT-based dimension reduction with a GP-based RSM yields the best-in-class surrogate model. These observations are further supported by the results presented in Table 3, where the MAPE is averaged over the torque signal.

Figure 7 shows the worst-case torque signal predictions for training dataset size $M_t = 1800$ and each combination of RSM and dimension reduction approach. The results for $M_t = 600$ and $M_t = 1200$ are available in Appendix (Figs. 12, 13). In all cases, it is clear that surrogates based on FNNs perform much worse than PCE and GP-based ones. The latter two RSMs do not result in significant differences and only a slight edge can be claimed for GP. Note that, in the case of either DFT or PCA-based dimension reduction, the observed deviations between original and predicted

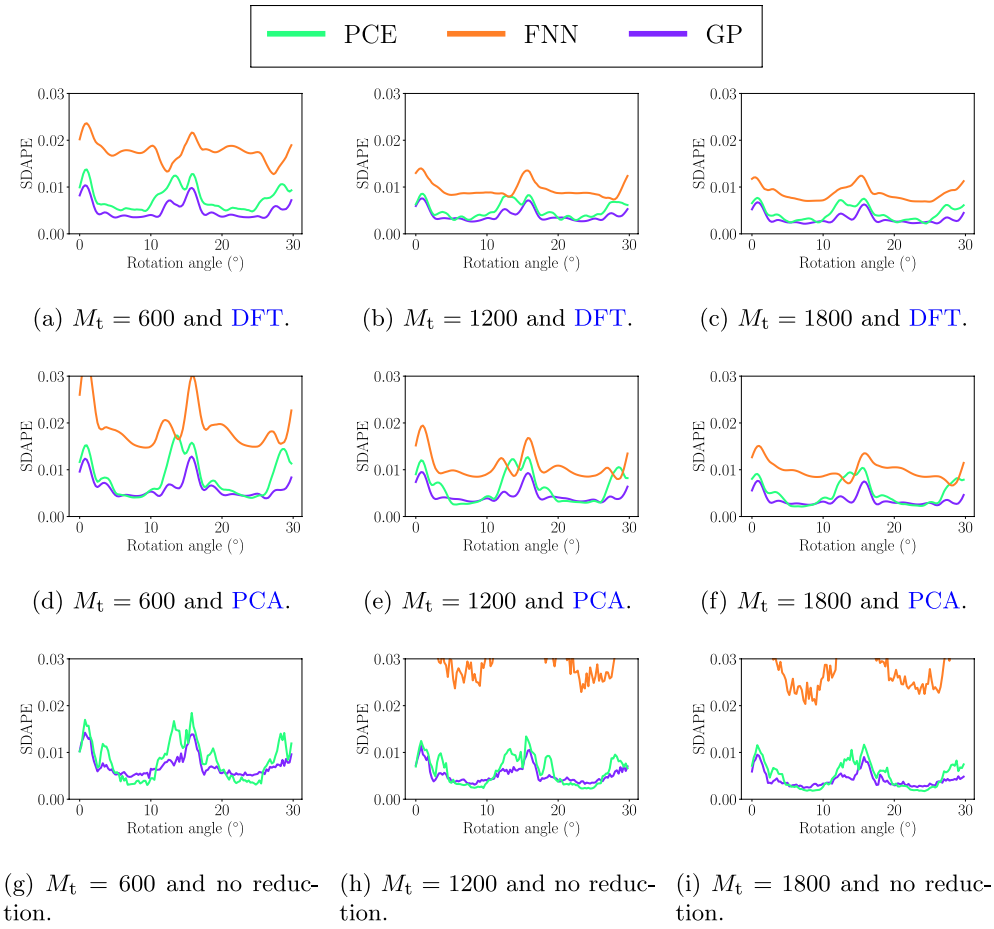

(a) $M_t = 600$ and DFT. (b) $M_t = 1200$ and DFT. (c) $M_t = 1800$ and DFT.

(d) $M_t = 600$ and PCA. (e) $M_t = 1200$ and PCA. (f) $M_t = 1800$ and PCA.

(g) $M_t = 600$ and no reduction. (h) $M_t = 1200$ and no reduction. (i) $M_t = 1800$ and no reduction.

Fig. 6 Standard deviation of absolute percentage error (SDAPE) for surrogates computed with different combinations of training dataset size, RSM, and dimension reduction approach

Table 3 Signal-averaged MAPE for surrogates computed with different combinations of training dataset size, RSM, and dimension reduction approach

	DFT			PCA			No reduction		
	PCE	FNN	GP	PCE	FNN	GP	PCE	FNN	GP
600	0.0085	0.0203	**0.0047**	0.0091	0.02407	0.0055	0.0091	0.0526	0.0068
1200	0.0057	0.0114	**0.0037**	0.0063	0.0127	0.0039	0.0063	0.0379	0.0049
1800	0.0044	0.0106	**0.0029**	0.0051	0.0110	0.0030	0.0050	0.0371	0.0038

The best-in-class error values for the same training dataset size are marked as bold

torque signal can be attributed solely to training data availability and RSM approximation capability (see Sect. 4.1).

Last but not least, it is worth noting that surrogate modeling accuracy and robustness clearly benefit from dimension reduction, both on average and in worst-case. This is particularly evident for the FNN, but actually applies to all RSMs. The enhancing effect of dimension reduction to the predictive accuracy of surrogate models is in fact known, e.g., considering time-series prediction [54] or deep operator learning [55]. For the use-case considered in this work, it is obvious both DFT and PCA provide this benefit, with a slight edge given to the former.

4.3 Surrogate-based uncertainty quantification

In this last numerical study, surrogates obtained with the suggested framework are used for UQ. In particular, they replace the original, high-fidelity PMSM model in Monte Carlo sampling, which aims at estimating torque statistics, specifically means and standard deviations of torque signals. Additionally, surrogate-based estimates are

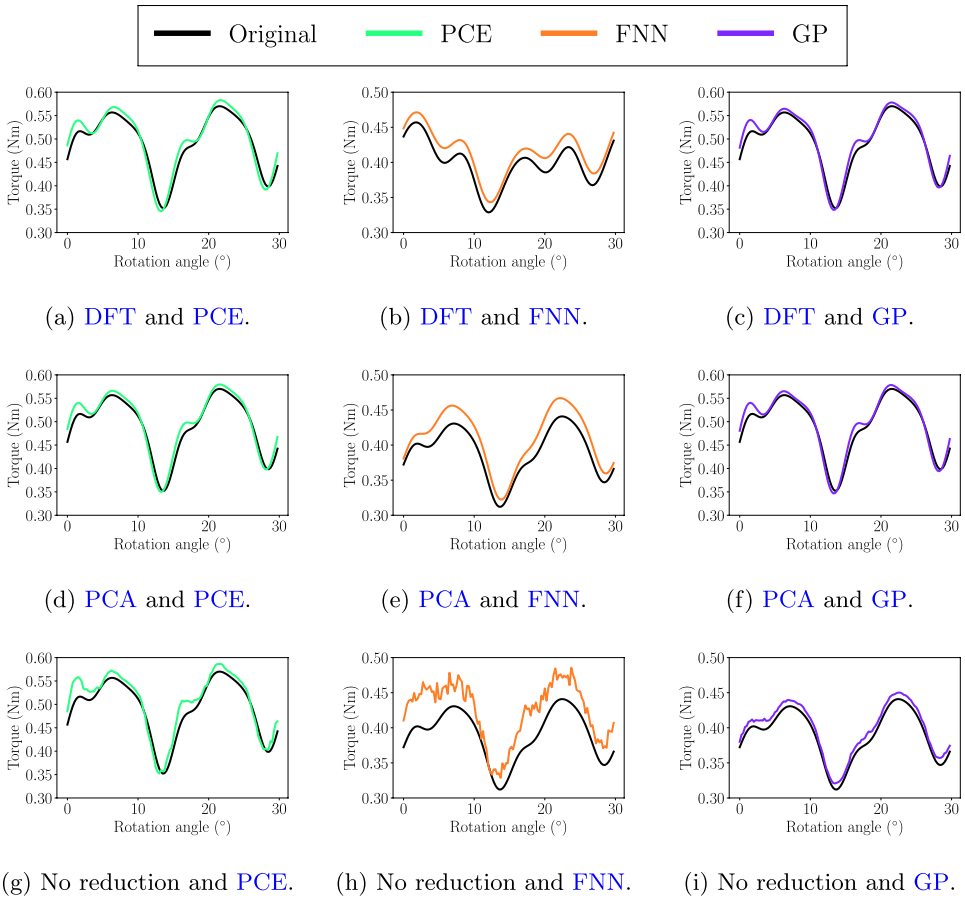

(a) DFT and PCE. (b) DFT and FNN. (c) DFT and GP.

(d) PCA and PCE. (e) PCA and FNN. (f) PCA and GP.

(g) No reduction and PCE. (h) No reduction and FNN. (i) No reduction and GP.

Fig. 7 Worst-case surrogate-based torque signal predictions for training dataset size $M_t = 1800$ and different combinations of RSM and dimension reduction approach

compared against reference statistics obtained by sampling the original model. The model, either original or surrogate, is sampled 10^4 times based on the considered parameter distributions. The sample mean and standard deviation of a torque value τ_β are denoted with $\hat{\mu}[\tau_\beta]$ and $\hat{\sigma}[\tau_\beta]$, respectively. In the following, we consider two distinct UQ scenarios concerning the parameter distributions, namely, uniform distributions in Sect. 4.3.1, and Gaussian distributions in Sect. 4.3.2. Section 4.3.3 discusses the computational costs and gains of the proposed method in the context of the aforementioned UQ studies.

Remark on modeling assumptions and UQ As previously noted in Sect. 2, only one-fourth of a two-dimensional cross-section of the PMSM is actually modelled, by assumption of rotational symmetry and translational invariance along the z-axis. In the UQ studies performed in the following, this assumption still holds, so that the PMSM model described in Sects. 2.1, 2.2, and 2.3 can be used. Nonetheless, we should note that this cannot be true in practice. For example, in a real-world setting, geometrical variations due to manufacturing tolerances will not be symmetrical.

4.3.1 Uniform parameter distributions

In the first UQ study, uniform distributions within the limits given in Table 1 are assumed for all design parameters. Figure 8 presents the sample mean and standard deviation of the torque signal for both the original and the surrogate models. Figure 9 shows the APE for the mean and standard deviation estimates relative to the reference statistics, i.e., the ones obtained by sampling the high-fidelity model. Additionally, signal-averaged APEs per training dataset size and RSM are reported in Table 4.

With respect to mean estimates, it is observed that GP is consistently the best RSM option, followed by PCEs and FNNs, irrespective of training dataset size. This aligns with our findings in Sect. 4.2 regarding surrogate modeling accuracy. In the case of standard deviation, however, no RSM choice consistently outperforms the others. One important advantage of GP is the comparatively smaller error fluctuation and the typically smaller maximum error values, with the only exception being the case for $M_t = 1200$ shown in Fig. 9e.

4.3.2 Gaussian parameter distributions

In the second UQ study, we assume a fixed nominal value for each parameter, which is equal to the middle value of its range as

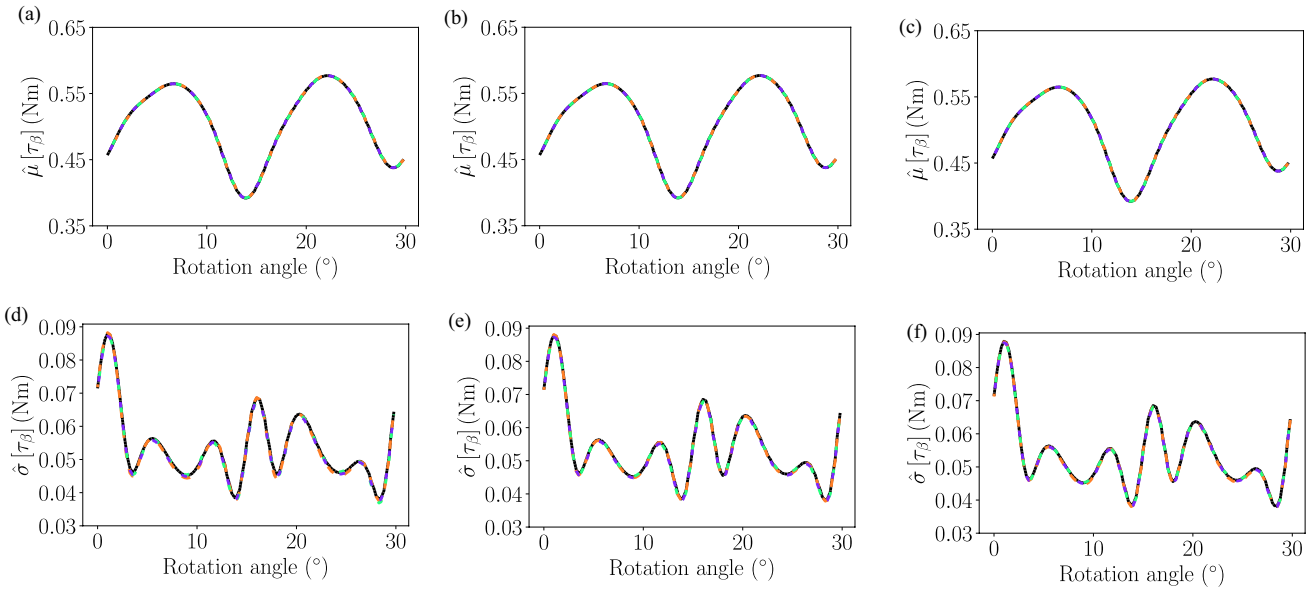

Fig. 8 Torque mean and standard deviation estimates for uniform parameter distributions, obtained by sampling the original, high-fidelity model and the framework-based surrogates, where the latter are based on different combinations of training dataset size and RSM

given in Table 1. Additionally, we consider an additive Gaussian noise term per parameter, with zero mean and standard deviation equal to 1% of the nominal parameter value. Similar to the case of uniform parameter distributions discussed in Sect. 4.3.1, we report sample means and standard deviations for both the original and the surrogate models in Fig. 10, while the corresponding APEs of the surrogate-based statistics estimates are shown in Fig. 11. Signal-averaged APEs per training dataset size and RSM are reported in Table 5.

Compared to the results reported in previous section for the case of uniform parameter distributions, important differences can be observed. First, the mean and standard deviation estimates are now visually discernible from the corresponding references, which was not the case before. Moreover, the errors presented in Fig. 11 and in Table 5 are significantly higher, in fact by approximately one order of magnitude. This deterioration in statistics estimation accuracy can be mainly attributed to the difference between the training and sampling distributions,

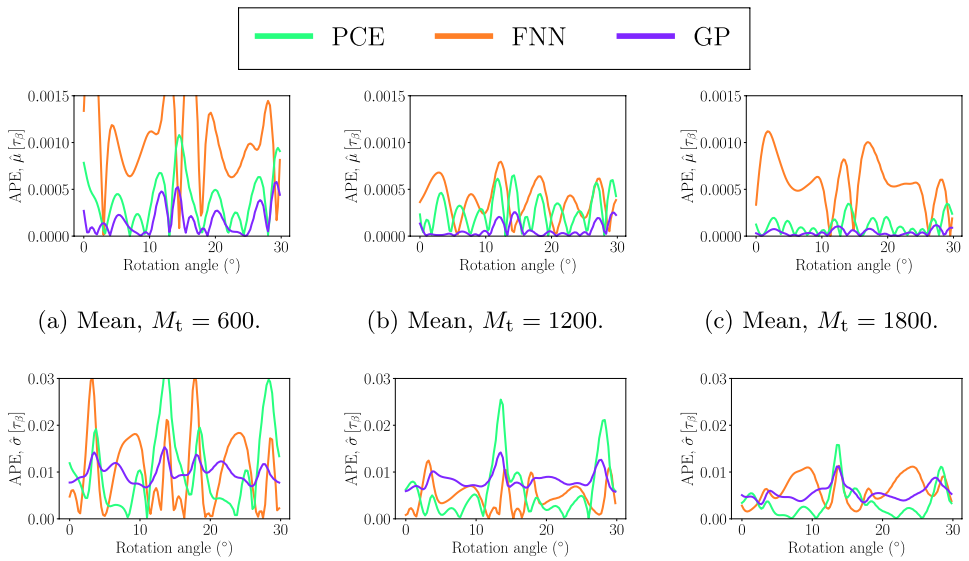

(a) Mean, $M_\mathrm{t} = 600$. (b) Mean, $M_\mathrm{t} = 1200$. (c) Mean, $M_\mathrm{t} = 1800$.

(d) Standard deviation, $M_\mathrm{t} = 600$. (e) Standard deviation, $M_\mathrm{t} = 1200$. (f) Standard deviation, $M_\mathrm{t} = 1800$.

Fig. 9 APE of torque mean and standard deviation estimates for uniform parameter distributions, obtained with framework-based surrogates computed with different combinations of training dataset size and RSM

Table 4 Signal-averaged APE of surrogate-based mean and standard deviation estimates for uniform parameter distributions

M_t	Mean			Standard deviation		
	PCE	FNN	GP	PCE	FNN	GP
600	0.0004	0.0011	**0.0002**	0.0106	0.0105	**0.0100**
1200	0.0003	0.0004	**0.0001**	0.0060	**0.0048**	0.0083
1800	0.0001	0.0006	**$4{\cdot}10^{-5}$**	**0.0035**	0.0063	0.0056

The surrogates are computed using the suggested framework, for different training dataset sizes and RSMs. The best-in-class error values for the same training dataset size are marked as bold

since the surrogate models are trained on uniform data and are then sampled for Gaussian parameter distributions. This *distributional shift*, as is the more common term in the context of ML, is known to deteriorate the performance of data-driven models in terms of their predictive ability [56]. Naturally, the reduced surrogate modeling accuracy affects the quality of the statistics estimates as well.

Additionally, looking at Table 5, the PCE seems to be the best RSM option concerning statistics estimation, with the only exception being the standard deviation estimate for $M_t = 1800$ training data points, in which case the FNN prevails. However, the signal-averaged errors can be quite misleading, as revealed by Fig. 11. In fact, the better signal-averaged errors of the PCE can be explained by very low errors in some parts for the signal, however, they are accompanied by large error fluctuations and high maximum errors. Once again, the GP yields consistently the lowest maximum errors and significantly less error fluctuation, and is clearly the best option in almost all cases. The only possible exception is the standard deviation estimate for $M_t = 1800$ training data points, in which case the FNN could also be selected.

4.3.3 Computational costs and gains

Last, we underline the computational gains achieved by using surrogate models for Monte Carlo-based statistics estimates. First, we note that a single high-fidelity PMSM model evaluation takes approximately 300 seconds on an up-to-date CPU. Accordingly, a Monte Carlo-based UQ study based on 10^4 random samples costs approximately 833 CPU-hours. In contrast, a surrogate-based UQ study using the same number of random samples can be performed in just 10 seconds, as the surrogate model is evaluated in 10^{-3} seconds at most.

However, in the latter case, one must also consider the so-called offline costs due to data acquisition and RSM training. Generating training datasets with sizes $M_t \in \{600, 1200, 1800\}$ amounts to (approximately) 50, 100, and 150 CPU-hours, respectively. The cost due to ML regression depends on the RSM and the training dataset size, as shown in Table 6. Even in the worst case, ML regression accounts for less than 1% of the total computational cost. Then, depending on the training dataset size, the computational cost of the surrogate-based UQ study is 16.7× – 5.5×

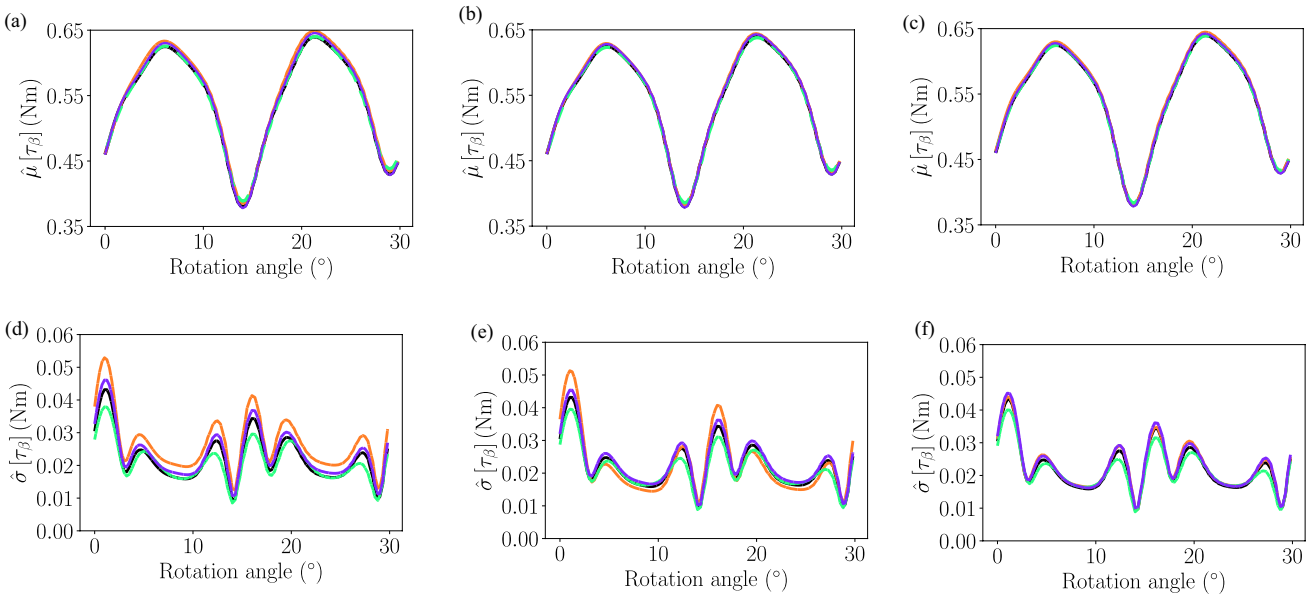

Fig. 10 Torque mean and standard deviation estimates for Gaussian parameter distributions, obtained by sampling the original, high-fidelity model and the framework-based surrogates, where the latter are based on different combinations of training dataset size and RSM

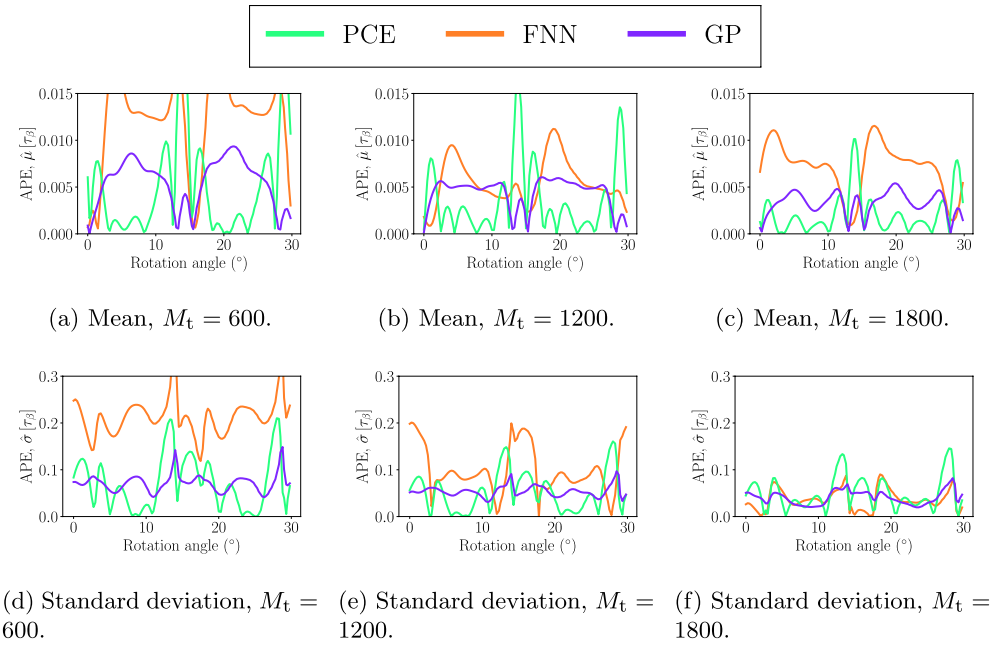

(a) Mean, $M_t = 600$. (b) Mean, $M_t = 1200$. (c) Mean, $M_t = 1800$.

(d) Standard deviation, $M_t = 600$. (e) Standard deviation, $M_t = 1200$. (f) Standard deviation, $M_t = 1800$.

Fig. 11 APE of torque mean and standard deviation estimates for Gaussian parameter distributions, obtained with framework-based surrogates computed with different combinations of training dataset size and RSM

Table 5 Signal-averaged APE of surrogate-based mean and standard deviation estimates for Gaussian parameter distributions

M_t	Mean			Standard deviation		
	PCE	FNN	GP	PCE	FNN	GP
600	**0.0052**	0.0125	0.0055	**0.071**	0.2140	0.0721
1200	**0.0037**	0.0055	0.0044	**0.0516**	0.0998	0.0541
1800	**0.0022**	0.0071	0.0033	0.0496	**0.0365**	0.0405

The surrogates are computed using the suggested framework, for different training dataset sizes and RSMs. The best-in-class error values for the same training dataset size are marked as bold

smaller compared to using the high-fidelity model, with statistics errors that can be considered more than acceptable.

Note that the computational gains could possibly be even greater, for example, considering scenarios with smaller training datasets or larger Monte Carlo samples. Moreover, it must be noted that the surrogate model is trained only once and can then be sampled for different distributions, as long as these distributions respect the given parameter ranges. That is, if new parameter distributions are to be considered, the data acquisition cost remains the same and only the marginal cost of sampling the surrogate model is added. This comes at the cost of possible deterioration in surrogate modeling accuracy, as shown in Sect. 4.3.2. On the contrary, in case of utilizing the original model within the UQ study, each change in the parameters' distributions would necessitate new costly model evaluations.

Table 6 Training times (in seconds) per RSM and training dataset size

M_t	PCE	FNN	GP
600	675	27	450
1200	2030	41	1780
1800	5100	51	4800

5 Conclusions

This work presented a novel, data-driven surrogate modeling framework, aiming at predicting the torque of a PMSM under geometric design variations cost-effectively, and utilizing these predictions for cost-efficient UQ studies. The framework consists of a reduced-order modeling and an inference part. The reduced-order modeling part combines DFT-based dimension reduction of torque signals with regression-based RSMs that approximate the functional dependency between the PMSM's design parameters and the reduced QoI. Three RSMs are tested, namely, PCE, FNN, and GP. In the inference part, the trained RSM is first evaluated for a set of new design parameters and subsequently an inverse DFT is applied to yield the predicted torque signal.

Our framework is applied for torque inference and UQ for a PMSM with 20 geometric design parameters. The torque signals contain 120 torque values evaluated upon equidistant rotation angles within a period. The numerical studies reveal a number of interesting observations. First, at least $R = 11$ frequency components must be kept in the ROM to ensure a sufficiently accurate torque signal reconstruction. Next, the best-in-class surrogate model in terms of prediction accuracy and robustness is obtained by combining DFT-based dimension reduction and GP-based RSMs. This result includes comparisons against non-framework-based surrogate models, which either replace DFT with PCA in the dimension reduction step or omit it altogether. Importantly, omitting the dimension reduction step leads to severe deterioration in prediction accuracy. In surrogate-based UQ by means of Monte Carlo sampling, the GP remains the best RSM choice concerning mean torque estimation accuracy. For standard deviation estimation, no RSM can claim best performance, however, GP offers the advantages of smaller maximum errors and less error fluctuation. Last, taking into account both the offline and online costs of the suggested framework, that is, due to data acquisition, ML regression, and model evaluation, surrogate-based UQ studies by means of Monte Carlo are several times more cost-efficient compared to sampling the original, high-fidelity PMSM model.

Overall, the suggested framework results in surrogate models that can reliably replace the high-fidelity PMSM model and enable otherwise computationally intractable parameter studies such as UQ, even outperforming commonly used approaches like direct QoI approximation or dimension reduction via PCA. The DFT-based dimension reduction is particularly suitable in this use case, given the physical characteristics of the QoI, mainly its periodicity. Periodicity is indeed present for a wide array of electric machine QoIs besides the torque. The method remains applicable for non-periodic QoIs as well, however, the extent of this applicability should be investigated further. Other extensions could consider a richer selection of RSM and dimension reduction possibilities, including nonlinear methods.

Appendix: Worst-case surrogates

See Figs. 12 and 13.

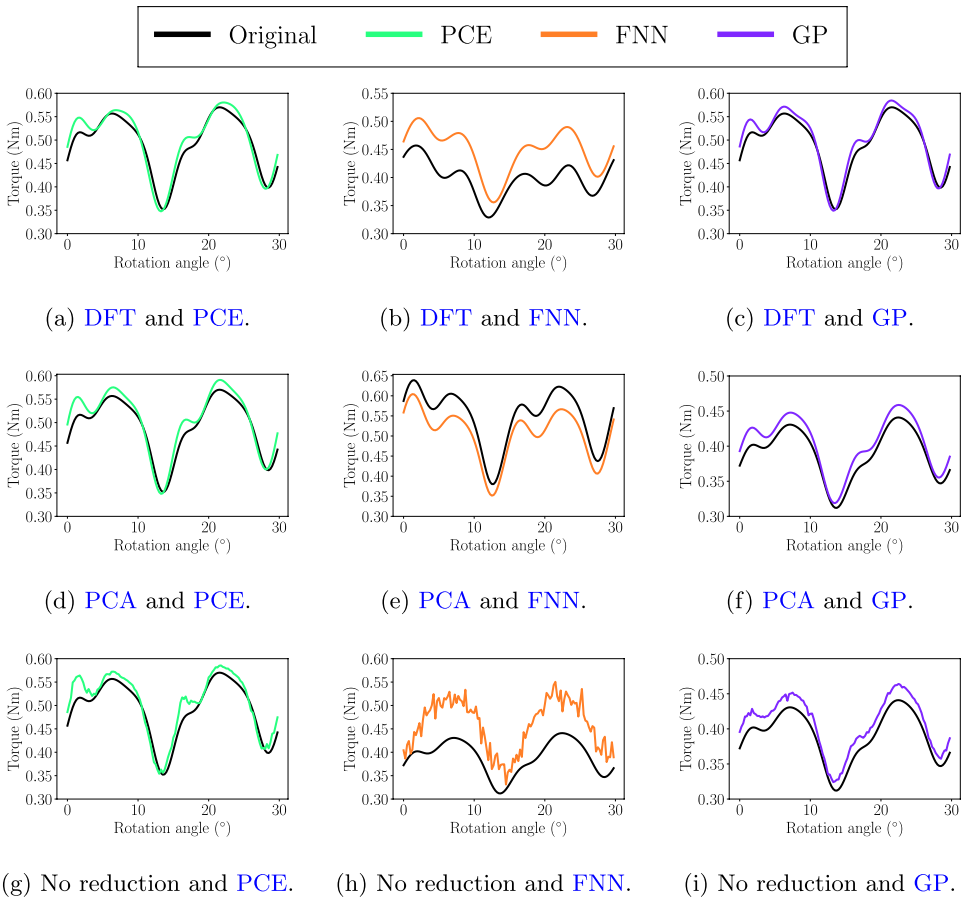

(a) DFT and PCE. (b) DFT and FNN. (c) DFT and GP.

(d) PCA and PCE. (e) PCA and FNN. (f) PCA and GP.

(g) No reduction and PCE. (h) No reduction and FNN. (i) No reduction and GP.

Fig. 12 Worst-case surrogate-based torque signal signal predictions for training dataset size $M_t = 600$ and different combinations of RSM and dimension reduction approach

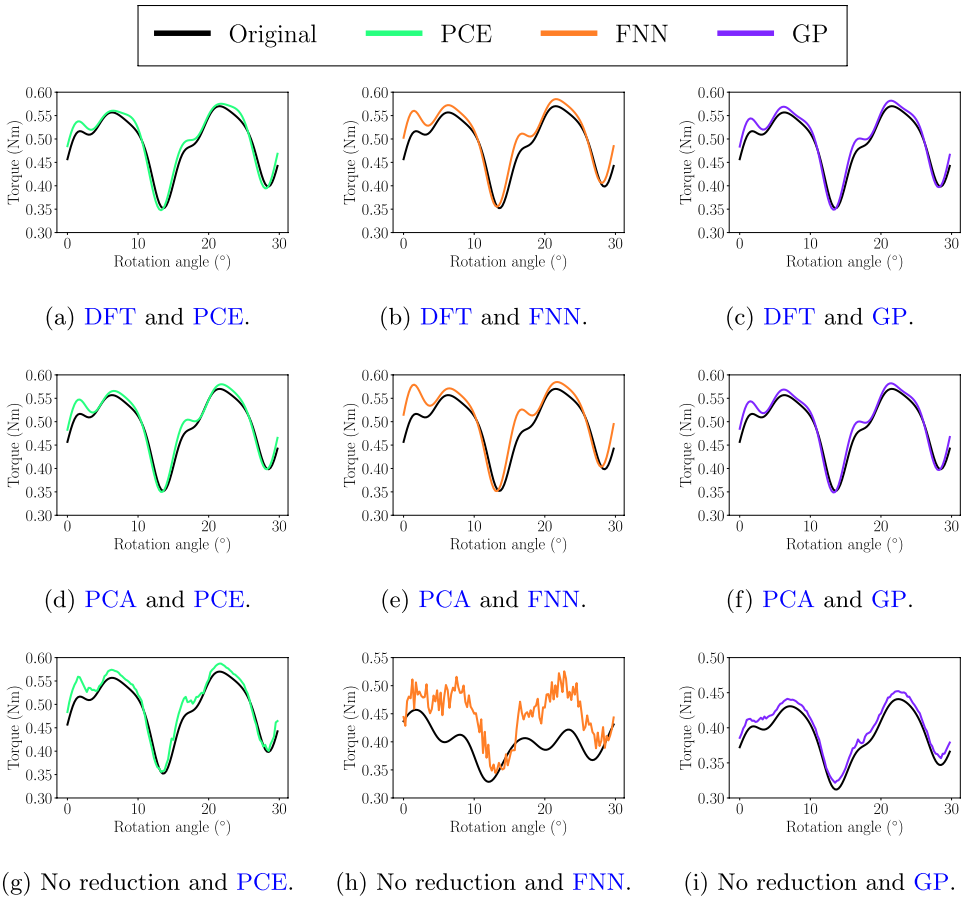

(a) DFT and PCE. (b) DFT and FNN. (c) DFT and GP.

(d) PCA and PCE. (e) PCA and FNN. (f) PCA and GP.

(g) No reduction and PCE. (h) No reduction and FNN. (i) No reduction and GP.

Fig. 13 Worst-case surrogate-based torque signal predictions for training dataset size $M_t = 1200$ and different combinations of RSM and dimension reduction approach

Acknowledgements The authors are partially supported by the joint DFG/FWF Collaborative Research Centre CREATOR (DFG: Project-ID 492661287/TRR 361; FWF: 10.55776/F90) at TU Darmstadt, TU Graz, and JKU Linz. The authors would like to thank Prof. Dr.-Ing. Herbert De Gersem for the fruitful exchanges during the revision of this work.

Author contributions A.P. implemented the software and performed all simulations and numerical studies. A.P. and D.L. analysed the numerical results. D.L. and S.Sc. conceived the original idea, acquired funding, and provided supervision. All authors contributed to writing and reviewing the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no competing interests.

References

1. Nadel S (2019) Electrification in the transportation, buildings, and industrial sectors: a review of opportunities, barriers, and policies. Current Sustainable/Renewable Energy Reports 6(4):158–168
2. Bramerdorfer G (2020) Effect of the manufacturing impact on the optimal electric machine design and performance. IEEE Trans Energy Convers 35(4):1935–1943
3. Jyrkama MI, Pandey MD (2016) On the separation of aleatory and epistemic uncertainties in probabilistic assessments. Nucl Eng Des 303:68–74
4. Smith RC (2024) Uncertainty quantification: theory, implementation, and applications. SIAM, Philadelphia
5. Zhang J (2021) Modern Monte Carlo methods for efficient uncertainty quantification and propagation: a survey. Wiley Interdiscip Rev Comput Stat 13(5):1539
6. Alizadeh R, Allen JK, Mistree F (2020) Managing computational complexity using surrogate models: a critical review. Res Eng Des 31(3):275–298

7. Kleijnen JPC (2015). In: Fu MC (ed) Response surface methodology. Springer, New York, pp 81–104
8. Hosder S (2012) Stochastic response surfaces based on non-intrusive polynomial chaos for uncertainty quantification. Int J Math Model Numer Optim 3(1–2):117–139
9. Thakre U, Mote RG (2022) Uncertainty quantification and statistical modeling of selective laser sintering process using polynomial chaos based response surface method. J Manuf Process 81:893–906
10. Beheshti Nezhad H, Miri M, Ghasemi MR (2019) New neural network-based response surface method for reliability analysis of structures. Neural Comput Appl 31:777–791
11. Xu Z, Zhang X, Wang S, He G (2022) Artificial neural network based response surface for data-driven dimensional analysis. J Comput Phys 459:111145
12. Costa NR, Lourenço J (2016) Gaussian process model - an exploratory study in the response surface methodology. Qual Reliab Eng Int 32(7):2367–2380
13. Chen C, Long J, Chen W, Liu Z, Guo J (2023) Modeling and prediction of spindle dynamic precision using the Kriging-based response surface method with a novel sampling strategy. Nonlinear Dyn 111(1):559–579
14. Benner P, Gugercin S, Willcox K (2015) A survey of projection-based model reduction methods for parametric dynamical systems. SIAM Rev 57(4):483–531
15. Chinesta F, Huerta A, Rozza G, Willcox K (2016) Model order reduction. Encyclopedia of computational mechanics, pp 1–36
16. Alla A, Kutz JN (2017) Nonlinear model order reduction via dynamic mode decomposition. SIAM J Sci Comput 39(5):778–796
17. Nouy A (2017). In: Ghanem R, Higdon D, Owhadi H (eds) Low-rank tensor methods for model order reduction. Springer, Cham, pp 857–882
18. Cicci L, Fresca S, Guo M, Manzoni A, Zunino P (2023) Uncertainty quantification for nonlinear solid mechanics using reduced order models with Gaussian process regression. Comput Math Appl 149:1–23
19. Drakoulas GI, Gortsas TV, Bourantas GC, Burganos VN, Polyzos D (2023) FastSVD-ML-ROM: a reduced-order modeling framework based on machine learning for real-time applications. Comput Methods Appl Mech Eng 414:116155
20. Conti P, Guo M, Manzoni A, Frangi A, Brunton SL, Nathan Kutz J (2024) Multi-fidelity reduced-order surrogate modelling. Proc R Soc A 480(2283):20230655
21. Wiesheu M, Komann T, Merkel M, Schöps S, Ulbrich S, Cortes Garcia I (2024) Combined parameter and shape optimization of electric machines with isogeometric analysis. Optim Eng 20:1–28
22. Hughes TJ, Cottrell JA, Bazilevs Y (2005) Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. Comput Methods Appl Mech Eng 194(39–41):4135–4195
23. Hou CKJ, Behdinan K (2022) Dimensionality reduction in surrogate modeling: a review of combined methods. Data Sci Eng 7(4):402–427
24. Cheng M, Zhao X, Dhimish M, Qiu W, Niu S (2024) A review of data-driven surrogate models for design optimization of electric motors. IEEE Trans Transp Electrif 2024:1
25. JSOL Corporation: JMAG-RT Model Library. https://www.jmag-international.com/modellibrary/
26. Wiesheu M (2023) MotorOptimization. https://doi.org/10.5281/zenodo.10160086
27. Vázquez R (2016) A new design for the implementation of isogeometric analysis in Octave and Matlab: GeoPDEs 3.0. Comput Math Appl 72(3):523–554
28. Piegl L, Tiller W (2012) The NURBS book. Springer, Berlin Heidelberg
29. Buffa A, Vázquez RH, Sangalli G, Veiga L (2015) Approximation estimates for isogeometric spaces in multipatch geometries. Numer Methods Partial Differ Equ 31(2):422–438
30. Egger H, Harutyunyan M, Löscher R, Merkel M, Schöps S (2022) On torque computation in electric machine simulation by harmonic mortar methods. J Math Ind 12(1):6
31. Bontinck Z, Corno J, Schöps S, De Gersem H (2018) Isogeometric analysis and harmonic stator-rotor coupling for simulating electric machines. Comput Methods Appl Mech Eng 334:40–55
32. Meeker DC (2019) Finite element method magnetics, version 4.2 (28feb2018 build). http://www.femm.info
33. Bontinck Z, De Gersem H, Schöps S (2015) Response surface models for the uncertainty quantification of eccentric permanent magnet synchronous machines. IEEE Trans Magn 52(3):1–4
34. Offermann P, Mac H, Nguyen TT, Clénet S, De Gersem H, Hameyer K (2015) Uncertainty quantification and sensitivity analysis in electrical machines with stochastically varying machine parameters. IEEE Trans Magn 51(3):1–4
35. Galetzka A, Bontinck Z, Römer U, Schöps S (2019) A multilevel Monte Carlo method for high-dimensional uncertainty quantification of low-frequency electromagnetic devices. IEEE Trans Magn 55(8):1–12
36. Beltran-Pulido A, Aliprantis D, Bilionis I, Munoz AR, Leonardi F, Avery SM (2020) Uncertainty quantification and sensitivity analysis in a nonlinear finite-element model of a permanent magnet synchronous machine. IEEE Trans Energy Convers 35(4):2152–2161
37. Ion IG, Bontinck Z, Loukrezis D, Römer U, Lass O, Ulbrich S, Schöps S, De Gersem H (2018) Robust shape optimization of electric devices based on deterministic optimization methods and finite-element analysis with affine parametrization and design elements. Electr Eng 100(4):2635–2647
38. Lei G, Bramerdorfer G, Ma B, Guo Y, Zhu J (2020) Robust design optimization of electrical machines: multi-objective approach. IEEE Trans Energy Convers 36(1):390–401
39. Bramerdorfer G (2020) Multiobjective electric machine optimization for highest reliability demands. CES Trans Electr Mach Syst 4(2):71–78
40. Xiu D, Karniadakis GE (2002) The Wiener–Askey polynomial chaos for stochastic differential equations. SIAM J Sci Comput 24(2):619–644
41. Feinberg J, Eck VG, Langtangen HP (2018) Multivariate polynomial chaos expansions with dependent variables. SIAM J Sci Comput 40(1):199–223
42. Lüthen N, Marelli S, Sudret B (2021) Sparse polynomial chaos expansions: literature survey and benchmark. SIAM/ASA J Uncertain Quantif 9(2):593–649
43. Blatman G, Sudret B (2011) Adaptive sparse polynomial chaos expansion based on least angle regression. J Comput Phys 230(6):2345–2367
44. Marelli S, Sudret B (2014) UQLab: a framework for uncertainty quantification in Matlab. In: Vulnerability, uncertainty, and risk: quantification, mitigation, and management, pp 2554–2563
45. Bottou L, Curtis FE, Nocedal J (2018) Optimization methods for large-scale machine learning. SIAM Rev 60(2):223–311
46. Baydin AG, Pearlmutter BA, Radul AA, Siskind JM (2018) Automatic differentiation in machine learning: a survey. J Mach Learn Res 18(153):1–43
47. Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M et al (2016) Tensorflow: large-scale machine learning on heterogeneous distributed systems. arXiv preprint arXiv:1603.04467
48. Williams CK, Rasmussen CE (2006) Gaussian processes for machine learning, vol 2. MIT Press, Cambridge
49. Hansen N, Ostermeier A (2001) Completely derandomized self-adaptation in evolution strategies. Evol Comput 9(2):159–195

50. Liu H, Cai J, Ong Y-S (2018) Remarks on multi-output Gaussian process regression. Knowl-Based Syst 144:102–121
51. Gandhi A, Corrigan T, Parsa L (2010) Recent advances in modeling and online detection of stator interturn faults in electrical motors. IEEE Trans Ind Electron 58(5):1564–1575
52. Handgruber P, Stermecki A, Biro O, Ofnery G (2013) Evaluation of interlaminar eddy currents in induction machines. In: IECON 2013-39th annual conference of the IEEE Industrial Electronics Society. IEEE, pp 2792–2797
53. Abdi H, Williams LJ (2010) Principal component analysis. Wiley Interdiscip Rev Comput Stat 2(4):433–459
54. Verleysen M, François D (2005) The curse of dimensionality in data mining and time series prediction. In: International work-conference on artificial neural networks. Springer, pp 758–770
55. Kontolati K, Goswami S, Karniadakis GE, Shields MD (2024) Learning nonlinear operators in latent spaces for real-time predictions of complex dynamics in physical systems. Nat Commun 15(1):5101
56. Mougan C, Broelemann K, Masip D, Kasneci G, Thiropanis T, Staab S (2023) Explanation shift: how did the distribution shift impact the model? arXiv preprint arXiv:2303.08081