Global and Local Sensitivity Analysis Applied to Electric Motor Design by Polynomial Chaos Expansion

Global and Local Sensitivity Analysis Applied to Electric Motor Design by Polynomial Chaos Expansion

Luca Giaccone, *Senior Member, IEEE*, Gianmarco Lorenti, *Member, IEEE*, Maurizio Repetto, and Luigi Solimene, *Member, IEEE*
Politecnico di Torino, Department of Energy "Galileo Ferraris", Torino, Italy, 10129

The design of electric motors increasingly requires the integration of multi-physics simulations, multi-objective optimization, and uncertainty quantification. In this work, we present a framework for global and local sensitivity analysis of a V-shaped Interior Permanent Magnet (IPM) motor using surrogate modeling and Polynomial Chaos Expansion (PCE). A neural network surrogate is trained on a dataset of high-fidelity finite element simulations to predict torque and temperature as functions of eight design parameters. Global sensitivity analysis is performed via PCE, enabling the computation of Sobol' indices and statistical metrics with negligible computational cost. The surrogate-based PCE results are validated against finite element-based model evaluations, confirming their accuracy and convergence. A local analysis is also performed along the Pareto front, where input parameters are perturbed within a narrow range to estimate confidence intervals around optimal solutions. This enables the quantification of robustness and highlights trade-offs between performance and sensitivity. The proposed approach offers an efficient method for integrating uncertainty-aware decision-making into the early stages of electric motor design.

Index Terms—**Electric motor design, Multi-objective optimization, Sensitivity Analysis, Polynomial Chaos Expansion, Surrogate modeling**

I. Introduction

SURROGATE models, also known as meta-models, in engineering are primarily designed to extract more information from analysis code runs. While a single simulation — such as one based on the Finite Element Method (FEM) — can be computationally expensive and provide accurate insight into a specific design configuration, engineers are often interested in mapping how output quantities respond across a broad region of the design space. Traditionally, surrogate modeling was used as a *curve fitting* technique, interpolating between configurations close in parameter space. Once trained, a surrogate model is significantly faster than full-order simulations, thus enabling the large-scale exploration of the design space and supporting tasks such as optimization and uncertainty quantification.

However, when surrogate models are integrated into optimization loops, the reliability of their predictions becomes a critical issue. Optimization routines often steer the search toward regions of the design space that are far from the initial samples, violating the assumptions underlying local interpolation and potentially reducing model accuracy. This has motivated the adoption of machine learning-based surrogates, which, unlike classical regression techniques, are better suited to generalization in high-dimensional spaces and can effectively capture nonlinear relationships across broad domains.

A robust surrogate model opens the door to modern design strategies, such as *Global Sensitivity Analysis* (GSA) and *Uncertainty Quantification* (UQ). GSA aims to determine how the uncertainty in each input parameter affects the variability of model outputs [1]. Among GSA techniques, *Polynomial Chaos Expansion* (PCE) is particularly attractive due to its spectral convergence, ability to produce analytical expressions for variance decomposition, and its compatibility with surrogate-based workflows [2], [3], [4]. Once a PCE model is constructed, statistical metrics and Sobol' sensitivity indices can be computed without the need for brute-force Monte Carlo simulations [5], [2], [3].

In recent years, PCE has been increasingly applied in many fields (e.g. design of complex energy systems, electromagnetics, bioelectromagnetics, circuit design) [6], [7], [8], However, its integration with neural surrogate models in the context of multi-physics electrical machine design remains an emerging topic. This paper contributes to this direction by presenting a comprehensive pipeline that combines neural network surrogates with PCE to analyze the design of a V-shaped Interior Permanent Magnet (IPM) motor.

The design problem considered here involves eight input parameters, covering both geometrical features and supply conditions. The quantities of interest are the *maximum torque* and *maximum temperature* under overload operation, obtained through high-fidelity FEM simulations. Artificial Neural Networks (ANNs) are trained on a dataset of 4096 simulations to reproduce these outputs with high accuracy [9].

In this study, we propose a methodology that integrates several elements into a unified framework, including:

- Validation of the surrogate-based PCE analysis against a reference PCE computed directly on the FEM dataset. This shows an excellent agreement between the results of the two methods in terms of mean, standard deviation, and Sobol' indices.
- Local sensitivity analysis close to the Pareto Front (PF) in the objective space and the corresponding Pareto Set (PS), obtained via multi-objective evolutionary optimization (NSGA-II) [10]. The robustness in the region close to PS is assessed by applying PCE on a $\pm 1\%$ perturbation neighborhood around each design point. This enables the definition of 95% confidence intervals for both objectives, revealing trade-offs between performance and robustness.

Corresponding author: Luigi Solimene (email: luigi.solimene@polito.it).

- A comparison between Sobol' indices and Pearson Correlation Coefficients (PCC) computed from the FEM dataset. The good agreement between the two independent methodologies answers the need for variance-based sensitivity measures in strongly coupled multi-physics problems.

The entire pipeline achieves a significant *computational speedup* (over 10^4-fold increase in speed compared to FEM analysis) making the approach practical for integration into design workflows. Moreover, the methodology allows for both global understanding and local refinement, thus supporting the identification of not only performant but also *robust* designs.

II. Dataset Generation and Surrogate Models

A. Dataset Generation

The design space considered in this work is defined by eight input variables, listed in Table I, which include both geometrical features of the stator and rotor as well as the current phase angle. These parameters were selected to cover a wide and realistic range of variation around a reference V-shaped Interior Permanent Magnet (IPM) motor topology [11].

TABLE I
Design parameters and variation ranges.

Parameter	**Range**	**Unit**
Barrier position d_α	[0.65, 0.85]	p.u.
Barrier width h_c	[0.3, 0.7]	p.u.
Rotor radius r	[60, 78]	mm
Tooth width w_t	[3.8, 6.3]	mm
Tooth length l_t	[15, 22.5]	mm
Slot opening w_o	[0.1, 0.4]	p.u.
Barrier shift $dxIB$	[–4, 6]	mm
Current phase angle γ	[30, 60]	°

Figure 1 illustrates the motor cross-section and highlights the seven geometrical design parameters considered in this study.

Sampling is a fundamental step in surrogate model building, especially in high dimensions, as the sampled points are the

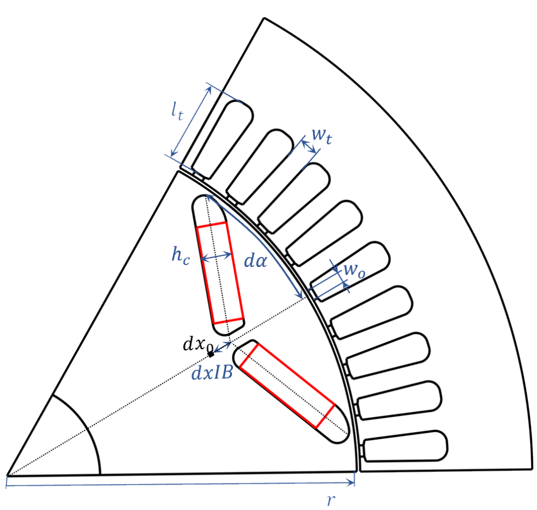

Fig. 1. Cross-section of the V-shaped IPM motor geometry with the seven geometrical design parameters [11].

hinges of the surrogate. Grid sampling systematically divides the space into equal intervals, ensuring regular coverage. However, its practicality decreases in high-dimensional spaces due to the exponential growth of the required number of points, a challenge known as *the curse of dimensionality* [12]. Sobol'-based quasi-random sampling, by contrast, uses deterministic low-discrepancy sequences that distribute points more uniformly across the space [5]. These sequences work by iteratively refining the placement of points through a mathematical process that ensures the points are spread evenly, thereby reducing the risk of large gaps or clusters. This method offers superior coverage compared to uniform distributions. The Sobol' sampling of the design parameter space allows to compute a dataset of points belonging to a p-dimensional real space $\mathbf{x} \in \mathcal{R}^p$ and to obtain a corresponding set of performance values by some analysis technique like FEM or others $\mathbf{y} \in \mathcal{R}^m$. With a Sobol'-based quasi-random sampling, a total of 4096 design configurations were generated. For each point, electromagnetic and thermal performance indicators were computed using finite element simulations, providing two key outputs: the maximum electromagnetic torque and the hotspot temperature in the winding region under overload conditions. The finite element simulations were performed using an in-house multi-physics workflow. The whole dataset generation workflow, including design parameter definition, FEM setup, and sampling methodology, is described in detail in [11]. The overall computational cost of the whole FEM campaign was approximately 122 hours of computation time (corresponding to about 440 000 seconds) and was performed on a workstation equipped with an Intel Xeon Cascade Lake processor (32 logical cores) and 128 GB of RAM, highlighting the need for a fast surrogate model.

B. Surrogate Modeling

To approximate the nonlinear relationship between the 8-dimensional input space and the two outputs (torque and temperature), we developed separate artificial neural networks (ANNs), one for each output variable. The models follow a standard feed-forward architecture with fully connected layers and hyperbolic tangent activation functions. The ANNs were implemented in Python using the `Sequential` model of the Keras library [13].

The structure of the network is different for each model output. A search is conducted among neural network structures (number of layers, number of neurons per layer) to minimize the reconstruction error. The optimal network for torque computation consists of three layers, with $(24, 64, 1)$ neurons, resulting in 1881 trainable parameters. For temperature, the optimal network consists of four layers, with $(8, 16, 32, 1)$, resulting in 793 trainable parameters. Models were trained using 55% of the dataset for training, 15% for validation, and 30% for testing.

The performance of the surrogate models on the test set is summarized in Table II. The mean absolute percentage error (MAPE) was below 0.6% and 0.05% respectively for torque and temperature, with maximum errors remaining under 7% and 1%. These results confirm the suitability of the surrogate models for both global and local sensitivity analysis.

TABLE II
SURROGATE MODEL ERRORS ON TEST SET.

Metric	Torque (%)	Temperature (%)
Mean error	0.534	0.041
95th percentile	1.526	0.103
Maximum error	6.725	0.869

The total computational time required to evaluate all 4096 configurations using the surrogate models was approximately 0.6 seconds, corresponding to a computational gain of over four orders of magnitude with respect to the full-order FEM simulations. This speed-up makes the surrogate models ideal for tasks requiring repeated evaluations, such as PCE-based uncertainty quantification and multi-objective optimization procedures.

III. Polynomial Chaos Expansion and Global Sensitivity Analysis

A. Overview of Polynomial Chaos Expansion

Polynomial Chaos Expansion (PCE) is a non-sampling-based method to determine the propagation of uncertainty in a nonlinear system when there is probabilistic uncertainty in the system parameters [2]. Given a variable $X(\omega)$ with finite variance where ω describes the randomness of the parameter X, one classical approach to evaluate the effect of ω through the nonlinear model is to use iterative sampling methods such as Monte Carlo. However, PCE provides a more efficient alternative by approximating the output as a spectral expansion over a suitable polynomial basis. Different uncertainty distributions (normal, uniform, etc.), defined *a priori* on the physical knowledge of the phenomenon, lead to different polynomial bases; for instance, a Gaussian distribution of ω is treated by Hermite polynomials. The PCE analysis is then based on the following steps:

1) Polynomial selection: once the variable (parameter) random distribution is defined, the related polynomial type is chosen. In this case, a uniform distribution is assumed. Thus, the Legendre polynomials are selected.
2) PCE accuracy order: define the *order* d for the polynomial expansion and all the combinations that lead to the multivariate polynomials.
3) PCE coefficients computation: calculate the PCE coefficients through the projection method by numerically integrating the performance function multiplied times the polynomials belonging to the basis. The higher the *level* of the numerical integration rule, the higher the number of function evaluations.
4) PCE Sensitivity indices: once the PC Model Output is obtained, mean and standard deviation values can be computed as a whole and as sensitivity indices with respect to each single parameter.

B. Implementation Details

The PCE analysis was performed using a dedicated open-source Python library developed by the authors (https://github.com/giaccone/pce). This library had already been tested and validated in other scientific works [14], [7]. The expansion coefficients are computed through spectral projection of the output function onto the polynomial basis. Given the eight input parameters uniformly distributed over the ranges defined in Table I, Legendre polynomials were adopted. The expansion was truncated to second order, which was found to be sufficient to ensure convergence of the output statistics and sensitivity indices (details will be provided in Section III-E). The quadrature rule was refined to accurately project the surrogate response up to second-order polynomial accuracy in the 8-dimensional input space. Thanks to the low computational cost of the surrogate model, the PCE coefficients were efficiently computed, enabling rapid and non-intrusive sensitivity analysis.

C. Global Sobol' Indices

Figure 2 reports the first-order Sobol' indices computed via surrogate-based PCE for both torque and temperature. The obtained Sobol' indices can be interpreted in light of electric motor physics. Torque is strongly influenced by the rotor radius, barrier position, and current phase angle, as these parameters directly affect the magnetic flux path and the electromechanical conversion process. Thermal performance, on the other hand, is primarily driven by rotor radius, tooth dimensions, and slot opening, which affect the copper volume and, thus, the computation of ohmic losses and the heat dissipation. These findings are consistent with electric motor design guidelines and confirm that the PCE-based sensitivity analysis captures physically meaningful dependencies.

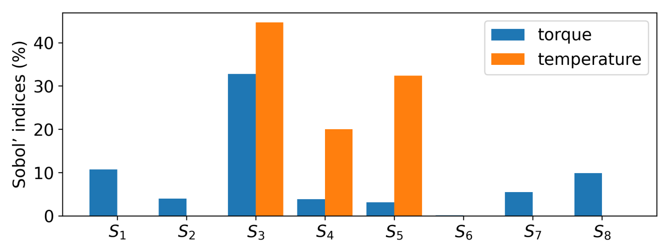

Fig. 2. First-order Sobol' indices for torque and temperature computed from surrogate-based PCE.

D. Comparison with Pearson Correlation Coefficients

As an additional interpretability tool, PCC were computed between each input parameter and the output quantities using the full FEM dataset. While these coefficients are straightforward to compute and often used in practice, they capture only linear dependencies.

Figure 3 shows the Pearson coefficients for torque and temperature, which can be compared with the first-order Sobol' indices in Figure 2. While the overall ranking of dominant variables is similar, the variance-based analysis provided by PCE reveals nonlinearities and interactions not visible from correlation analysis. Although only the first-order Sobol' indices are shown in Figure 2 for clarity, the full PCE model also enables the computation of higher-order indices, which reveal contributions due to interactions between input parameters. These contributions are absent in the PCC analysis.

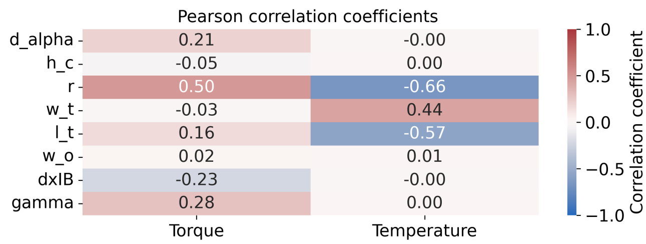

Fig. 3. Pearson correlation coefficients for torque and temperature.

E. Convergence Analysis and Validation Against FEM-Based PCE

The convergence behavior of the method was evaluated by comparing surrogate-based results with those obtained using the FEM-based approach. In particular, Figure 4 shows the evolution of the third Sobol' index S_3 with increasing PCE order and integration rule level. Similarly, Figures 5 and 6 show the convergence of the output mean and standard deviation, respectively, for both torque and temperature.

The convergence behavior can be explained by the fact that the system outputs are already well approximated by a quadratic expansion. Increasing the polynomial order increases the number of coefficients to be determined. If the number of quadrature points is not scaled accordingly, the projection of higher-order terms becomes noisy and may appear less accurate, as observed in Figures 4, 5, and 6. This does not imply that higher-order expansions are intrinsically worse, but rather that they provide no additional benefit when the true response is already well captured by a quadratic model. In practice, second-order PCE ensures both accuracy and stability with a limited computational cost and was therefore adopted throughout the study.

All metrics show rapid convergence, with negligible differences between surrogate- and FEM-based PCE results already at polynomial order $d = 2$ and refined integration. This confirms that, for the model under analysis, surrogate-based PCE achieves the same accuracy as full-model analysis at a fraction of the computational cost, making it highly suitable for global sensitivity analysis.

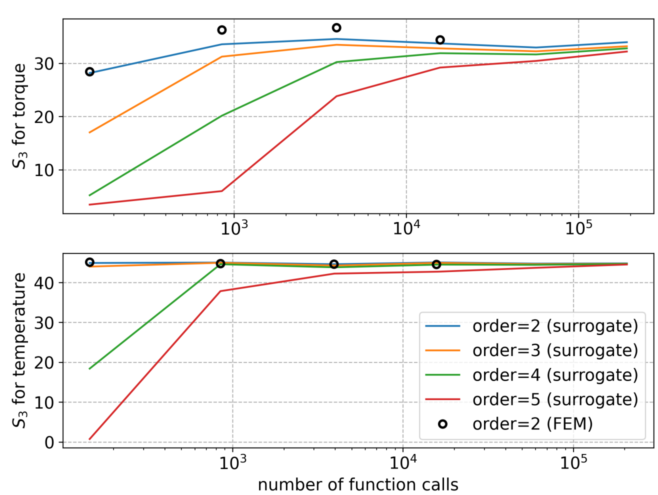

Fig. 4. Convergence of Sobol' index S_3 for torque as a function of PCE order and integration rule level. Comparison of convergence for surrogate-based and FEM-based approaches.

Fig. 5. Convergence of mean values and standard deviations for torque as a function of PCE order and integration rule level. Comparison of convergence for surrogate-based and FEM-based approaches.

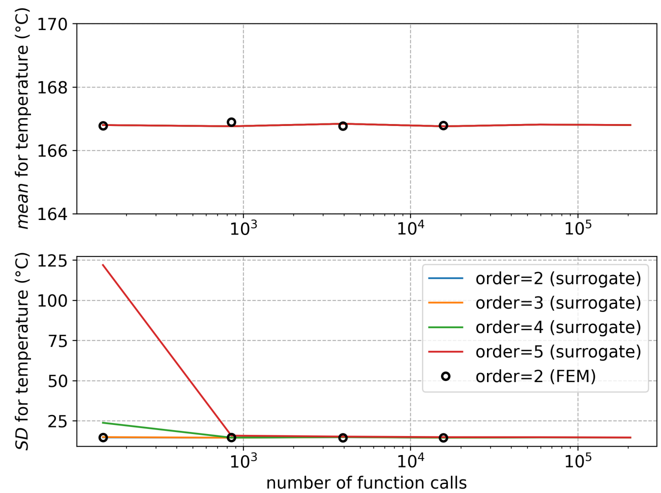

Fig. 6. Convergence of mean values and standard deviations for temperature as a function of PCE order and integration rule level. Comparison of convergence for surrogate-based and FEM-based approaches.

IV. Local Sensitivity Analysis on the Pareto Front

In addition to the global sensitivity analysis over the entire design space, a local sensitivity analysis was conducted to investigate the robustness of optimal solutions along the most interesting region for design, that is the Pareto front. PF represents the trade-off between two conflicting objectives: the maximization of torque and the minimization of temperature under overload conditions.

A. Pareto Front Identification

The PF was identified using the NSGA-II algorithm [10], executed over 10 independent runs using the `pymoo` Python library [15]. Since NSGA-II is formulated as a minimization algorithm, torque maximization is equivalently expressed as minimization of its negative value ($-\mathrm{torque}$). Each run generated a set of non-dominated solutions in the 8-dimensional input space. The union of these fronts was then processed to extract a refined, high-quality approximation of the overall PF.

To select a representative subset of points for local analysis, a distance-based filtering strategy was applied in objective space. Specifically, 48 points were selected to ensure coverage of the entire front. Each point was treated as a nominal design around which local variability was assessed.

B. Local PCE Setup and Confidence Interval Estimation

For each selected Pareto-optimal point, a local PCE analysis was performed. The input parameters were perturbed uniformly within a $\pm 1\%$ range around their nominal values, constrained to remain within the original design bounds. A second-order PCE model was then constructed around each point, using the surrogate model to evaluate the outputs.

This allowed the computation of local statistical descriptors—mean and standard deviation—for both objectives, as well as the definition of confidence intervals. In particular, the 95% confidence interval was computed for each point, providing insight into the expected variability in performance resulting from small input fluctuations.

C. Results and Discussion

Figure 7 provides a comprehensive visualization of the local variability around the Pareto-optimal designs. The black curve represents the PF, obtained from surrogate predictions. The red markers indicate the mean values of torque and temperature computed from the local PCE analysis at each selected point. As already mentioned, the torque values are shown as negative in the figure.

The light blue cloud corresponds to the ensemble of PCE evaluations sampled around all 48 Pareto points, while the orange points refer specifically to the evaluations generated around the 25th point, defined by sorting the solutions from left to right in the plot (i.e., by decreasing torque). For this point, the red open circle represents the expected (mean) output, while the red cross denotes the 95% confidence interval in both objectives.

The inset in Figure 7 highlights how small perturbations in the input space can cause the output distribution to deviate from the expected PF. In this case, the mean output (red circle) lies off the Pareto front, and the confidence interval indicates non-negligible uncertainty in both torque and temperature. This shows that some design points can be more robust than others.

This visualization highlights the benefit of incorporating local uncertainty quantification into multi-objective design. By identifying designs with small confidence regions and aligned mean outputs, designers can prioritize solutions that

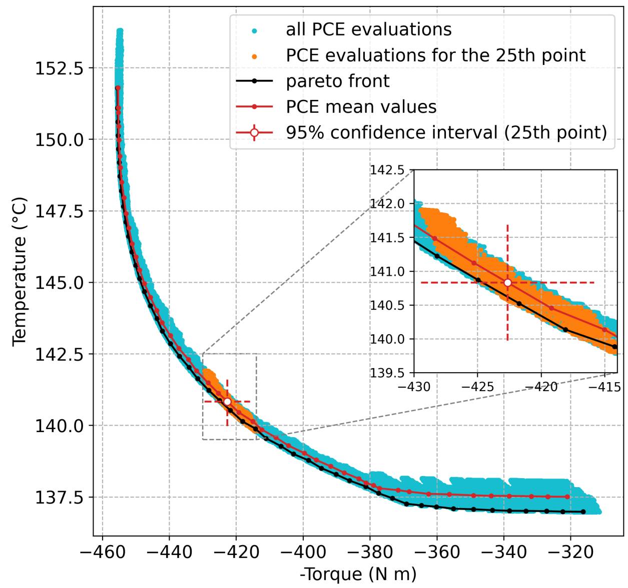

Fig. 7. Pareto front with local uncertainty information from PCE. Torque is shown as negative due to the minimization-based formulation.

are not only high-performing but also stable under parameter uncertainty.

It is also important to note that, although the confidence region around a given point may visually extend toward the utopia point (i.e., toward simultaneously higher torque and lower temperature), such combinations are not attainable through local input perturbations. The feasible outcomes remain constrained by the shape of the global PF, which acts as a natural performance boundary. This is evident in the figure, where the cloud of local evaluations never crosses the front, i.e. all points tried by PCE analysis lead to dominated solutions. In other words, uncertainty may lead to degraded performance, but improvements beyond the nominal Pareto point are not observed, highlighting the Pareto front's role as a frontier of best achievable trade-offs under the given design constraints.

Figure 8 shows the mean values and 95% confidence intervals for torque and temperature at each of the 48 selected Pareto-optimal points. The top panel refers to torque (plotted with a negative sign due to the minimization formulation), and the bottom panel to temperature. In both plots, the red lines indicate the mean values computed from the local PCE analysis, while the vertical error bars represent the corresponding confidence intervals. The points are indexed from left to right in order of decreasing torque. It can be observed that the amplitude of the confidence intervals varies significantly along the front. In general, configurations at the extremities of the front tend to exhibit smaller uncertainty, while intermediate points show more pronounced variability, particularly in torque. This visualization complements the global analysis by highlighting the variation in the magnitude of local output fluctuations across the Pareto front. The information provided by these confidence intervals can support further evaluation of trade-offs between nominal performance and sensitivity to

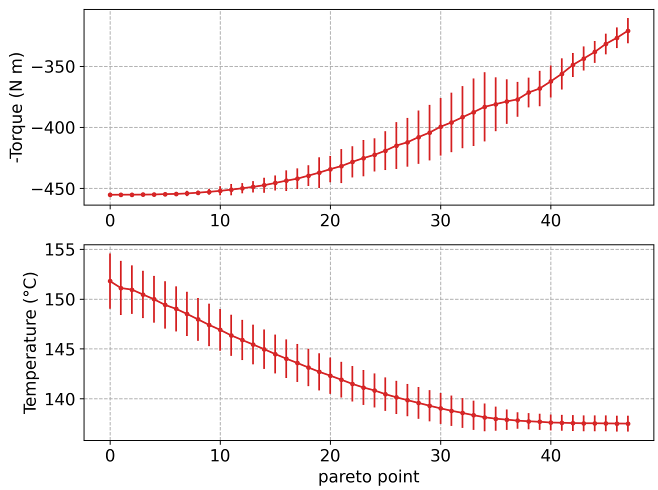

Fig. 8. Mean values and 95% confidence intervals for torque (top) and temperature (bottom) at each point along the Pareto front, computed via local PCE. Points are ordered by decreasing torque.

input variations.

This local sensitivity analysis enables a robustness-aware interpretation of the PF. Rather than selecting optimal points exclusively based on nominal performance, designers can incorporate confidence information into decision-making. For example, among multiple solutions with similar torque levels, the one with lower temperature uncertainty may be the preferred choice. Moreover, the methodology is computationally efficient thanks to the surrogate models: the local PCE analysis for all 48 points was completed with almost negligible additional cost. This makes the proposed approach suitable for integration into surrogate-assisted multi-objective optimization frameworks, where robustness and performance must be jointly evaluated.

V. Conclusion

This paper presented a comprehensive framework for global and local sensitivity analysis in the design of electric motors, based on surrogate modeling and Polynomial Chaos Expansion. Starting from a dataset of high-fidelity finite element simulations of a V-shaped IPM motor, artificial neural networks were trained to predict torque and temperature over a high-dimensional design space. The resulting surrogate models enabled efficient uncertainty quantification via PCE with negligible computational costs.

The global sensitivity analysis identified the most influential input parameters and validated the reliability of the surrogate-based approach by comparison with full-model PCE results. The analysis demonstrated that second-order expansions, combined with refined quadrature rules, were sufficient to achieve convergence in statistical indicators and Sobol' indices.

A local sensitivity analysis was also carried out along the PF of optimal solutions, where input parameters were perturbed within a narrow range around each design point. This enabled the estimation of confidence intervals for each objective and revealed the presence of non-uniform robustness across the front. In particular, while some designs exhibited limited variability, others showed significant output dispersion, especially in terms of torque. Although global and local analyses are presented separately, they are complementary. The global PCE identifies parameter importance across the entire design domain, while the local PCE, along the Pareto front, quantifies robustness around optimal solutions. Overall, the proposed methodology offers an efficient and accurate tool for embedding sensitivity and robustness considerations into the design process of electrical machines. Beyond the methodological contribution, the analysis provides physical insight into motor behavior. The identified links between geometry and performance, such as the effect of barrier position on torque and slot opening on thermal stress, demonstrate that the proposed framework is not only numerically efficient but also physically interpretable, thereby supporting engineers in making informed design choices under uncertainty. Future developments may extend the framework to include robustness-aware optimization strategies or Bayesian surrogate models for uncertainty-aware design refinement.

References

[1] A. Saltelli *et al.*, *Global Sensitivity Analysis for Importance Assessment*. John Wiley & Sons, Ltd, 2002, ch. 2, pp. 31–61. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/0470870958.ch2

[2] B. Sudret, "Global sensitivity analysis using polynomial chaos expansions," *Reliability Engineering & System Safety*, vol. 93, no. 7, pp. 964–979, jul 2008.

[3] M. Eldred, "Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design," in *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics, may 2009.

[4] A. Kaintura, T. Dhaene, and D. Spina, "Review of polynomial chaos-based methods for uncertainty quantification in modern integrated circuits," *Electronics*, vol. 7, no. 3, p. 30, feb 2018.

[5] I. M. Sobol', "On the distribution of points in a cube and the approximate evaluation of integrals," *Zh. Vychisl. Mat. Mat. Fiz.*, vol. 7, no. 4, pp. 784–802, 1967.

[6] E. Chiaramello *et al.*, "Assessment of fetal exposure to 4g LTE tablet in realistic scenarios: Effect of position, gestational age, and frequency," *IEEE Journal of Electromagnetics, RF and Microwaves in Medicine and Biology*, vol. 1, no. 1, pp. 26–33, jun 2017.

[7] L. Giaccone, "Uncertainty quantification in the assessment of human exposure to pulsed or multi-frequency fields," *Physics in Medicine & Biology*, vol. 68, no. 9, p. 095001, apr 2023. [Online]. Available: https://dx.doi.org/10.1088/1361-6560/acc924

[8] P. Manfredi and R. Trinchero, "Nonparametric formulation of polynomial chaos expansion based on least-square support-vector machines," *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108182, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0952197624003403

[9] L. Solimene *et al.*, "Data-driven approaches for electromagnetic analysis of traction electrical motors: A proposal for a benchmark problem," in *2024 International Conference on Electrical Machines (ICEM)*, 2024, pp. 1–7.

[10] K. Deb *et al.*, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.

[11] S. Ferrari *et al.*, "A multiphysics dataset generation procedure for the data-driven modeling of traction electric motors," *IEEE Access*, vol. 13, pp. 54 534–54 546, 2025.

[12] F. Y. Kuo and I. H. Sloan, "Lifting the curse of dimensionality," *Notices of the American Mathematical Society*, vol. 52, no. 11, pp. 1320–1328, 2005.

[13] "Keras: Deep Learning for humans." [Online]. Available: https://keras.io/

[14] L. Giaccone, P. Lazzeroni, and M. Repetto, "Uncertainty quantification in energy management procedures," *Electronics*, vol. 9, no. 9, 2020. [Online]. Available: https://www.mdpi.com/2079-9292/9/9/1471

[15] J. Blank and K. Deb, "Pymoo: Multi-objective optimization in python," *IEEE Access*, vol. 8, pp. 89 497–89 509, 2020.