Surrogate-Assisted Many-Objective Optimization for Traction Electric Motor Design

Gianmarco Lorenti*, Luigi Solimene*, Maurizio Repetto*,

* Politecnico di Torino, Department of Energy *Galileo Ferraris* (DENERG), Torino, Italy (e-mail: name.surname@polito.it)

***Abstract*—This paper investigates surrogate-assisted many-objective optimization for traction electric motor design under computationally expensive multiphysics evaluations. The case study is a V-shaped interior permanent magnet motor described by eight design variables and assessed through electromagnetic, thermal, and structural metrics. We optimize four objectives (torque, power factor, copper mass, and magnet mass) under mechanical-stress and winding-temperature constraints. Starting from a public dataset, we train a multi-output multi-layer perceptron and embed it within NSGA-III to enable a large evaluation budget. Multiple optimization runs are used to create a best-so-far Pareto set, from which a compact, well-spread subset is selected via reference-direction proximity, which is then re-evaluated with the high-fidelity FEM workflow. Results show that, although FEM re-evaluation reveals some surrogate discrepancies, especially near constraint boundaries, the adopted workflow provides high-quality Pareto designs at a reduced computational cost.**

***Index Terms*—design optimization, electric traction motors, pareto optimization, multi-layer neural networks, multiphysics.**

I. Introduction

Traction electric motor design is inherently multidisciplinary: electromagnetic performance, thermal behavior, and mechanical integrity must be jointly satisfied under strict packaging and efficiency constraints [1]–[4]. Candidate designs are increasingly evaluated against multiple and often conflicting key performance indicators (KPIs), naturally leading to many-objective formulations. For instance, increasing torque usually requires higher current, increasing losses and temperature; higher speed amplifies rotor stresses, creating a trade-off between electromagnetic performance and structural integrity [5]. In practice, each design evaluation typically relies on computationally expensive multiphysics simulations, which make direct simulation-driven optimization challenging [6].

Furthermore, as the dimensionality of the objective space increases, achieving a well-distributed Pareto-front approximation typically requires a huge number of function evaluations [7]. Many-objective variants of classical multi-objective evolutionary algorithms incorporate mechanisms to explicitly promote diversity in objective space (e.g., via reference points [8]). However, when each evaluation relies on computationally expensive simulations, the required evaluation budget often remains prohibitive. Surrogate-assisted optimization addresses this limitation by learning computationally inexpensive models that approximate the mapping from design variables to multiphysics KPIs, enabling broader exploration with fewer high-fidelity evaluations [9], [10]. Earlier work relied on interpolation-driven techniques [11] while now machine learning is exploited [12], [13].

In this paper, we investigate surrogate-assisted many-objective optimization for the design of an interior permanent-magnet (IPM) traction motor, in which multiphysics performance indicators must be jointly traded off. We consider a parametrized motor design problem with four objectives and two constraints spanning electromagnetic performance, thermal behavior, mechanical constraints, and material usage [5]. To enable tractable search under limited simulation budgets, high-fidelity evaluations during optimization are replaced by a data-driven surrogate.

Rather than proposing a new optimization algorithm or surrogate model, this work focuses on the integration and validation of a practical workflow for surrogate-assisted many-objective motor design under coupled multiphysics constraints (summarized in Fig. 1). Specifically, a single multi-output multi-layer perceptron (MLP) predicting all KPIs is coupled with NSGA-III; multiple runs are aggregated into a best-so-far non-dominated set, and a representative subset is re-evaluated via FEM for high-fidelity validation and trade-off analysis.

The main contributions of this work are:

- a surrogate-assisted many-objective workflow for traction motor design based on six multiphysics KPIs;
- a single multi-output MLP used within NSGA-III to predict all KPIs;
- FEM re-evaluation of selected Pareto candidates for high-fidelity validation and trade-off analysis.

The remainder of the paper is organized as follows: Section II describes the motor design problem, Section III presents the workflow, Section IV reports the results, and Section V concludes the paper.

Code/data availability: the work uses an open dataset [14], while source code is freely available at https://github.com/cadema-PoliTO/SurrogateManyOpt-Motor.

II. Problem definition

We consider the design of an IPM traction motor representative of full-electric vehicle applications (rated torque 236 $\mathrm{N\,m}$, rated power 120 kW, maximum speed 15000 rpm at the rotor shaft) [5]. The following subsections introduce the motor parametrization, the multiphysics performance indicators, and the high-fidelity FEM dataset used in the subsequent analyses.

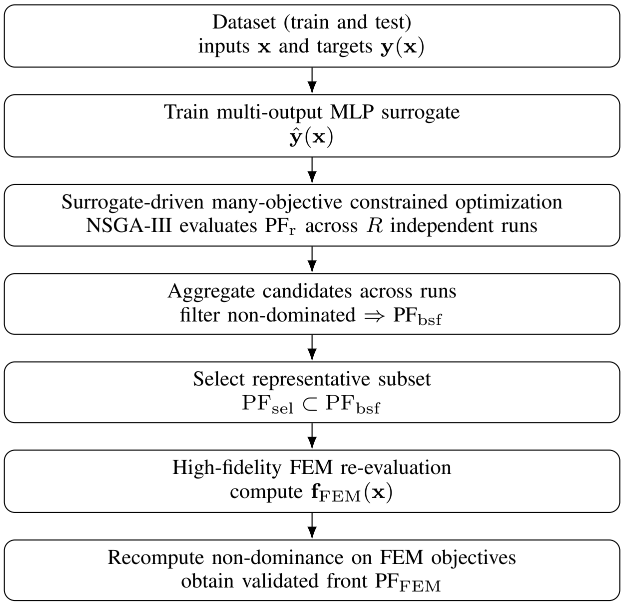

Fig. 1: Overview of the surrogate-assisted many-objective optimization workflow with multi-run aggregation and final FEM validation.

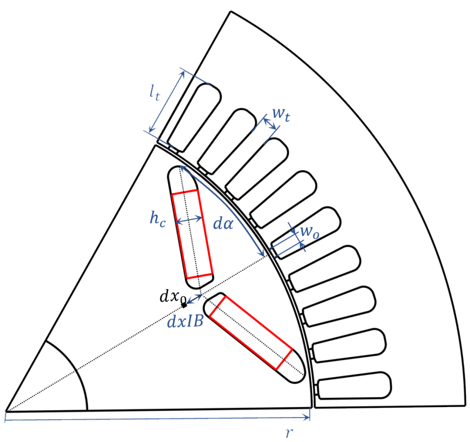

Fig. 2: Parametric definition of the V-shaped IPM traction motor geometry [5].

A. Parametric motor design

Candidate motor geometries are described through a set of continuous geometric parameters defining a single-barrier V-shaped rotor topology (see Fig. 2). Each candidate is also evaluated at a prescribed current phase angle, included among the design parameters. Overall, each design is specified by 8 parameters: seven geometric variables and one operating parameter (Table I). Geometric feasibility is enforced through space-related bounds, implemented by adapting the limits of the per-unit quantities in the table.

B. Multiphysics KPIs and dataset

The multiphysics performance of the motor is assessed through the KPIs in Table II: average torque T and power factor $\cos\varphi$ (electromagnetic); von Mises equivalent stress VM (structural); and maximum winding temperature Temp (thermal). Material usage is quantified by the copper mass m_{Cu} and the permanent-magnet mass m_{mag}.

A public dataset [14] provides high-fidelity multiphysics evaluations over the design space defined in Table I. The dataset consists of $2^{14} = 16384$ designs, generated via a Sobol' low-discrepancy (quasi-random) sampling of the design space, and the associated KPIs computed through a FEM-based multiphysics workflow (electromagnetic, thermal, and structural analyses) [4], [5], [15]. The dataset is used to train/test the surrogate and as a high-fidelity reference for the final FEM-based comparison.

III. Methodology

We adopt the workflow in Fig. 1, combining surrogate training/testing, surrogate-driven many-objective optimization, and final FEM re-evaluation of selected Pareto candidates.

A. Many-objective optimization

Let $\mathbf{x} \in \mathbb{R}^d$ denote the decision vector; we consider a constrained many-objective problem with m objectives and K inequality constraints:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{f}(\mathbf{x}) = \left[f_1(\mathbf{x}), \ldots, f_m(\mathbf{x})\right]^\top \\ \text{s.t.} \quad & \mathbf{x}^L \leq \mathbf{x} \leq \mathbf{x}^U \\ & g_k(\mathbf{x}) \leq 0, \quad k = 1, \ldots, K. \end{aligned} \tag{1}$$

In our case, the optimized objectives are $[-T, -\cos\varphi, m_{\mathrm{Cu}}, m_{\mathrm{mag}}]$ $(m = 4)$, with T and $\cos\varphi$ taken with negative sign as they are to be maximized. The lower and upper bounds of the decision variables ($\mathbf{x}^L$ and $\mathbf{x}^U$, respectively) assume the values shown in Table I. Feasibility is enforced by limiting mechanical stress and winding temperature:

$$g_1(x) = \mathrm{VM} - 450 \quad [\mathrm{MPa}]\,, \tag{2}$$

$$g_2(x) = \mathrm{Temp} - 180 \quad [^\circ\mathrm{C}]\,. \tag{3}$$

TABLE I: Parametric motor design space.

Symbol	Description	Min	Max	Unit
r	Rotor radius	60	78	mm
d_α	Barrier position	0.65	0.85	p.u.
h_c	Barrier width	0.30	0.70	p.u.
$d_{x,\mathrm{IB}}$	Barrier shift	-4	6	mm
w_t	Tooth width	3.8	6.3	mm
l_t	Tooth length	15.0	22.5	mm
w_o	Slot opening	0.10	0.40	p.u.
γ	Current phase angle	30	60	°

TABLE II: Multiphysics KPIs.

Symbol	Description	Unit
T	Average torque	N m
$\cos\varphi$	Power factor	–
m_{Cu}	Copper mass	kg
m_{mag}	Magnet mass	kg
VM	Von Mises equivalent stress	MPa
Temp	Max winding temperature	°C

The two constraints are aggregated into a constraint-violation score (CV), defined as follows:

$$\mathrm{CV}(\mathbf{x}) = \max(0, g_1(\mathbf{x})) + \max(0, g_2(\mathbf{x})) \tag{4}$$

A design is feasible when $\mathrm{CV}(\mathbf{x}) = 0$ and infeasible when $\mathrm{CV}(\mathbf{x}) > 0$.

a) Reference directions: To preserve diversity in the many-objective regime we adopt NSGA-III [8], which uses H reference directions $\{\mathbf{r}_h\}_{h=1}^{H}$ on the m-dimensional unit simplex. At each generation, objectives are normalized before associating individuals to directions. Ideal and nadir points estimated from the current population are denoted by $\mathbf{z}^*$ and $\mathbf{z}^{\mathrm{nad}}$ (objective-wise best and worst values across the non-dominated points); the normalized objectives are defined as:

$$\tilde{f}_j(\mathbf{x}) = \frac{f_j(\mathbf{x}) - z_j^*}{z_j^{\mathrm{nad}} - z_j^*}, \quad j = 1, \ldots, m. \tag{5}$$

Individuals in the last accepted front are selected based on the perpendicular distance from the directions in the normalized objective space, prioritizing underrepresented ones [8].

b) Multi-run aggregation and assessment: NSGA-III is repeated for R independent runs with different initial populations to reduce stochastic effects. Let PF_r denote the final feasible non-dominated set of run r. All run-wise fronts are pooled to obtain a global best-so-far approximation:

$$\mathrm{PF}_{\mathrm{bsf}} = \mathrm{ND}\left(\cup_{r=1}^{R} \mathrm{PF}_r\right), \tag{6}$$

where $\mathrm{ND}(\cdot)$ applies non-dominance among feasible solutions. Convergence and diversity are assessed through inverted generational distance (IGD), hypervolume (HV), and spacing on normalized objectives [8], using $\mathrm{PF}_{\mathrm{bsf}}$ as reference; we also report the per-run contribution to $\mathrm{PF}_{\mathrm{bsf}}$ ($N_{\mathrm{PF}_{\mathrm{bsf}}}$). Finally, a compact subset $\mathrm{PF}_{\mathrm{sel}}$ is obtained from $\mathrm{PF}_{\mathrm{bsf}}$ for high-fidelity FEM re-evaluation, by keeping only the closest point per reference direction.

B. Surrogate model

The expensive mapping $x \mapsto y(x)$ is approximated with a surrogate $\hat{y}(x)$, where y denotes the multiphysics KPIs used to construct the objective vector $f(x)$ and the constraint functions $g(x)$. A single multi-output MLP is trained to jointly predict all targets. The network has 8 input neurons, 6 output neurons, and a variable hidden-layer architecture selected during model selection.

a) Normalization: Since they span different ranges, inputs and targets are min–max scaled to $[0, 1]$ using statistics computed on the training set:

$$x_i^{\mathrm{s}} = \frac{x_i - x_i^{\min}}{x_i^{\max} - x_i^{\min}}, \qquad y_j^{\mathrm{s}} = \frac{y_j - y_j^{\min}}{y_j^{\max} - y_j^{\min}}. \tag{7}$$

b) Prediction error: While the model is trained by minimizing the MSE in normalized space, accuracy is reported through relative errors in physical units to ease interpretation across heterogeneous KPIs. Therefore, for a target y and its prediction $\hat{y}$, we report signed relative error (in %) as:

$$\varepsilon_{\%} = 100 \frac{\hat{y} - y}{y}. \tag{8}$$

C. High-fidelity FEM re-evaluation

Selected candidates $\mathrm{PF}_{\mathrm{sel}}$ are re-evaluated using the original FEM-based multiphysics workflow to obtain high-fidelity $\mathbf{y}(\mathbf{x})$. We then (i) re-check feasibility using the FEM constraints, (ii) recompute non-dominance in FEM objective space to obtain the validated set $\mathrm{PF}_{\mathrm{FEM}}$, and (iii) quantify surrogate discrepancies by computing (8) on Pareto candidates. Hence, the surrogate covers the full optimization budget, whereas high-fidelity FEM is required only for the high-fidelity samples used in this study and for the final validation subset.

IV. Results

A. Surrogate model training and accuracy assessment

a) Surrogate model setup: We used `keras` and `sklearn` to implement, train, and test the MLP. We used only 25% of the available dataset for training (4096 designs), for two reasons: (i) to assess the accuracy when training with fewer data (a realistic case when simulations are expensive), and (ii) to retain many previously unseen designs for a post-optimization comparison. Table III summarizes the main surrogate-model settings. The final MLP architecture was selected through a small grid search over the number of hidden layers and their widths, retaining the configuration that provided small train/test error with stable statistics.

b) Prediction errors: Table IV reports statistics of the error $\varepsilon_{\%}$ from (8) for each KPI, computed on the training set and the held-out test set. The close agreement between train and test error percentiles suggests that the selected configuration generalizes well over the available design space, with no evident overfitting. Overall, the electromagnetic/material objectives (T, $\cos\varphi$, m_{Cu} and m_{mag}) show limited dispersion. Temperature predictions remain highly accurate, while the

TABLE III: Main surrogate-model settings.

MLP architecture	
Input layer width	8 neurons
Hidden-layers widths	64, 48, 48 neurons
Output layer width	6 neurons
Hidden activation function	ReLU
Output activation function	Linear
Training	
Optimizer	Adam
Loss function	MSE
Batch size	32
Maximum epochs	2000
Validation split	10%

TABLE IV: Prediction error statistics over the dataset (in %).

		T	m_{Cu}	m_{mag}	$\cos\varphi$	VM	Temp
Train	$5\mathrm{p}^{\mathrm{th}}$	-1.20	-0.66	-1.33	-0.53	-4.60	-0.25
	median	-0.25	0.07	-0.22	0.17	0.20	-0.02
	$95\mathrm{p}^{\mathrm{th}}$	0.64	0.84	1.11	0.94	4.80	0.24
Test	$5\mathrm{p}^{\mathrm{th}}$	-1.23	-0.66	-1.33	-0.55	-4.82	-0.24
	median	-0.27	0.07	-0.24	0.15	0.12	-0.02
	$95\mathrm{p}^{\mathrm{th}}$	0.66	0.85	1.14	0.95	4.86	0.24

widest spread is observed for the mechanical stress. This is critical near feasibility boundaries, where surrogate discrepancies may change the feasibility status after FEM re-evaluation.

B. Optimization results

a) Optimization setup: The many-objective optimization was performed using `pymoo`'s NSGA-III implementation. The main optimization settings are summarized in Table V. Variation operators not explicitly specified in the setup were left to their default values (SBX crossover with $\eta_c = 30$ and $p_c = 1.0$, and polynomial mutation with $\eta_m = 20$). Although using more directions than individuals ($H > N$) is not recommended, as it may slow down convergence, preliminary tests with $H = N = 100$ showed poor coverage because several directions lie outside the projection of the feasible objective region. The value $H = 400$ was set by trial-and-error to ensure that about N directions are effectively covered, while the remaining ones correspond to infeasible regions.

b) Convergence and run-to-run variability: Fig. 3 shows the convergence history across the independent runs, in terms of IGD, which is computed against $\mathrm{PF_{bsf}}$. The plot reports the minimum, median, and maximum values at each generation: IGD drops quickly within the first 200 generations; beyond this point, the median curve remains close to a plateau, although it exhibits some fluctuations. The figure also reports the median of the population maximum CV across runs: CV quickly drops to zero, indicating that the population becomes fully feasible early, according to the surrogate. Table VI summarizes the final-generation performance in terms of $N_{\mathrm{PF_{bsf}}}$, IGD, HV (computed using Nadir as reference point), and spacing, showing relatively narrow indicator ranges.

c) Selected Pareto set: The best-so-far Pareto front $\mathrm{PF_{bsf}}$ includes 768 non-dominated solutions pooled from the final populations of the ten runs. To obtain a compact yet well-spread subset for high-fidelity re-evaluation, we built

TABLE V: Main NSGA-III optimization settings.

Independent runs R	10
Population size N	100
Reference directions H	400
Reference-direction method	Riesz s-energy method [16]
Maximum generations G	1000

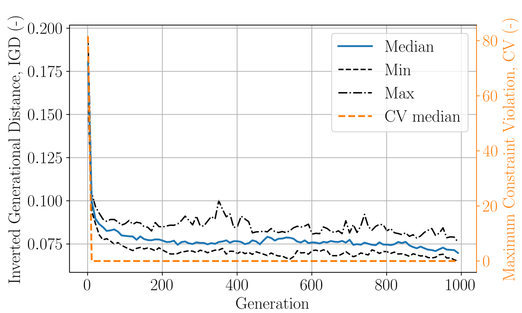

Fig. 3: NSGA-III convergence history (IGD and CV) over ten independent runs.

TABLE VI: Statistics of PF quality indicators across runs.

	$\mathrm{N_{PF_{bsf}}}$	IGD (-)	HV (-)	Spacing (-)
Worst run	68	0.081	0.465	0.107
Median	76	0.078	0.473	0.083
Best run	86	0.068	0.488	0.071

$\mathrm{PF_{sel}} \subset \mathrm{PF_{bsf}}$ by assigning each solution to its closest reference direction and retaining, for each direction, the candidate with minimum distance (i.e., the best-aligned solution), resulting in 101 points. Fig. 4 reports a parallel-coordinates plot (PCP) of $\mathrm{PF_{bsf}}$ (grey background) and $\mathrm{PF_{sel}}$ (blue), where each objective axis is scaled between the corresponding ideal and nadir values; accordingly, lower positions indicate better performance and higher positions worse performance.

C. High-fidelity validation

The selected subset $\mathrm{PF_{sel}}$ was re-evaluated using the high-fidelity FEM-based multiphysics workflow. Table VII reports the surrogate prediction error on the FEM re-evaluation of the selected candidates $\mathrm{PF_{sel}}$. Compared to Table IV, here the range error is wider, although median values remain close to zero for all KPIs; the largest increase is observed for VM.

As a consequence, 10 candidates were found infeasible after FEM re-evaluation due to violations of the VM constraint from (2). After updating the Pareto ranking using the FEM-evaluated objectives, one further design became dominated and was discarded, resulting in a final validated set $\mathrm{PF_{FEM}}$ comprising 90 feasible non-dominated solutions.

D. Post-optimization analysis

The set of FEM-validated solutions ($\mathrm{PF_{FEM}}$) allows to analyze relationships between design variables and multiphysics

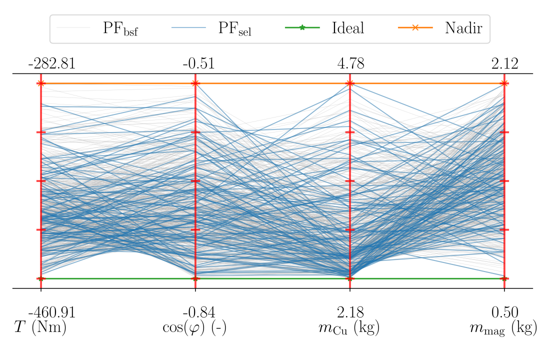

Fig. 4: PCP of the surrogate best-so-far Pareto set $\mathrm{PF_{bsf}}$ and the selected subset $\mathrm{PF_{sel}}$, with ideal and nadir points.

TABLE VII: Prediction error statistics over $\mathrm{PF_{sel}}$ (in %).

		T	m_{Cu}	m_{mag}	$\cos\varphi$	VM	Temp
$\mathrm{PF_{sel}}$	$5\mathrm{p^{th}}$	-1.02	-0.92	-1.82	-0.47	-5.59	-0.27
	median	0.15	-0.16	-0.39	0.13	1.21	0.01
	$95\mathrm{p^{th}}$	1.29	1.08	1.13	1.08	7.73	0.35

KPIs within Pareto-optimal configurations. We focus on trade-offs, proximity to feasibility limits, and recurring variable patterns across the Pareto set, in the spirit of *innovization* [17].

a) Trade-offs between objectives: Fig. 5 highlights the five best solutions for each objective to emphasize trade-offs that are difficult to discern in the full PCP. Maximizing torque generally correlates with higher active-material usage, although different allocations between copper and magnet mass can achieve similar torque levels. Maximizing $\cos\varphi$ is typically associated with higher magnet mass and lower copper mass (a trend that is not necessarily causal [5]). Moreover, minimizing copper mass appears less conflicting with the other objectives, especially the electromagnetic KPIs, whereas reducing magnet mass usually comes at the expense of electromagnetic performance.

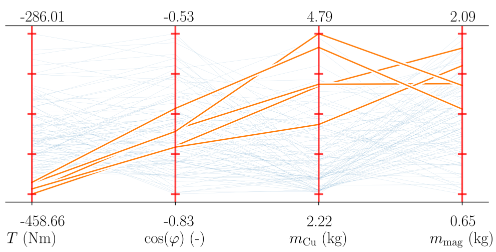

(a) best T

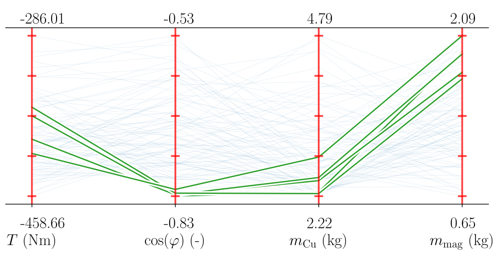

(b) best $\cos\varphi$

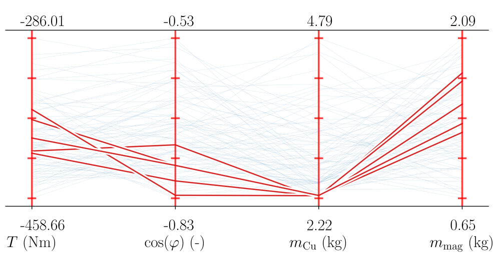

(c) best m_{Cu}

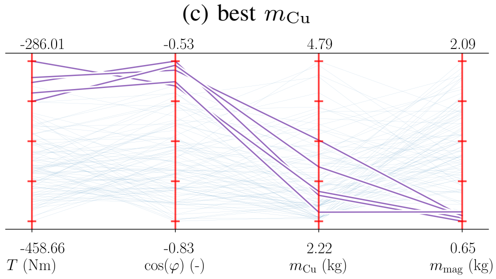

(d) best m_{mag}

Fig. 5: PCPs of the five best solutions for each objective, on top the full $\mathrm{PF}_{\mathrm{FEM}}$.

b) Constraints satisfaction: Fig. 6 shows VM and Temp distributions: both are skewed toward their limits, indicating that Pareto designs tend to operate close to feasibility boundaries. Ten designs exceed the threshold of VM, i.e., those discarded after FEM re-evaluation due to stress infeasibility.

c) Innovization insights: The distributions in Fig. 7 show that Pareto-optimal designs occupy a limited portion of the design space. Several variables are highly skewed towards the upper (e.g., w_{o} and γ), or lower bound (e.g., h_{c}) or, anyway, within a narrow sub-range (e.g., d_α). Conversely, variables with broader or multimodal distributions (e.g., dxIB and l_t) can be used to tune trade-offs among objectives.

E. Assessment against dataset

A full FEM-driven many-objective optimization would be required to rigorously benchmark the quality of the obtained Pareto set, but this would require $R \times G \times N = 10^6$ high-fidelity evaluations. Given that the FEM workflow required on the order 3 hours per 100 designs (in our setting), such a benchmark would imply a computational cost on the order of 3×10^4 hours, and is therefore impractical.

However, $\mathrm{PF}_{\mathrm{FEM}}$ can still be compared against the available high-fidelity dataset. Although the dataset is generated by space-covering sampling rather than optimization, it still provides a useful reference to assess whether the proposed procedure identifies more Pareto-relevant designs. We extract the locally non-dominated feasible designs from the training and test sets, denoted as $\mathrm{PF}_{\mathrm{train}}$ and $\mathrm{PF}_{\mathrm{test}}$, yielding 187 out of 4096 designs for train and 340 out of 12288 for test. We then merge $\mathrm{PF}_{\mathrm{FEM}}$, $\mathrm{PF}_{\mathrm{train}}$, and $\mathrm{PF}_{\mathrm{test}}$ and recompute non-dominance to define a common best-so-far front $\mathrm{PF}_{\mathrm{bsf}}$, used as reference in the indicator comparison.

Since front-quality indicators depend on the set size, we mimic the outcome of an optimization run with N individuals by uniformly subsampling N designs from $\mathrm{PF}_{\mathrm{train}}$ and $\mathrm{PF}_{\mathrm{test}}$ and computing IGD/HV/spacing with respect to $\mathrm{PF}_{\mathrm{bsf}}$, to keep the comparison consistent. As shown in Table VIII, $\mathrm{PF}_{\mathrm{FEM}}$ achieves consistently better indicator values than the size-matched dataset subsets, suggesting that the surrogate-assisted workflow identifies and validates solutions in more Pareto-relevant regions of the design space.

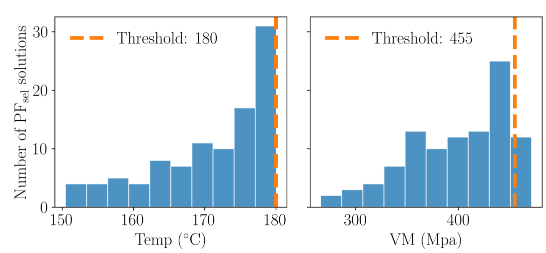

Fig. 6: Distribution of temperature and von Mises stress over FEM-validated selected candidates, with feasibility thresholds.

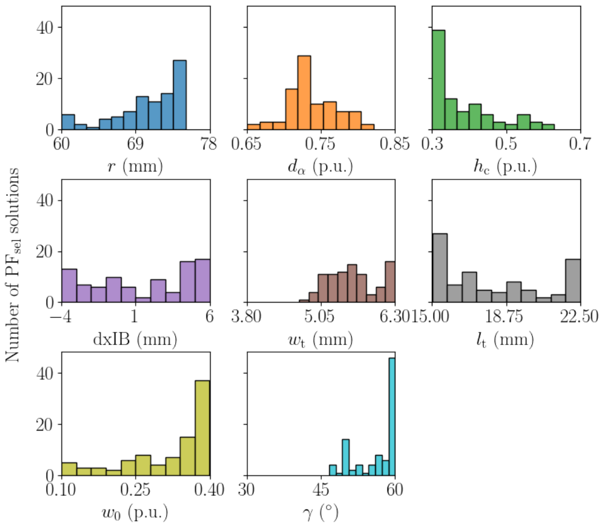

Fig. 7: Distributions of the design variables over FEM-validated selected candidates.

TABLE VIII: Quality indicators computed against the global Pareto front.

	IGD	HV	Spacing	$N_{\mathrm{PF}_{\mathrm{bsf}}}$
$\mathrm{PF}_{\mathrm{sel}}$	**0.071**	**0.575**	**0.063**	**87**
Test	0.077	0.503	0.087	59
Train	0.085	0.482	0.078	42

V. Conclusion

This work investigated a surrogate-assisted many-objective optimization workflow for traction electric motor design under computationally expensive multiphysics evaluations. Four conflicting objectives (torque, power factor, copper mass, magnet mass) were optimized under mechanical-stress and winding-temperature constraints. A multi-output MLP surrogate trained on a high-fidelity public dataset was embedded into NSGA-III, to enable a large optimization budget (10^6 evaluations) at negligible marginal cost. Multiple runs were aggregated into a best-so-far set, from which a compact and well-spread subset was selected via reference-direction proximity and then re-evaluated with the high-fidelity FEM workflow. The workflow relies on about 4200 high-fidelity FEM evaluations in this study (about 120 hours of FEM time), whereas a FEM-only optimization loop of the same scale would require on the order of 3×10^4 hours.

High-fidelity validation confirms the practical value of the adopted approach: out of 101 surrogate-selected candidates, 90 designs remained feasible and non-dominated, while 10 were discarded due to von Mises stress violations. Compared with the available dataset, the validated set achieves better front-quality indicators (lower IGD, higher HV, improved spacing), indicating that surrogate-driven exploration can effectively identify Pareto-relevant regions. The FEM-validated front also enables post-optimization insights (trade-offs, constraint proximity, recurring variable patterns) to support engineering decisions.

Future work will address: (i) fair benchmarking against FEM-driven NSGA-III under the same FEM budget, (ii) improved surrogate accuracy, especially near feasibility boundaries, (iii) online-learning infill to refine the surrogate in Pareto-relevant regions, and (iv) robustness/generalization across operating points and design domains.

References

[1] M. Gobbi *et al.*, “Traction motors for electric vehicles: Maximization of mechanical efficiency – a review,” *Applied Energy*, vol. 357, p. 122496, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261923018603

[2] C. De Gréef *et al.*, “Time-efficient multi-physics optimization approaches for the design of synchronous reluctance motors,” in *2021 IEEE Energy Conversion Congress and Exposition (ECCE)*, 2021, pp. 4197–4204.

[3] V. P, S. V. Nair, and K. Hatua, “Electromagnetic, structural and thermal analysis of interior permanent magnet synchronous motor for electric vehicle application,” in *2023 11th International Conference on Power Electronics and ECCE Asia (ICPE 2023 - ECCE Asia)*, 2023, pp. 148–153.

[4] R. Torchio *et al.*, “Design and experimental validation of a multiphysics twin of a high-voltage ev motor,” *IEEE Transactions on Transportation Electrification*, vol. 11, no. 1, pp. 3287–3297, 2025.

[5] 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.

[6] B. Praslicka, C. Ma, and N. Taran, “A computationally efficient high-fidelity multi-physics design optimization of traction motors for drive cycle loss minimization,” *IEEE Transactions on Industry Applications*, vol. 59, no. 2, pp. 1351–1360, 2023.

[7] C. A. Coello Coello *et al.*, “Evolutionary multiobjective optimization: open research areas and some challenges lying ahead,” *Complex & Intelligent Systems*, vol. 6, pp. 221–236, 2020. [Online]. Available: https://doi.org/10.1007/s40747-019-0113-4

[8] K. Deb and H. Jain, “An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: Solving problems with box constraints,” *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, 2014.

[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] L. Giaccone *et al.*, “Global and local sensitivity analysis applied to electric motor design by polynomial chaos expansion,” *IEEE Transactions on Magnetics*, pp. 1–1, 2025.

[11] P. Alotto *et al.*, “A multiquadrics-based algorithm for the acceleration of simulated annealing optimization procedures,” *IEEE Transactions on Magnetics*, vol. 32, no. 3, pp. 1198–1201, 1996.

[12] M. Cheng *et al.*, “A review of data-driven surrogate models for design optimization of electric motors,” *IEEE Transactions on Transportation Electrification*, vol. 10, no. 4, pp. 8413–8431, 2024.

[13] Y. Sakamoto *et al.*, “Multi-objective motor design optimization with physics-assisted neural network model,” in *2023 IEEE International Electric Machines & Drives Conference (IEMDC)*, 2023, pp. 1–7.

[14] S. Ferrari *et al.*, “Multiphysics Datasets for the Data-Driven Modeling of Traction Electric Motors,” 2025. [Online]. Available: https://zenodo.org/doi/10.5281/zenodo.17648857

[15] S. Ferrari and G. Pellegrino, “SyR-e Synchronous Reluctance – evolution design,” https://github.com/SyR-e/syre_public, 2025, [Online; accessed 03-Dec-2025].

[16] J. Blank *et al.*, “Generating well-spaced points on a unit simplex for evolutionary many-objective optimization,” *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 1, pp. 48–60, 2021.

[17] K. Deb and A. Srinivasan, “Innovization: innovating design principles through optimization,” in *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO ’06. New York, NY, USA: Association for Computing Machinery, 2006, pp. 1629—-1636.