Surrogate Modeling of Synchronous Reluctance Machines for Electric Vehicle Traction Using a Multilayer Perceptron

Andrea Credo
DIIIE
University of L'Aquila
L'Aquila, Italy
andrea.credo@univaq.it

David R. Stewart
SPARK Lab, Pigman Eng. College
University of Kentucky
Lexington, KY, USA
david.stewart@uky.edu

Marco Villani
DIIIE
University of L'Aquila
L'Aquila, Italy
marco.villani@univaq.it

Dan M. Ionel
SPARK Lab, Pigman Eng. College
University of Kentucky
Lexington, KY, USA
dan.ionel@ieee.org

***Abstract*—This work presents a surrogate modeling framework for the prediction of the d- and q-axis inductances of a synchronous reluctance machine (SynRM) based on a multi-output multilayer perceptron (MLP). A parametrically defined SynRM geometry, including rotor flux-barrier topology and stator slot variables, is optimized across multiple operating points to generate a dataset of over 3000 high-fidelity FEA evaluations. Two performance-relevant conditions are considered: maximum-torque point and the field-weakening operation. Hyperparameter tuning and an accelerated random forest based (ARF) algorithm are employed to serve as validation for the chosen model and architecture. Results show excellent agreement between predicted and FEA-calculated inductances; therefore, the proposed surrogate model allows for a rapid and accurate evaluation of SynRM designs, significantly reducing the computational burden associated with iterative finite-element-based optimization.**

I. Introduction

In the last few years synchronous reluctance motors (SynRMs) have attracted attention for electric vehicle (EV) traction applications, driven by the need to reduce reliance on rare-earth permanent magnets while ensuring good efficiency, cost effectiveness, and supply-chain robustness. Despite their lower torque and power density compared to the permanent-magnet machines, SynRMs offer a competitive compromise for automotive drivetrains, particularly in terms of efficiency especially at high-speed. Their high efficiency at elevated speeds is a key advantage for EV traction applications and primarily stems from the absence of a permanent-magnet flux linkage, [1]. This feature significantly reduces iron losses at high speed and eliminates rotor magnet losses, thereby simplifying rotor thermal management and improving overall high-speed robustness [2].

Recent studies have tackled these issues from complementary perspectives, ranging from analytical assessments of high-speed operating limits and flux-barrier design criteria to rotor-topology optimization [3]–[6]. For traction applications, sleeve-rotor and carbon-fiber-wrapped solutions have also been proposed to enlarge the admissible speed range while preserving rotor integrity [7], [8]. However, the extent of the constant-power field-weakening region is inherently limited by the achievable machine anisotropy, which cannot be excessively increased at high rotational speeds due to the presence of inner rotor ribs required to ensure mechanical integrity. This represents a limitation for high-speed EV traction operation [3], [4], [9]. Moreover, increased saturation and geometric constraints often lead to higher torque ripple, negatively impacting NVH performance [10], [11].

Another challenge for automotive SynRM adoption is related to the large number of flux barriers typically required to achieve high saliency ratios. While beneficial from an electromagnetic standpoint, this increases rotor-stamping complexity, production cost, and sensitivity to manufacturing tolerances in high-volume EV applications [6], [12]. The resulting design space can encompass ten or more independent rotor geometric variables per topology [13], making conventional FEM-based design workflows computationally demanding and poorly suited to automotive development timelines.

For these reasons, AI- and machine-learning- (ML) based surrogate models are increasingly being explored as an effective tool for electric motor design, enabling rapid performance prediction from a limited set of high-fidelity simulations while still capturing key nonlinear electromagnetic characteristics [14], [15]. Since the predictive capability of these models strongly depends on architectural choices and training settings, hyperparameter tuning plays a key role in improving accuracy, convergence behavior, and generalization performance, while reducing the risk of overfitting on the training data [16].

The aim of this paper is to propose a baseline optimization of the SynRM using conventional design and analysis techniques, followed by the introduction of an AI/ML-driven surrogate modeling framework to accelerate performance prediction and design-space exploration. This approach demonstrates how a reduced set of detailed simulations can be leveraged to enable faster machine evaluation and a more efficient overall optimization workflow. The paper is organized as follows: Section II presents the machine modeling approach and the geometric constraints adopted for the optimization and data preparation stages. Section III presents the optimization results together with a preliminary analysis of the generated dataset.

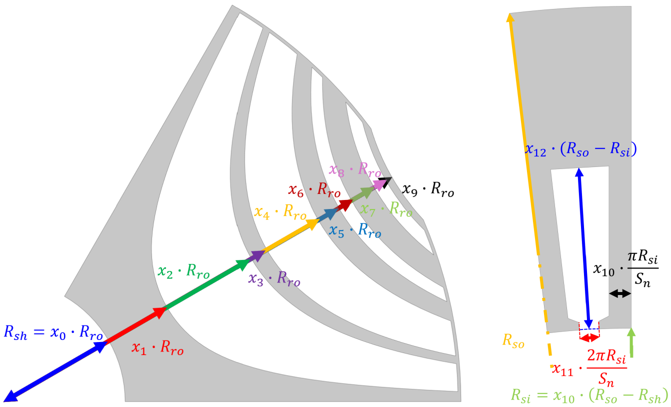

(a) (b)

Fig. 1. Geometric variables of the synchronous reluctance machine with (a) rotor Variables and (b) stator variables.

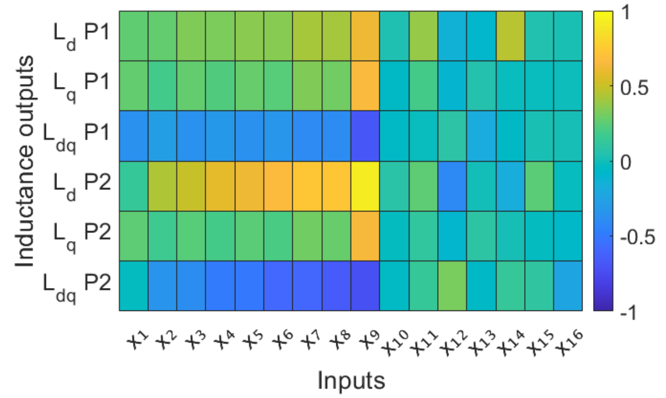

Fig. 2. Pearson correlation between geometric/control input variables and the main performance indicators.

Section IV describes in detail the adopted model architecture and training procedure, including the main modeling choices and the corresponding prediction results. Finally, Section V presents the conclusions.

II. Modeling Framework and Geometric Constraints

The SynRM geometry is fully parameterized using 13 per-unit geometrical design variables describing the radial coordinates of the four flux barriers, the notch position, the rotor radius, and the stator tooth/slot proportions, together with two control variables: the current magnitude and current angle.

The rotor flux barriers are defined using curves derived from the Joukowsky equations, providing smooth and physically consistent geometries. Once the radial coordinates of the barrier-defining lines are specified, the barrier shapes are uniquely determined. Enforcing each radial coordinate to be strictly greater than the previous one ensures a non-intersecting and solver-safe geometry without the need for iterative geometric corrections. The radial position as a function of the angular coordinate δ (measured from the d-axis) is:

$$A_i = \left(\frac{x_i}{x_0}\right)^p - \left(\frac{x_0}{x_i}\right)^p$$

$$r_i(\delta) = R_s\sqrt{A_i + \sqrt{A_i^2 + 4\sin^2(p\delta)}}.$$

where $r_i(\delta)$ is the radial position of the i-th Joukowsky line, x_i is the per-unit radial coordinate of that line, x_0 is the per-unit shaft radius, and A_i is an auxiliary geometric coefficient introduced to compact the barrier representation.

The design variables for a four-barrier-plus-notch, six-pole, 54-slot configuration are shown in Fig. 1a and Fig. 1b. All geometric parameters are expressed in per-unit (p.u.) form, ensuring generality and compatibility with a wide range of optimization tools. Simple linear constraints ensure manufacturability by enforcing minimum barrier separations and feasible slot geometry, enabling safe and solver-robust large-scale optimization without the need for geometric post-processing. The descriptions and ranges for the geometric design variables are expressed in the following:

- **x_1–x_2, x_3–x_4, x_5–x_6, x_7–x_8:** lower and upper radial coordinates of the four flux-barrier layers (relative to the rotor outer radius), with box constraints $0.322 \le x_i \le 0.985$ for $i = 1, \dots, 8$.
  Bounds selected to satisfy shaft radius + minimum carrier thickness, and to keep the final notch sufficiently wide.
- **x_9:** radial coordinate of the notch line with box constraint $0.322 \le x_9 \le 0.985$.
  (Same as above).
- **x_{10}:** rotor outer radius in p.u. relative to the design-space thickness (difference between stator outer radius and shaft radius), with constraint $0.775 \le x_{10} \le 0.900$.
  Practical rotor-sizing range within the available space.
- **x_{11}:** tooth width relative to slot pitch, constrained by $0.25 \le x_{11} \le 0.75$.
  Typical tooth-width ratios.
- **x_{12}:** slot opening relative to pitch, with $0.2 \le x_{12} \le 0.5$.
  Minimum ensures feasible winding insertion.
- **x_{13}:** slot height relative to stator radial thickness, with allowed range $0.60 \le x_{13} \le 0.85$.
  Typical slot depth-radial thickness ratios.

To ensure manufacturability and avoid geometric overlap, a set of additional linear constraints is imposed on the rotor variables. The minimum separation thresholds s_k represent the minimum admissible thickness between consecutive flux barriers. They are expressed in per-unit of the rotor outer diameter and decrease linearly from $s_1 = 0.032$ to $s_8 = 0.014$, reflecting the assumption that flux-barrier thickness decreases from the shaft region toward the air gap. The resulting geometric constraints can be expressed through the minimum barrier–separation condition and the slot–opening limitation, reported together below:

$$x_{k+1} \ge x_k + s_k \quad (k = 1, \dots, 8) \qquad\qquad x_{11} + x_{12} \le 1$$

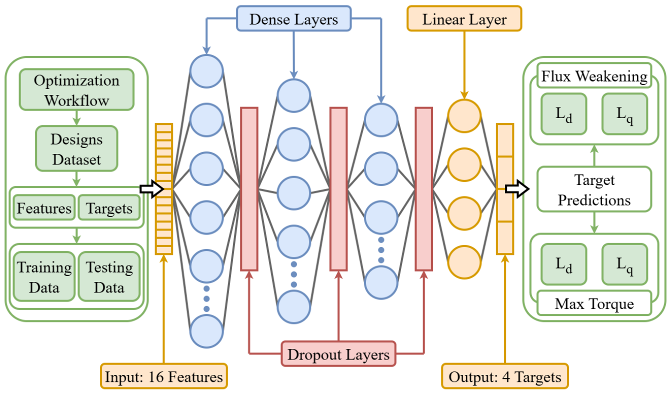

Fig. 3. Overview of the proposed surrogate model framework and MLP architecture, illustrating the workflow in which an optimized design dataset is used to train the MLP for rapid prediction of selected performance metrics.

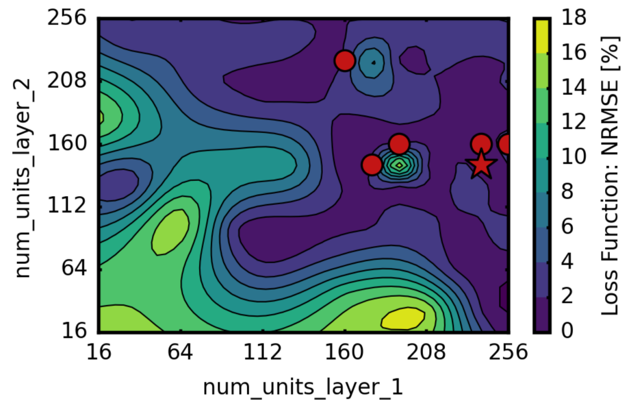

(a)

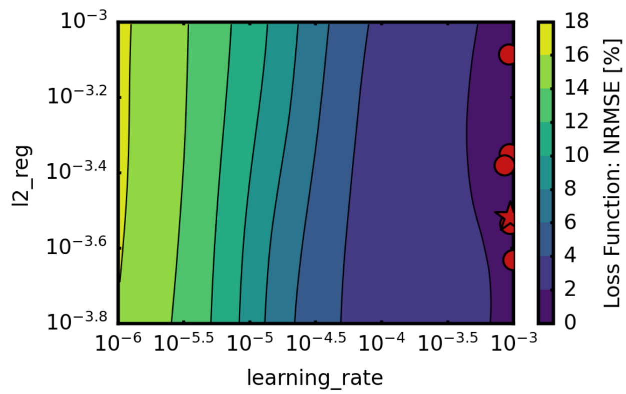

(b)

Fig. 4. Contour plots of explored hyperparameters for two cases: (a) the number of neurons in layer 1 versus layer 2 and (b) the learning rate versus L2 regularization. The marked points indicate lowest NRMSE percentage among all trials with the star representing the best case trial used in the final training step of the model.

These linear inequalities can be handled directly by most optimization algorithms and guarantee that all generated geometries remain manufacturable, mechanically robust, and electromagnetically valid.

To reduce computational burden while maintaining accurate performance prediction, the electromagnetic analysis evaluates one third of the electrical period. This setup provides reliable estimates of the quantities required for optimization, including peak voltage, average torque, torque ripple, flux linkages, and d-q axis inductances. Longer simulations for loss evaluation may be enabled when needed, but they are omitted from the main dataset to preserve computational efficiency.

In this work, the adopted FEA model serves as a validated benchmark for the study, since it was previously calibrated by applying the procedure proposed in [12] to the results presented in [17], which were obtained for a motor with dimensions comparable to those of the machine considered in this study. This calibration accounts for manufacturing effects, thereby allowing the numerical model to reproduce the actual electromagnetic behavior of the machine with improved accuracy. On this basis, the FEA results are used as the reference dataset for both the optimization process and the training and validation of the proposed surrogate model.

III. Optimization Results and Preliminary Data Analysis

A preliminary dataset of approximately 3000 evaluated designs was generated through the optimization workflows. The optimization was carried out using a population-based genetic algorithm. To obtain a sufficiently broad dataset for the subsequent machine-learning stage, the first 1500 samples were generated through random exploration of the admissible design space. This initial sampling strategy was adopted to avoid an excessive concentration of designs around high-performing solutions only, thus providing a more representative dataset for surrogate-model training. Two operating points were considered in the process: (i) the base-speed, maximum-torque condition at 6000 rpm, evaluated with a fixed current magnitude of 800 A and with the current-angle control ε included as an additional design variable (x_{14}); this operating point represents machine behavior under strong magnetic saturation. (ii) the field-weakening condition at 18 000 rpm, evaluated with both the current-angle control and the current magnitude treated as additional variables (x_{15} and x_{16}, respectively); this operating point corresponds to deep flux-weakening operation. This dual-point formulation enables the optimization to balance performance between nominal-speed and high-speed operation, reflecting typical design trade-offs of SynRM machines. Since these two operating points describe substantially different operating saturation behavior, a surrogate model that can represent both with good accuracy may be regarded as comparatively robust.

In addition to the geometric constraints required to avoid designs that could not be reliably simulated and to prevent manufacturing-critical geometries, nonlinear electromagnetic constraints were also adopted. These included a current-density constraint at the maximum-torque, maximum-current

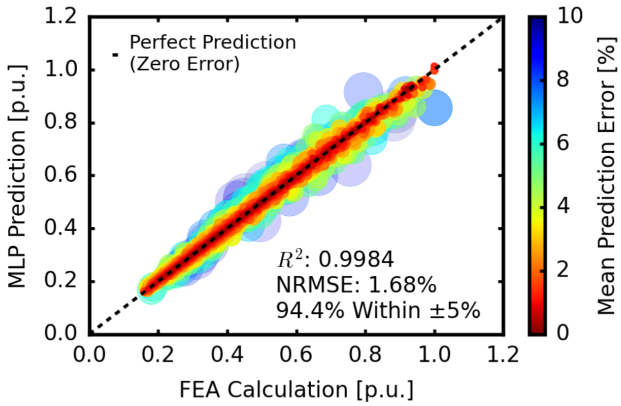

(a)

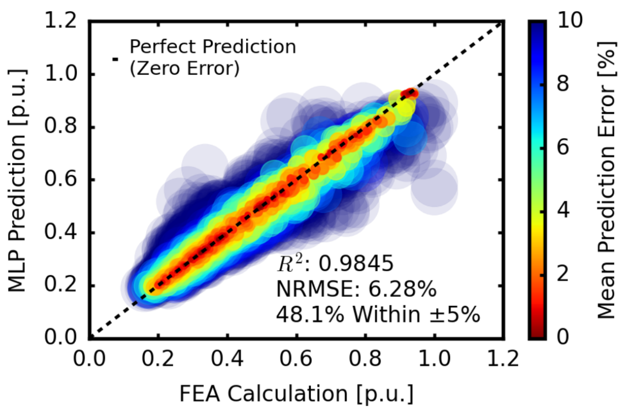

(b)

Fig. 5. Normalized per-unit regressions plots for (a) the MLP model and (b) the ARF model trained on the full dataset where the color and size of each point is proportional to the mean prediction error percentage across all targets. The MLP substantially outperforms the ARF model with nearly double the amount of designs within the ±5% error threshold.

operating point, introduced to satisfy the thermal limits imposed by the adopted cooling system, as well as line-to-line RMS voltage constraints at both operating points, required to ensure that the resulting solutions correspond to physically feasible operating conditions. For the voltage constraints, a 10% margin was imposed with respect to the theoretical maximum value obtainable from the DC bus, converted into the corresponding equivalent RMS voltage, in order to preserve a suitable control margin.

To further investigate the structure of the generated dataset, the linear relationships between the 16 geometric and control input variables and the electromagnetic performance indicators were analyzed through the Pearson correlation coefficient. As expected for synchronous reluctance machines, most of the correlations are weak to moderate, confirming the highly nonlinear nature of the geometry–performance relationship and further emphasizing the need for evolutionary or ML-based surrogate models.

Among the geometric inputs, the variables associated with the rotor—particularly the flux-barrier radial coordinates and the notch position—exert the strongest influence on torque, inductances, and flux linkages. The notch radial position (x_9) is especially impactful, as it determines the final radial extension of the outermost magnetic barrier. When placed near mid-radius (around 0.5 p.u.), the outer rotor region becomes predominantly air, increasing the effective q-axis airgap and reducing the available iron path. This leads to a significant reduction in saliency and flux-carrying capability. Extreme notch placements further distort the barrier distribution, yielding suboptimal electromagnetic performance.

The current angle control (x_{14}) at operating point 1 shows a clear positive correlation with the d-axis inductance. This behavior is physically consistent: as the current angle control increases, the magnetic saturation level of the machine is reduced, leading to an increase in the d-axis inductance.

TABLE I
SURROGATE PERFORMANCE METRICS ACROSS DIFFERENT TRAINING DATA SIZES

Num. Designs	Target	R^2	NRMSE	Within ±5% Err.
585	LdP1	0.97	0.0347	77.2
	LqP1	0.95	0.0408	80.2
	LdP2	0.98	0.0340	88.6
	LqP2	0.95	0.0366	77.7
1171	LdP1	0.98	0.0243	85.0
	LqP1	0.98	0.0276	90.4
	LdP2	0.98	0.0266	90.3
	LqP2	0.97	0.0256	87.5
1756	LdP1	0.99	0.0194	89.4
	LqP1	0.98	0.0187	92.8
	LdP2	0.99	0.0228	93.6
	LqP2	0.98	0.0197	91.9
2342	LdP1	0.99	0.0169	93.2
	LqP1	0.99	0.0165	94.5
	LdP2	0.99	0.0197	96.0
	LqP2	0.99	0.0140	93.7

IV. Machine Learning based Surrogate Model

A. Model Architecture and Training

Artificial neural networks (ANNs) have proven highly effective as meta and surrogate models for electric machines capable of capturing nonlinear relationships between design variables and performance metrics [18]–[20]. A multilayer perceptron (MLP), a class of feedforward ANN, was implemented as a SynRM surrogate model using TensorFlow Keras API [21]. A large-scale optimization dataset with two distinct operating points is used to train the model to simultaneously predict d- and q-axis inductances based on a set of 16 geometric and control variables. Prior to training, all input features and output targets were scaled using standard normalization methods to ensure consistent scaling and numerical stability.

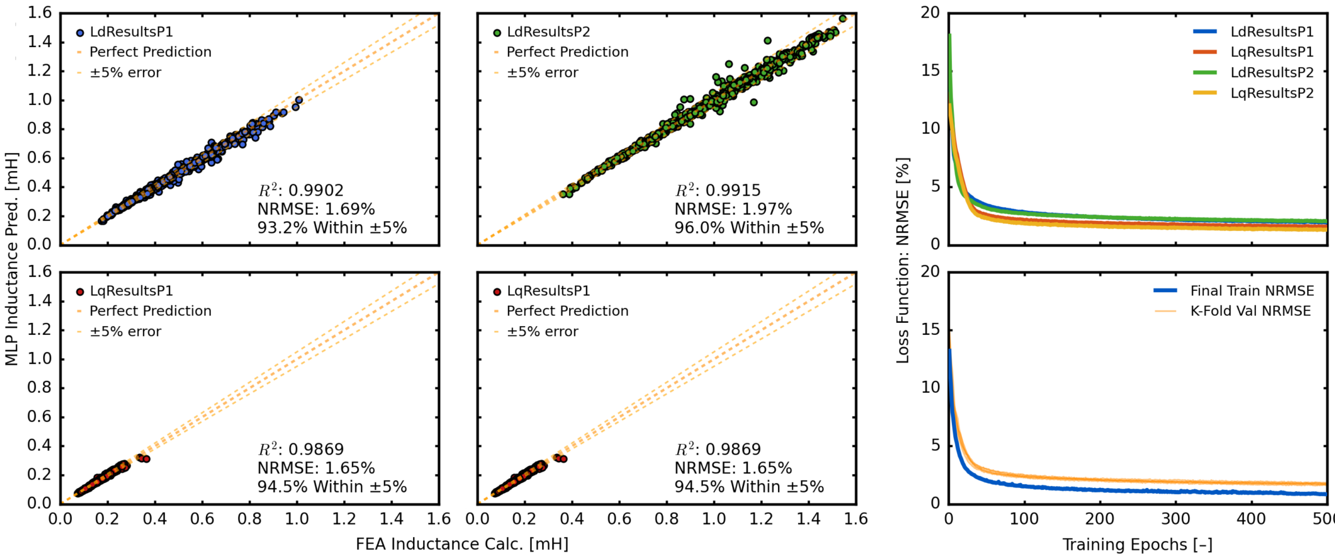

(a) (b)

Fig. 6. MLP surrogate model (a) per-target regression plots showing individual prediction performance for each inductance component and (b) individual target and overall NRMSE convergence across 500 training epochs trained on 100% of the available training data.

The MLPs architecture, depicted in Fig. 3, and hyperparameters are chosen via a tuning algorithm developed with Optuna [22] and employed to minimize the overall normalized root mean squared error (NRMSE). A search space is first defined for the chosen hyperparameters and then sampled using a tree-structured Parzen estimator (TPE) for the number of trials selected. A model is then built and trained using the sampled hyperparameters, as shown in the example contour plots in Fig. 4, with the lowest NRMSE trial being selected as the best case. This ensures that an optimal architecture is used for the final training step.

The hyperparameters chosen for this model consist of an input interface followed by three dense hidden layers with 240, 144, and 224 neurons, respectively. Each hidden layer utilizes the rectified linear unit (ReLU) activation function to introduce nonlinearity and enable the network to approximate complex relationships between the design variables and performance indicators. L2 weight regularization and dropout are applied to each fully connected layer to mitigate overfitting and improve generalization capability. The output layer is a linear dense layer where the number of assigned neurons is dependent on the number of predicted performance metrics. For this study 4 neurons are used in the output layer to represent the d- and q-axis inductances for both operating points.

Model training is performed using the Adam optimizer with early stopping enabled, based on validation loss, to ensure efficient convergence. A 5-fold cross-validation method is used to randomly partition the dataset into five equally sized subsets or “folds” to ensure that each data-point contributes to both training and validation. The resulting MLP acts as a computationally efficient surrogate model, enabling rapid exploration of the design space without requiring additional FEA simulations.

B. *Results and Discussion*

The performance of the proposed MLP surrogate model was assessed using both aggregate multi-output and individual per-target metrics for the four predicted inductance quantities across two points of operation for the SynRM. In effort to further justify the chosen methodology and machine learning model, an additional study using an accelerated random forest (ARF) algorithm is presented alongside the MLP results. The mean prediction error across all targets for the MLP model, depicted in Fig. 5(a), indicates that over 94% of all design point predictions fall within $\pm 5\%$ of their calculated FEA value, nearly doubling that of the alternative ARF model, shown in Fig. 5(b).

The MLP surrogate achieves a combined coefficient of determination (R^2) of 0.99 and a NRMSE of 0.0168 when trained with the entirety of the available training data. Additional studies, depicted in Table I, show that for the same test set, MLP models trained on as little as 25% of the available data demonstrate strong predictive capability even surpassing that of alternative methods trained on the entirety of the available data. This behavior suggests that the chosen surrogate methodology can generalize effectively across the design space without requiring extensive FEA simulations.

The per-target regression metrics, displayed in Fig. 6(a), further illustrate the model’s suitability as a surrogate capable of achieving near perfect alignment with FEA. The overall and individual target loss function convergence behavior, depicted in Fig. 6, shows a smooth and stable decrease of NRMSE with stabilization early in the training process, indicative of an efficient model architecture. These results indicate the network is successful in capturing the nonlinear dependence of the d- and q-axis inductances on the SynRM geometric and control variables at both operating points.

V. Conclusion

This study presents a surrogate modeling framework for EV traction synchronous reluctance machines based on a multi-output MLP capable of predicting the d- and q-axis inductances at two points of operation representative of the machine's full operating envelope: maximum torque and deep field-weakening. The parameterized SynRM model and the large-scale optimization workflow enable the creation of a rich dataset that captures the nonlinear dependence of key electromagnetic performance metrics on the geometric and control variables.

The surrogate model demonstrates high predictive accuracy across all outputs, with an R^2 exceeding 0.99 and over 94% of the predictions falling within $\pm 5\%$ of FEA results. Comparison against an ARF baseline confirms that the proposed MLP architecture generalizes well across the design space even for reduced training set sizes. These findings indicate that a properly trained neural-network surrogate can effectively replace repeated finite-element simulations within optimization loops, offering enhanced computational efficiency necessary for supporting automotive development timelines.

Acknowledgment

At University of Kentucky the work was supported by the National Science Foundation (NSF) Graduate Research Fellowship under Grant No. 2239063. The support of ANSYS Inc. and University of Kentucky L. Stanley Pigman Chair in Power endowment is also gratefully acknowledged. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsoring organizations.

References

[1] M. Murataliyev, M. Degano, M. Di Nardo, N. Bianchi, and C. Gerada, "Synchronous reluctance machines: A comprehensive review and technology comparison," *Proceedings of the IEEE*, vol. 110, no. 3, pp. 382–399, 2022.

[2] K. B. Tawfiq, M. Güleç, P. Sergeant, H. Zeineldin, M. N. Ibrahim, A. Al-Durra, and E. F. El-Saadany, "A comprehensive study of thermal and performance characteristics in revamped multiphase synchronous reluctance machines," *IEEE Transactions on Transportation Electrification*, vol. 11, no. 1, pp. 4717–4729, 2025.

[3] G. Gallicchio, M. D. Nardo, M. Palmieri, A. Marfoli, M. Degano, C. Gerada, and F. Cupertino, "High speed synchronous reluctance machines: Modeling, design and limits," *IEEE Transactions on Energy Conversion*, vol. 37, no. 1, pp. 585–597, 2022.

[4] C. Babetto, G. Bacco, and N. Bianchi, "Synchronous reluctance machine optimization for high-speed applications," *IEEE Transactions on Energy Conversion*, vol. 33, no. 3, pp. 1266–1273, 2018.

[5] A. Belahcen, M. A. Sitnikov, and S. A. Galunin, "Optimization of a high-speed synchronous reluctance machine's rotor topology," in *2021 International Conference on Electrotechnical Complexes and Systems (ICOECS)*, 2021, pp. 279–284.

[6] G. Bacco and N. Bianchi, "Design criteria of flux-barriers in synchronous reluctance machines," *IEEE Transactions on Industry Applications*, vol. 55, no. 3, pp. 2490–2498, 2019.

[7] P. B. Reddy, K. Grace, and A. M. El-Refaie, "Conceptual design of sleeve rotor synchronous reluctance motor for traction applications," *IET Electric Power Applications*, vol. 10, no. 5, pp. 368–374, 2016.

[8] K. Grace, S. Galioto, K. Bodla, and A. M. El-Refaie, "Design and testing of a carbon-fiber-wrapped synchronous reluctance traction motor," *IEEE Transactions on Industry Applications*, vol. 54, no. 5, pp. 4207–4217, 2018.

[9] A. Credo, M. Villani, G. Fabri, and M. Popescu, "Adoption of the synchronous reluctance motor in electric vehicles: A focus on the flux weakening capability," *IEEE Transactions on Transportation Electrification*, vol. 9, no. 1, pp. 805–818, 2023.

[10] Y. Xu, Z. Xu, H. Cao, and W. Liu, "Torque ripple suppression of synchronous reluctance motors for electric vehicles based on rotor improvement design," *IEEE Transactions on Transportation Electrification*, vol. 9, no. 3, pp. 4328–4338, 2023.

[11] V. Varvolik, G. Buticchi, D. Prystupa, S. Wang, A. Aboelhassan, S. Peresada, and M. Galea, "Comparative study on torque ripple reduction considering minimum losses for synchronous reluctance motor drives," *IEEE Transactions on Transportation Electrification*, vol. 10, no. 3, pp. 6527–6538, 2024.

[12] A. Credo and P. Pescetto, "A systematic method for the calibration of fea model of synchronous reluctance machines considering manufacturing effects," *IEEE Open Journal of Industry Applications*, vol. 6, pp. 539–550, 2025.

[13] Y. Wang, D. M. Ionel, M. Jiang, and S. J. Stretz, "Establishing the relative merits of synchronous reluctance and pm-assisted technology through systematic design optimization," *IEEE Transactions on Industry Applications*, vol. 52, no. 4, pp. 2971–2978, 2016.

[14] M. Cheng, X. Zhao, M. Dhimish, W. Qiu, and S. Niu, "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.

[15] Y. Gong, A. Gneiting, S. Weigel, N. Parspour, and Z. An, "Surrogate model-based drive cycle modeling and optimization of synchronous reluctance machines for electric vehicles," *IEEE Transactions on Magnetics*, vol. 61, no. 9, pp. 1–6, 2025.

[16] S. Shekhar, A. Bansode, and A. Salim, "A comparative study of hyper-parameter optimization tools," in *2021 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, 2021, pp. 1–6.

[17] A. Credo, I. Petrov, J. Pyrhönen, and M. Villani, "Impact of manufacturing stresses on multiple-rib synchronous reluctance motor performance," *IEEE Transactions on Industry Applications*, vol. 59, no. 2, pp. 1253–1262, 2023.

[18] O. A. Badewa and D. M. Ionel, "Hybrid fea and meta-modeling for de optimization of a pm stator-excited motor with a reluctance rotor," in *2025 IEEE/AIAA Transportation Electrification Conference and Electric Aircraft Technologies Symposium (ITEC+EATS)*, 2025, pp. 1–6.

[19] M. Tahkola, J. Keränen, D. Sedov, M. F. Far, and J. Kortelainen, "Surrogate modeling of electrical machine torque using artificial neural networks," *IEEE Access*, vol. 8, pp. 220 027–220 045, 2020.

[20] M. Vatani, D. R. Stewart, D. D. Lewis, and D. M. Ionel, "Design optimization and scaling of coreless afpm machines using hybrid fea-based differential evolution and machine learning," in *2025 IEEE Energy Conversion Conference Congress and Exposition (ECCE)*, 2025, pp. 1–7.

[21] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/

[22] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.