



Proceedings of the
Estonian Academy of Sciences
2026, **75**, 4, 325–335

<https://doi.org/10.3176/proc.2026.4.06>

www.eap.ee/proceedings
Estonian Academy Publishers

ELECTRICAL ENGINEERING

RESEARCH ARTICLE

Received 1 July 2026
Accepted 16 September 2026
Available online 21 October 2026

Keywords:

power electronics, inductor design,
magnetic inductor, reduced-order
modeling, sparse measurements, singular
value decomposition, QR decomposition,
nonlinear magnetic materials

Corresponding author:

Kamil Gruszczyk
kamil.gruszczyk@polsl.pl

Citation:

Gruszczyk, K., Lasek, P., Stępień, M. and
Habelok, K. 2026. Sparse-sensing-based
reduced modeling for rapid
characterization of nonlinear magnetic
inductors. *Proceedings of the Estonian
Academy of Sciences*, **75**(4), 325–335.
<https://doi.org/10.3176/proc.2026.4.06>

Sparse-sensing-based reduced modeling for rapid characterization of nonlinear magnetic inductors

Kamil Gruszczyk , Paweł Lasek , Mariusz Stępień  and
Krzysztof Habelok 

Electrical Drives and Robotics, Department of Power Electronics, Silesian University of
Technology, Krzywoustego 2, 44-100 Gliwice, Poland

ABSTRACT

Rapid prototyping of magnetic inductors requires the evaluation of inductance over a wide range of turn counts and operating currents. This task becomes computationally or experimentally demanding when nonlinear magnetic materials are considered because each candidate design may require iterative solution of the magnetic circuit model or repeated evaluations. This paper proposes a sparse-sensing framework for reconstructing inductance curves using a small number of strategically selected turn counts. A reduced basis is constructed by singular value decomposition (SVD), and informative sensing turn counts are selected by QR decomposition with column pivoting. The resulting model reconstructs the full inductance–turn relationship from only a few evaluations. The method is demonstrated using simulated nonlinear toroidal inductor data and then incorporated into a multi-objective design workflow to select candidate cores from a commercial database. The results show that five selected turn counts are sufficient to reconstruct the inductance curve with a relative error below 0.1% in the validation cases under consideration. Combined with the non-dominated sorting genetic algorithm II (NSGA-II) optimization, the reduced model enables rapid screening of feasible core, material, and stacking configurations while preserving accuracy sufficient for engineering design.

1. Introduction

Magnetic inductors are among the fundamental components of power-electronic converters, filters, and energy-conversion systems. In many practical designs, the magnetic core is intentionally or unavoidably operated outside the strictly linear region [1,2]. This makes the inductance dependent on the operating current, core geometry, and material properties. Although this nonlinearity complicates modeling and design, it can also be exploited to advantage. Partially saturated inductors have recently been used in boost converters, together with nonlinear model predictive control, in which magnetic saturation and losses are included in the inductor model [1]. Nonlinear inductors have also been investigated in inverter systems, where the current-dependent inductance affects dynamic behavior, stability boundaries, and medium-frequency oscillations [2].

The accurate characterization of nonlinear magnetic inductors remains challenging. Classical toroidal-core models already show that inductors made from nominally linear ferrites may exhibit measurable nonlinear behavior [3]. More detailed approaches have also been proposed. Finite-element-based procedures compute the inductance of toroidal ferrite cores from the linear region up to saturation [4]. Equivalent circuit models reproduce voltage and current waveforms while accounting for saturation and core losses [5]. Mixed magnetic–electronic time-domain simulation techniques address nonlinear and hysteretic magnetic systems coupled to drive electronics [6]. For laminated magnetic cores supplied by pulse-width-modulated converters, additional phenomena such as hysteresis, eddy-current losses, excess losses, and skin effect may also need to be represented in compact simulation models [7]. These studies show that nonlinear magnetic components require models that are accurate enough for design yet sufficiently inexpensive to evaluate repeatedly.

A typical inductor design task requires evaluating the inductance over a range of turn counts, operating currents, magnetic materials, and core geometries. Direct numerical or experimental characterization of all combinations may therefore become expensive, particularly when nonlinear material properties require iterative solution procedures. Reduced-order modeling provides a natural way to exploit the structure of such parametric data. Proper orthogonal decomposition (POD), introduced by Sirovich [8], represents complex datasets using a small number of dominant modes. Related ideas appear in compressed sensing, where sparse signals can be reconstructed from limited measurements [9,10]. In reduced-order modeling, sparse-sampling strategies such as the discrete empirical interpolation method (DEIM) and QR-based selection operators make it possible to reconstruct parameter-dependent quantities from a small number of informative sensing turn counts [11].

This paper applies sparse-sensing-based reduced modeling to the rapid characterization of nonlinear magnetic inductors. Instead of determining the full inductance–turn characteristic for every candidate design, a reduced basis is constructed using singular value decomposition (SVD). The number of retained modes is selected from the singular value spectrum, with reference to optimal hard thresholding concepts [13]. The sensing turn counts are then selected using QR decomposition with column pivoting. The resulting framework reconstructs the full inductance curve from a small number of inductance values and is subsequently combined with multi-objective optimization. The final design stage uses the non-dominated sorting genetic algorithm II (NSGA-II) [14], a fast elitist multi-objective method with crowding-distance preservation to identify Pareto-optimal inductor configurations with respect to winding length, core losses, and core volume.

The purpose of the present study is to demonstrate the benefits of combining SVD with complex multi-objective algorithms such as NSGA-II for the fast and efficient selection of electrical components and devices. The method introduces a new approach to the selection of nonlinear Sendust toroidal cores and, at the same time, provides a software environment that can be applied to other design problems involving other machines and devices. The designer saves time and selects products of a given manufacturer more easily among the many items available on the market, using their own criteria, including economic ones.

2. Magnetic circuit model

Digital toroidal cores were generated and used to construct the snapshot data for reduced-order modeling. A toroidal magnetic circuit is considered, for which the nonlinear magnetic reluctance is defined as

$$\mathcal{R}(B) = \frac{l_c}{A_c \mu_0 \mu_r(B)}, \quad (1)$$

where $\mathcal{R}$ is the magnetic reluctance, l_c is the magnetic path length, A_c is the core cross-sectional area, μ_0 is the vacuum permeability, B is the magnetic flux density in the core, and $\mu_r(B)$ is the relative nonlinear permeability.

The core dimensions and nominal material data used in this study are based on the Micrometals MS Sendust toroidal-core catalog [15]. Sendust is an iron–silicon–aluminum alloy powder material available in several nominal relative-permeability classes. The dimensions and the nominal permeability classes are catalog data, whereas the relationship between μ_r and B is a model. This catalog was selected because of its availability and large set of data on nonlinear toroidal cores. The Sendust data file contains nine nominal permeability classes, of which six were selected for the present model: 60, 75, 90, 125, 147, and 160. These catalog values are treated as permeability classes rather than constant permeabilities because the effective permeability varies with magnetic loading due to saturation.

The relative permeability of the Sendust material is represented using a smooth saturation function,


$$\mu_r(B) = \frac{\mu_{r0}}{1 + \exp(a(B - B_s))}, \quad (2)$$

where μ_{r0} is the initial relative permeability, a is the saturation-sharpness coefficient, and B_s is the saturation flux density.

A detailed list of the magnetic permeabilities used in the calculations is given in Table 1. The permeability values were computed according to Eq. (2).

Table 1. Parameters of the saturation function described by Eq. (2)

μ_{r0} (–)	a (1/T)	B_s (T)
60	18	0.28
75	20	0.27
90	25	0.25
125	26	0.25
147	18	0.22
160	20	0.20

**Fig. 1.** Inductance as a function of the number of turns for two example current levels. Calculated using Eqs (1)–(5) for the MS-184 core with a material class of $\mu_{r0} = 60$.

The saturation function in Eq. (2) is an approximation, not a measured curve or one derived from the manufacturer's data. The values of a and B_s in Table 1 were assumed for the purposes of the experiment. The inductance values given refer to the material model that was adopted, not to a core from the catalog. The L_2 error is calculated relative to the reference solution from the same model. It characterizes the accuracy of the reconstruction, not its agreement with a physical core.

The magnetic circuit equation is written as

$$NI = \mathcal{R}(B)\Phi, \quad (3)$$

where N is the number of turns, I is the current, and Φ is the magnetic flux. The equation is solved numerically for the flux density, and the inductance is computed as

$$L = \frac{N\Phi}{I}. \quad (4)$$

For $I = 0$, the ratio in Eq. (4) is indeterminate. In this case, the inductance is evaluated from its small-current limit,

$$L_{I \rightarrow 0} = \frac{N^2 \mu_0 \mu_{r0} A_c}{l_c}. \quad (5)$$

For $I > 0$, the flux density is obtained by solving Eq. (3) numerically, and the inductance follows from Eq. (4).

Figure 1 illustrates the dependence of inductance on the number of turns for two operating currents.

The curves differ because the effective permeability changes with the operating point. Consequently, the inductance cannot be described by a single linear expression over the full current range, and iterative calculations are required for each combination of core geometry, material, winding

turns, and current. The plot is based on synthetic data and shows the differences in the growth of the individual curves as the current changes, which justifies the use of SVD in order to avoid iterative procedures that are very time-consuming.

3. Reduced-basis construction

To analyze the structure of the inductance curves, a dataset of magnetic configurations was generated using randomly selected combinations of cores, materials, and operating currents. The snapshot set consists of 100 configurations. Each configuration comprises one of nine core geometries, one of the permeability classes listed in Table 1, and an operating current value within the range of 0 to 100 A. Each of these configurations was selected at random. A fixed random seed of 0 was used, and the resulting training currents lie between 3.203 A and 97.898 A. The resulting inductance values span the range from approximately 0.12 μH to 22 115 μH . The study is focused on a turn-count range of $N = 200$. Each column of the snapshot matrix $\mathbf{X}$ corresponds to a complete inductance curve sampled over the considered turn-count range:

$$\mathbf{X} = \begin{bmatrix} L_{1,1} & L_{1,2} & \cdots & L_{1,m} \\ L_{2,1} & L_{2,2} & \cdots & L_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ L_{n,1} & L_{n,2} & \cdots & L_{n,m} \end{bmatrix} \in \mathbb{R}^{n \times m}, \quad (6)$$

where n is the number of sampled turn counts, m is the number of magnetic configurations, and $L_{i,j}$ denotes the inductance at the i -th turn count for the j -th configuration. The snapshot matrix can be extended at every iteration. The number m is the number of columns and therefore the number of distinct cases and core configurations, which differ in dimensions, materials, magnetic permeability, and other parameters. This allows the algorithm to learn and to recognize patterns. In the present case, based on synthetic data, m reflects the number of combinations contained in the snapshot matrix. The knowledge base can be extended at any time, which will improve the accuracy obtained when a new core is analyzed. The value of m follows from the combinations produced by the script; it was not imposed in the experiment but was determined by other factors such as the choice of the number of cores, materials, and other parameters taken from the synthetic database. The number n is the number of turns, which was set to $N = 200$. It should be added that, for the algorithm to compute each combination correctly, every sampling point up to the 200th turn must be evaluated.

SVD is applied to the snapshot matrix,

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (7)$$

where $\mathbf{U}$ contains the left singular vectors, $\mathbf{\Sigma}$ contains the singular values, and $\mathbf{V}$ contains the right singular vectors. In this formulation, the columns of $\mathbf{U}$ are orthonormal modes defined over the turn-count coordinate N . Each mode describes a characteristic shape of the inductance curve, while the corresponding singular value quantifies its contribution to the variability of the snapshot dataset. The right singular vectors in $\mathbf{V}$ contain the coefficients associated with individual magnetic configurations.

This decomposition is equivalent to POD. Rapid decay of the singular values indicates that most inductance curves in the dataset can be represented accurately in a low-dimensional subspace. The reduced basis $\mathbf{\Psi}$ is therefore defined as

$$\mathbf{\Psi} = \mathbf{U}_r = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_r], \quad (8)$$

where $\mathbf{U}_r$ denotes the matrix formed by retaining the first r columns of $\mathbf{U}$. Thus, $\mathbf{\Psi} \in \mathbb{R}^{n \times r}$ spans the reduced space used to approximate the full inductance vector over all considered turn counts. Here, n is the number of sampled turn-count values, and $r \ll n$ is the reduced dimension. The matrix $\mathbf{\Psi}$ is the reduced basis obtained by truncating the redundant columns of $\mathbf{U}$. This operation is required in order to reduce the cost of the algorithm and to relieve the computer of unnecessary additional calculations.




Fig. 2. Selection of the reduced basis dimension using singular value decomposition (SVD): (a) singular value spectrum, (b) cumulative energy of the SVD modes, and (c) relative L_2 error as a function of the number of retained modes. The dashed red line indicates the selected truncation level $r = 5$.

The reduced dimension r is selected from the singular-value spectrum. Since the singular values are sorted in descending order, the leading columns of $\mathbf{U}$ contain the most informative modes.

The cumulative energy shown in Fig. 2b is defined as

$$E_r = \frac{\sum_{i=1}^r \sigma_i^2}{\sum_{i=1}^k \sigma_i^2} \cdot 100\%, \quad (9)$$

where σ_i are the singular values contained in $\mathbf{\Sigma}$, and k is their total number.

The optimal singular-value threshold proposed by Gavish and Donoho [13] provides a principled rule for choosing the truncation level.

Figure 2c shows the relative L_2 reconstruction error obtained for each number of retained modes on the validation set of 40 configurations. The blue curve represents the mean error and the red curve

the worst case. Each point is a separate experiment carried out in order to obtain the error statistics. A steep decrease is visible at the beginning, where the mean error falls from 3.81% at $r = 1$ to 0.065% at $r = 4$, after which the curve flattens out. This indicates that increasing the number of modes further is not worthwhile for the present experiment because the gain in accuracy is small compared with the additional computational cost. Despite the flattening already visible at $r = 4$, it is only at $r = 5$, with an L_2 error of 0.088%, that the error falls below 0.1%, which justifies the use of $r = 5$ modes.

4. Assumptions for sparse-sensing reconstruction and results

4.1. Reconstruction procedure

Instead of evaluating inductance at all turn counts, the full curve can be reconstructed from evaluations at a small number of selected turns. Let $\mathbf{C}$ be a sampling matrix that selects the evaluated entries from the full inductance vector. If $\mathbf{y}$ denotes the vector of evaluated inductances, the reconstructed full inductance curve is obtained as

$$\mathbf{L} = \Psi(\mathbf{C}\Psi)^{-1}\mathbf{y}, \quad (10)$$

where $\mathbf{L}$ is the reconstructed inductance vector, Ψ is the reduced basis, and $\mathbf{C}\Psi$ is the sensing matrix.

The sensing turn counts are selected using QR decomposition with column pivoting applied to the reduced basis matrix [11]. The reduced basis Ψ is obtained from the matrix $\mathbf{U}$. The pivoted QR decomposition is applied to the transpose Ψ^T , of size $r \times n$, which returns a permutation vector $\mathbf{p}$; its first r entries, sorted in ascending order, are taken as the sensing indices. The pivoting order indicates in which step the algorithm selected each point. In this way, the sensing turn counts are obtained, at which the inductance has to be evaluated in order to reconstruct the whole curve. For the considered dataset, the algorithm selects the following sensing turn counts:

$$\mathcal{N}_s = \{10, 28, 60, 122, 199\}. \quad (11)$$

Only the inductance values at these five turns are required to reconstruct the complete inductance curve over the considered range. The values given above refer to the physical numbering of the turns and not to the positions in the data table of the script; in the code, the sequence therefore runs from 0 to 199 rather than from 1 to 200. The computations were carried out in Python using standard libraries and their built-in functions: NumPy 2.5.2, SciPy 1.18.1, and Matplotlib 3.11.1. The SVD was computed with `scipy.linalg.svd`, the pivoted QR decomposition with `scipy.linalg.qr`, and the reconstruction system of Eq. (10) was solved with `numpy.linalg.solve`. The condition number of the sensing matrix, evaluated with `numpy.linalg.cond`, is 2.4458.

4.2. Validation results

The accuracy of the reconstruction was determined using the L_2 error over the entire curve and the relative point error, taking its sign into account:

$$\varepsilon_{L_2} = \frac{\|\mathbf{L}_{\text{rec}} - \mathbf{L}_{\text{ref}}\|_2}{\|\mathbf{L}_{\text{ref}}\|_2} \cdot 100\%, \quad \varepsilon_i = \frac{L_{\text{rec},i} - L_{\text{ref},i}}{L_{\text{ref},i}} \cdot 100\%, \quad (12)$$

where $\mathbf{L}_{\text{ref}}$ is the reference inductance vector obtained by solving Eqs (1)–(4) at every turn count, $\mathbf{L}_{\text{rec}}$ is the vector reconstructed from Eq. (10), and i runs over the sampled turn counts. Both measures are normalized by the reference value, so ε_{L_2} is a single number per configuration, and ε_i is a curve.

The method was validated on a non-overlapping set of 40 magnetic configurations, which were not used to construct the reduced basis or to select the sensing turn counts. Inductance values were evaluated only at the selected sensing turn counts, and the full inductance curve was reconstructed using the reduced-order model. Across the entire validation set, the relative L_2 error from Eq. (12) has a mean of 0.0388%, a median of 0.0354%, and a worst-case value of 0.0880%. An example reconstruction is shown in Fig. 3 for a configuration whose L_2 error corresponds to the median of this distribution. The reconstructed inductance curve corresponds to the reference solution, with an L_2 error of 0.0358% for the case shown. The largest pointwise relative error occurs at very low inductance values, where relative error measures are more sensitive to small absolute deviations: the maximum of 39.021% is reached at $N = 1$, where $L_{\text{ref}} = 0.2015 \mu\text{H}$, and the corresponding absolute deviation is $0.0786 \mu\text{H}$.




Fig. 3. Validation of the sparse-sensing model for a test configuration with $r = 5$: (a) the reconstructed inductance curve $L(N)$ plotted against the reference curve and (b) pointwise relative error of the reconstruction.

Table 2. Comparison of the accuracy and runtime of the proposed method and the reference methods on a dataset comprising 40 configurations for $r = 5$

Method	Sensing turn counts (-)	Mean L_2 (%)	Runtime (μ s)
SVD with pivoted QR (proposed)	10, 28, 60, 122, 199	0.0388	8.1
SVD with classical DEIM	11, 24, 71, 127, 199	0.0413	8.3
Cubic spline, same points	10, 28, 60, 122, 199	0.1576	166.3
Polynomial of degree 4, same points	10, 28, 60, 122, 199	0.5367	51.4
Cubic spline, equidistant	1, 51, 100, 149, 199	0.4380	159.8
Polynomial of degree 4, equidistant	1, 51, 100, 149, 199	0.3985	51.5
Quadratic model $L \sim N^2$	10, 28, 60, 122, 199	13.7890	43.6

4.3. Comparison with reference methods

For reference purposes, computations were also carried out using cubic spline interpolation, polynomial approximation of degree four, and the classical DEIM [12]. Table 2 additionally includes variants of the spline and polynomial methods that use the same number of sensing turn counts at equal intervals, as well as the quadratic model $L \sim N^2$.

The L_2 error is described by Eq. (12) and defines the relative error of the curve reconstruction with respect to the reference solution calculated for the full range. As shown in Table 2, the largest L_2 error occurred in the quadratic method and amounted to 13.7890%. Only in the case of the SVD-based methods did the error fall below 0.1%. Nevertheless, SVD with pivoted QR proved to be the most accurate. Furthermore, in terms of running time, the method described in this article proved to be the fastest of all the reference methods.

The data presented in this paper are synthetic. They were generated numerically and no physical measurements were performed at any stage. The results are presented for data describing a simulated toroidal inductor with a nonlinear characteristic. The data shown in Figs 1–5 were generated synthetically and do not correspond to real measurements.

5. Inductor design as a multi-objective optimization

Designing a magnetic inductor that satisfies requirements at both low and high operating currents is a multi-criteria optimization problem. In this study, the target low-current inductance is approximately 550 μ H at $I = 0$ A, while the high-current inductance should not fall below approximately 200 μ H at $I = 20$ A. Here, $L(I = 0)$ is understood in the sense of the small-current limit of Eq. (5), introduced in Section 2. A tolerance of ± 50 μ H is allowed for both target values. The resulting bound of $200 - 50 = 150$ μ H follows from the adopted tolerance conditions and is acceptable for the purposes of this experiment.

At the same time, three competing objectives are minimized: winding length, core losses, and core volume. Winding length is used as a geometry-based measure of the winding requirement, rather than winding resistance. Unlike resistance, it does not depend on the selected conductor cross section and is therefore not directly affected by application-specific assumptions concerning root-mean-square (RMS) current, permissible current density, or cooling conditions. This is appropriate for the present formulation, where the inductor candidates are primarily constrained by peak-current inductance requirements, while the final wire size may be selected later based on thermal and RMS-current criteria.

Core losses are included as an approximate indicator of magnetic loss performance. The loss values are estimated using the available Sendust core-loss model, which is sufficient for comparing candidate configurations within the same design workflow. However, because the material response is strongly nonlinear, the estimated losses should not be interpreted as exact values. In practice, they may depend on the excitation waveform, frequency, local flux density distribution, and thermal operating point.

Core volume is minimized to promote compact, lightweight, and cost-effective designs. This objective is necessary because the inductance constraints alone could often be satisfied by selecting an excessively large core. Including the core volume therefore penalizes oversized solutions and makes the resulting Pareto front more relevant from a practical engineering perspective.

The search space is based on the Micrometals MS Sendust toroidal-core catalog, an excerpt of which is shown in Table 3. The catalog contains 545 entries describing toroidal cores with different dimensions, magnetic path lengths, effective cross-sectional areas, volumes, and inductance factors [15]. Of these 545 records, 356 correspond to one of the six permeability classes listed in Table 1. Combined with the three stacking variants, corresponding to one, two, or three cores, this gives $356 \times 3 = 1068$ combinations that were checked.

In the first stage, candidates were pre-filtered to ensure that they satisfy the inductance requirements. The initial filtering retained 185 combinations, which correspond to 106 entries in the catalog, taking into account that a single core can be stacked. Verification of the 638 variations in the number of turns contained within the 185 combinations rejected 18 of them and retained 167 feasible designs. After taking into account constraints relating to volume, length, and core losses, 91 cases searched by the NSGA-II algorithm remained in the candidate space. Figure 4 shows the distribution of outer diameters and core volumes for the 167 feasible designs. The widespread occurrence of both quantities indicates that many geometrically distinct solutions can satisfy the same electrical constraints, motivating an additional optimization step.

A key advantage of the sparse-sensing approach is that the reduced basis Ψ and the sampling matrix $\mathbf{C}$ are determined once from the snapshot matrix and then reused for the core–material pairs in the catalog. This eliminates the need to perform full numerical calculations for each candidate individually.

After pre-filtering, a three-dimensional objective space is constructed for all feasible candidates using core volume, winding length, and core losses. This space is shown in Fig. 5. The objectives are mutually conflicting: configurations with a smaller core volume may require longer windings or exhibit higher losses. Therefore, there is no unique optimum, and the design problem is naturally formulated in terms of Pareto optimality.

Table 3. Excerpt from the Micrometals MS series core catalog [15]

Part No.	Core name	μ_r	A_L (nH/turns ²)	OD (mm)	ID (mm)	H (mm)	A_e (cm ²)	l_e (cm)	V_e (cm ³)
MS-014014-8	MS-014	14	3	3.56	1.78	1.52	0.0137	0.817	0.0107
MS-014026-8	MS-014	26	5.5	3.56	1.78	1.52	0.0137	0.817	0.0107
MS-014040-8	MS-014	40	9	3.56	1.78	1.52	0.0137	0.817	0.0107
MS-014060-8	MS-014	60	13	3.56	1.78	1.52	0.0137	0.817	0.0107
...	...	...	...	...	...	...	...	...	...
MS-775060-2	MS-775	60	88	196.85	146.05	25.4	6.26	53.9	337




Fig. 4. Distribution of geometrical properties for the 167 feasible designs retained after pre-filtering: (a) core outer diameter and (b) magnetic core volume V_c . The candidates are sorted independently in ascending order for each quantity. The horizontal axis is the index of the feasible candidate after sorting, not a catalog part number.

The NSGA-II algorithm is used to identify non-dominated solutions. NSGA-II is a fast elitist multi-objective evolutionary algorithm proposed by Deb et al. [14]. It evolves a population of candidate solutions using Pareto-dominance ranking and a crowding-distance measure to maintain diversity along the Pareto front. This is suitable for the considered inductor design problem because winding length, core losses, and core volume are mutually competing objectives; therefore, no single scalar optimum is expected. The settings of the NSGA-II algorithm are as follows: a single decision variable, namely the index of the candidate configuration; three objectives, namely winding length, core losses, and core volume; a search space of 91 candidates; and a population of 20 individuals evolved over 25 generations, which gives 500 evaluations in total. The default operators of the Optuna implementation were used. The use of an evolutionary algorithm is justified by the need to identify the candidate that is optimal with respect to the applied criteria. When the number of combinations grows, or when further objective functions are introduced, such an algorithm can help to isolate the key configurations. Work on the use of the NSGA-II algorithm will be presented in subsequent publications. Figure 5 shows the resulting Pareto front, where the non-dominated solutions are marked separately. The objectives correspond to winding length, core losses, and core volume.




Fig. 5. Multi-objective optimization of feasible inductor candidates: (a) a space representing 91 candidates in terms of volume, winding length, and losses and (b) a space representing the Pareto front from NSGA-II. The objectives are winding length (mm), estimated core losses (mW), and core volume (mm^3).

Table 4. Pareto-optimal inductor configurations returned by the SVD and NSGA-II workflow

ID	Core	μ_r	Turns	Cores	$L(I_{\min})$ (μH)	$L(I_{\max})$ (μH)	l_w (mm)	P_c (mW)	V_c (mm ³)
13	MS-157	60	82	1	571.11	153.80	3666.2	0.2611	11465.3
19	MS-171	75	71	1	593.85	154.48	3507.4	0.3202	15097.6
28	MS-185	125	40	2	582.08	155.70	3160.0	0.3765	32579.4
36	MS-184	125	31	2	551.53	171.11	2936.6	0.4975	45381.5
46	MS-200	125	43	2	573.35	155.87	3134.3	0.3821	33248.9
54	MS-225	147	39	2	570.90	158.49	3021.3	0.4815	43923.4
60	MS-226	90	52	1	544.42	164.48	3184.5	0.5088	30757.8
63	MS-226	147	29	2	543.82	171.24	2659.9	0.6298	61515.7

The complete SVD-based workflow, combined with NSGA-II, yields eight Pareto-optimal configurations, listed in Table 4, also marked in red in Fig. 5. The reported inductance values correspond to the two design current levels used in the feasibility constraints: $I_{\min} = 0$ A and $I_{\max} = 20$ A. The solutions are dominated by cores from the MS-157 to MS-226 range, with nominal permeability classes between 60 and 147 and stacks of one or two cores. The number of turns ranges from 29 to 82. The variation in inductance at $I_{\min}$, equal to 544–594 μH , and at $I_{\max}$, equal to 154–171 μH , reflects the trade-off between the individual design objectives and allows the final configuration to be selected according to application-specific priorities. These configurations have already been selected from among many candidates. It is the role of the designer to choose those that best match the requirements of the particular design, including the economic ones.

6. Conclusions

The conducted study confirms that singular value decomposition (SVD) is an effective tool for reconstructing full inductance characteristics from a small number of selected evaluations. For the considered nonlinear inductor dataset, defined over the range from 1 to 200 turns, five dominant modes are sufficient to reconstruct the characteristic $L(N)$ with an L_2 error below 0.1% in the validation cases under consideration. The largest pointwise relative errors occur for the lowest numbers of turns, where the absolute inductance values are very small. In this region, even a small absolute reconstruction mismatch leads to a relatively large percentage error because the reference inductance appears in the denominator of the relative-error expression. Therefore, the high local relative error at low turn counts should not be interpreted as a substantial degradation of the reconstructed inductance curve. The same procedure can be extended to a wider range of turn counts by constructing a new snapshot dataset over the desired design interval and recomputing the reduced basis and sensing turn counts.

The proposed approach also avoids the need to explicitly linearize the nonlinear permeability characteristic $\mu_r(B)$. Instead, the nonlinearity is represented implicitly in the snapshot matrix, and its dominant effects are captured by the SVD basis. This leads to substantial data compression while preserving the shape of the inductance curve over the turn-count range. The reduced basis reproduces the nonlinearity of the model, not the characteristics of the actual core. Measurement verification will be addressed in further work.

The objective functions used here are simplified: winding length, core losses, and core volume. They were chosen so as to identify the core with the lowest possible losses, which is essential in the design of electrical devices. The use of these objective functions follows from the mechanisms by which losses arise in the objects under study; these losses depend on the frequency, the waveform, the flux swing, and the temperature. The winding length relates to the cost of the material required to wind the core. The reason is purely economic since this may affect the overall cost. In addition, the other two objective functions have design-related considerations. A compact size is desirable, hence the use of the core volume criterion, and the amount of material required to wind the core should be limited, hence the use of the winding length criterion.

The main upfront cost is constructing the snapshot matrix, which requires full numerical calculations across a representative set of core configurations, materials, and operating currents. For very

large datasets, direct SVD may also become computationally expensive. Nevertheless, once the basis and sensing turn counts are available, the combination of sparse sampling and NSGA-II provides a practical workflow for reducing the effort per candidate to a few evaluations, or measurements, at the selected turn counts, while maintaining accuracy suitable for engineering design.

Future work will focus on validating the procedure using measured inductance data, extending snapshot generation to more detailed finite-element models, and incorporating thermal constraints and loss models directly into the multi-objective optimization loop.

Data availability statement

The data and scripts used to generate the results are available from the corresponding author upon reasonable request.

Acknowledgment

The publication costs of this article were covered by the Estonian Academy of Sciences.

References

1. Firpo, P., Ravera, A., Oliveri, A., Lodi, M. and Storace, M. Use of a partially saturating inductor in a boost converter with model predictive control. *Electronics*, 2023, **12**(14), 3013. <https://doi.org/10.3390/electronics12143013>
2. Wang, F., Zhou, H., Zhang, J. and Cao, H. Medium-frequency oscillation in voltage-controlled full-bridge inverter with a non-linear inductor. *Energies*, 2021, **14**(13), 3848. <https://doi.org/10.3390/en14133848>
3. Watson, J. On the nonlinearities of inductors using linear ferrite toroidal cores. *IEEE Transactions on Magnetics*, 1981, **17**(3), 1320–1325. <https://doi.org/10.1109/TMAG.1981.1061205>
4. Salas, R. A. and Pleite, J. Simple procedure to compute the inductance of a toroidal ferrite core from the linear to the saturation regions. *Materials*, 2013, **6**(6), 2452–2463. <https://doi.org/10.3390/ma6062452>
5. Salas, R. A. and Pleite, J. Simulation of waveforms of a ferrite inductor with saturation and power losses. *Materials*, 2014, **7**(3), 1850–1865. <https://doi.org/10.3390/ma7031850>
6. Brown, A. D., Ross, J. N. and Nichols, K. G. Time-domain simulation of mixed nonlinear magnetic and electronic systems. *IEEE Transactions on Magnetics*, 2001, **37**(1), 522–532. <https://doi.org/10.1109/20.914373>
7. Rasilo, P., Martinez, W., Fujisaki, K., Kyyrä, J. and Ruderman, A. Simulink model for PWM-supplied laminated magnetic cores including hysteresis, eddy-current, and excess losses. *IEEE Transactions on Power Electronics*, 2019, **34**(2), 1683–1695. <https://doi.org/10.1109/TPEL.2018.2835661>
8. Sirovich, L. Turbulence and the dynamics of coherent structures. I. Coherent structures. *Quarterly of Applied Mathematics*, 1987, **45**(3), 561–571. <https://doi.org/10.1090/qam/910462>
9. Donoho, D. L. Compressed sensing. *IEEE Transactions on Information Theory*, 2006, **52**(4), 1289–1306. <https://doi.org/10.1109/TIT.2006.871582>
10. Candès, E. J., Romberg, J. K. and Tao, T. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 2006, **59**(8), 1207–1223. <https://doi.org/10.1002/cpa.20124>
11. Drmač, Z. and Gugercin, S. A new selection operator for the discrete empirical interpolation method – improved a priori error bound and extensions. *SIAM Journal on Scientific Computing*, 2016, **38**(2), A631–A648. <https://doi.org/10.1137/15M1019271>
12. Chaturantabut, S. and Sorensen, D. C. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 2010, **32**(5), 2737–2764. <https://doi.org/10.1137/090766498>
13. Gavish, M. and Donoho, D. L. The optimal hard threshold for singular values is $4/\sqrt{3}$. *IEEE Transactions on Information Theory*, 2014, **60**(8), 5040–5053. <https://doi.org/10.1109/TIT.2014.2323359>
14. Deb, K., Pratap, A., Agarwal, S. and Meyerarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 2002, **6**(2), 182–197. <https://doi.org/10.1109/4235.996017>
15. Micrometals. MS – Sendust. Micrometals alloy powder cores catalogue data. <https://www.micrometals.com/products/materials/ms> (accessed 2026-05-25).

Hõreandmestikul põhinev vähendatud mudel mittelineaarsete magnetiliste induktorite kiireks iseloomustamiseks

Kamil Gruszczyk, Paweł Lasek, Mariusz Stępień ja Krzysztof Habelok

Artiklis esitatakse meetod mittelineaarsete magnetiliste induktorite omaduste kiireks hindamiseks ja sobiva induktori valimiseks. Mittelineaarsete magnetmaterjalide puhul sõltub induktiivsus südamiku omadustest, keerduke arvust ja töövoolest, mistõttu võib erinevate lahenduste võrdlemine nõuda suurt hulka arvutusi. Väljapakutud meetod kasutab singulaarväärtuste lahutust (SVD), et kirjeldada induktiivsuse muutumist vähendatud mudeliga. QR-lahutuse abil valitakse seejärel kõige informatiivsemad arvutuspunktid. Uuritud juhul piisas kogu 1–200 keeruga induktiivsuskõvera taastamiseks ainult viiest punktist. Valideerimisel 40 erineva konfiguratsiooniga jäi suurim suhteline L_2 -viga alla 0,1%, kusjuures meetod oli võrreldud lahendustest nii täpsem kui ka kiirem. Vähendatud mudel ühendati NSGA-II optimeerimisalgoritmiga, et valida kommertsandmebaasist sobivaid induktorisüdamikke, arvestades samal ajal mähise pikkust, südamikukadusid ja ruumala. Tulemused näitavad, et meetod võimaldab oluliselt vähendada projekteerimiseks vajalikku arvutusmahtu, säilitades samal ajal piisava täpsuse. Uuring põhineb praegu arvutuslikult genereeritud andmetel ning edasises töös on kavas meetodit kontrollida ka mõõteandmetega.