applied sciences

MDPI

Article

A Reduced-Order Equivalent-Dipole Model for DC Stray Magnetic Fields

Carlo Silano [1], Raffaele Fresa [2], Vincenzo Paolo Loschiavo [1,*] and Antonio Quercia [3]

[1] Dipartimento di Ingegneria, Università degli Studi del Sannio, Corso Garibaldi 107, 82100 Benevento, Italy; c.silano@studenti.unisannio.it
[2] Dipartimento di Ingegneria, Università della Basilicata, Via dell'Ateneo Lucano 10, 85100 Potenza, Italy; raffaele.fresa@unibas.it
[3] Dipartimento di Ingegneria Elettrica e delle Tecnologie dell'Informazione, Università degli Studi di Napoli Federico II, Via Claudio 21, 80125 Napoli, Italy; aquercia@unina.it
* Correspondence: loschiavo@unisannio.it

Abstract

Accurate characterization of stray magnetic fields is essential in high-field devices, where external fields may affect diagnostics and personnel safety. Although finite-element models provide high-fidelity solutions, they may require extensive preprocessing and discretization of large source-free regions, making repeated evaluations costly when the magnetic source is unchanged. This paper proposes a reduced-order equivalent-dipole model that provides an analytical representation of the three-dimensional background stray field after one-time offline calibration. Dipoles are placed at Gauss–Legendre nodes within an auxiliary volume, while symmetry and anti-symmetry conditions are embedded in the source mapping. Their effective moments are identified from magnetic-flux-density data through a regularized linear inverse problem. The method is assessed for the SHiP spectrometer magnet at CERN against a high-fidelity finite-element reference. A controlled comparison of five model orders selects the 125-dipole configuration, reducing the relative vector L^2 error on a common three-dimensional evaluation grid from 21.2% for 27 dipoles to 4.59%. An independently generated dense three-dimensional post-selection dataset confirms the accuracy of the selected model, giving a relative vector L^2 error of 1.77% for $d_{\mathrm{aux}} \geq 0.50$ m. The complementary dense mid-plane assessment gives 0.95% over the same validity region.

Keywords: equivalent sources; Gauss–Legendre source placement; inverse problems; magnetic dipoles; magnetostatics; particle-accelerator magnets; stray magnetic field

1. Introduction

Accurate prediction and reconstruction of DC stray magnetic fields are essential in the design, commissioning, and operation of large magnetic systems. External fields may perturb nearby diagnostics, magnetic sensors, detector electronics, and other field-sensitive equipment. In fusion devices, for example, they can also affect the transport of charged particles within diagnostic systems [1]. In addition, magnetic installations may be subject to occupational-exposure requirements and facility-specific safety rules [2–5]. A representative example is the dipole magnet developed for the Scattering and Neutrino Detector of the SHiP (Search for Hidden Particles) experiment at CERN, which is required to produce a magnetic flux density above 1.2 T over a large instrumented volume while maintaining a very low field in the surrounding region [6]. These requirements make the

Academic Editor: Roberto Zivieri

Received: 14 August 2026
Revised: 15 September 2026
Accepted: 16 September 2026
Published: 19 September 2026

accurate and computationally efficient evaluation of the external field an important part of the magnet-design workflow.

Magnetostatic numerical analyses are commonly performed by means of the finite-element method (FEM), finite-difference schemes, boundary-element methods (BEM), or coupled formulations. FEM is particularly effective for complex geometries and nonlinear ferromagnetic materials. However, the treatment of an unbounded exterior domain requires either a sufficiently large truncated air region, special infinite-domain elements, or a coupling with an exterior-field formulation. Several advanced approaches have been developed to reduce this burden, including infinite-domain finite elements and FEM–BEM or FEM–Green couplings [7,8]. Consequently, the motivation for reduced field models is not that full-order methods are intrinsically unsuitable, but that repeated design iterations, dense three-dimensional field maps, uncertainty studies, and real-time or near-real-time evaluations may remain computationally expensive even when an efficient full-order formulation is available.

In a homogeneous source-free region, the magnetostatic field satisfies Maxwell's equations and can be represented through harmonic scalar or vector potentials. This structure has motivated several physics-consistent field-reconstruction techniques. For accelerator magnets, BEM formulations have been used to reconstruct two-dimensional field maps from measured boundary data, avoiding nonphysical artifacts that may arise from direct interpolation [9] in regions such as the surrounding air, where conventional FE or FD approaches may employ relatively coarse meshes. More recently, local three-dimensional field representations derived from solutions of Laplace's equation have been identified from distributed rotating-coil measurements through Bayesian inversion, thereby incorporating regularization and providing a route toward uncertainty quantification [10]. These studies show that a compact field representation can complement full-order simulation when the quantity of interest is the field in a prescribed source-free region rather than the complete internal electromagnetic state of the device.

A related family of techniques is based on equivalent or auxiliary sources. In the Method of Auxiliary Sources and in the closely related Method of Fundamental Solutions, the unknown field is approximated by a weighted superposition of elementary solutions whose singularities are placed outside the region in which the field is evaluated [11–13]. Such representations satisfy the governing differential equation exactly away from the auxiliary sources and provide inexpensive field evaluation once the source amplitudes have been determined [14]. Their numerical performance, however, depends strongly on source placement, sampling geometry, and regularization: the resulting matrices are generally dense and may be severely ill-conditioned [15,16]. The same issue arises in inverse magnetostatic identification from external field measurements, for which truncated singular-value decomposition, Tikhonov-type methods, and related regularization strategies are commonly required [17]. Therefore, the positivity of a quadrature rule or the analytical form of the elementary field does not, by itself, guarantee stability of the associated inverse problem.

For configurations that also include ferromagnetic materials, reduced source representations have been formulated in terms of equivalent currents, uniformly magnetized prisms, and magnetic dipoles. Analytical brick and dipole models have been investigated for the efficient evaluation of the contribution of ferromagnetic components in tokamaks [18], while equivalent-current formulations, also reported as stick models [19,20], have been proposed for rapid field computation in inductors and permanent magnets [21] and even to mimic the plasma behavior in a tokamak [22]. Accordingly, the methodological contribution of the present work does not lie in the dipole kernel, Gauss–Legendre source placement, symmetry reduction, or regularized inverse identification taken separately. Rather, it lies

https://doi.org/10.3390/app16189298

in their integration into a symmetry-constrained, data-calibrated reduced-order formulation in which the physical source, auxiliary source domain, and observation region are explicitly separated, while model order, calibration clearance, sampling distribution, and regularization are assessed within a common identification and validation framework.

The objective of this work is not to replace the full-order magnetostatic model but to avoid its repeated use whenever the magnetic source remains unchanged. For this reason, the field produced by the original device is represented by a finite set of equivalent magnetic dipoles. Their positions are prescribed, whereas their moments are identified from magnetic-flux-density data obtained from a high-fidelity simulation or, in principle, from measurements. Once this offline calibration has been completed, the external field is evaluated online by direct superposition of closed-form dipole fields. The resulting model is therefore analytical and mesh-free during the evaluation stage.

Figure 1 summarizes the intended use of the method. The full-order model or a set of field measurements is used only in the offline stage to identify the equivalent dipole moments. The calibrated model then provides a reusable background field for repeated evaluations. In a future coupled setting, this background contribution may be superposed with the field computed from a local numerical model, provided that the latter does not significantly alter the magnetic state of the source represented by the equivalent dipoles.

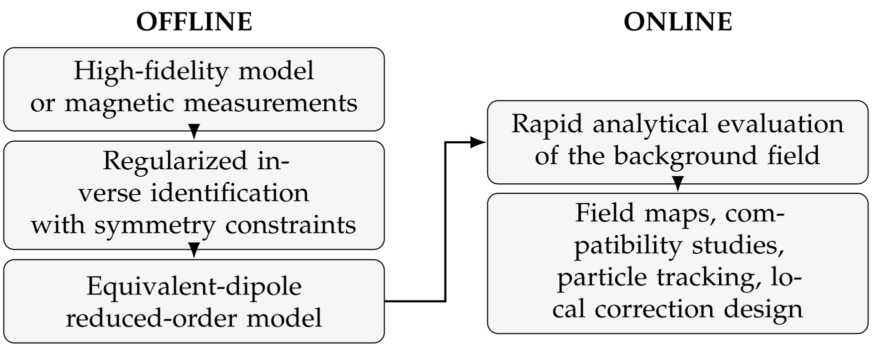

Figure 1. Conceptual workflow of the proposed reduced-order equivalent-source model.

The proposed formulation belongs to the broad family of equivalent-source and auxiliary-source methods. The use of Gauss–Legendre quadrature to replace a distributed magnetic source by an ensemble of equivalent point magnetic dipoles also has a clear precedent in potential-field modeling [23]. However, in the present work, the equivalent sources are not placed on the real geometry of the magnet. Instead, they are located at tensor-product Gauss–Legendre nodes [24–26] of a regular auxiliary region, represented by a parallelepiped in the present implementation, which encloses the magnetic assembly. This distinction is essential. The auxiliary domain is not a geometric approximation of the magnet, and the identified dipole moments are not samples of the local magnetization. The parallelepiped provides a deterministic, regular, and symmetry-compatible support for an equivalent-source expansion. Its boundary also defines the natural limit of the intended reconstruction region.

Let Ω_{mag} denote the physical magnetic assembly and Ω_{aux} the enclosing auxiliary parallelepiped, with $\Omega_{\mathrm{mag}} \subset \Omega_{\mathrm{aux}}$. All equivalent dipoles are contained in Ω_{aux}, and the reconstructed field is intended for points in

$$\Omega_{\mathrm{obs}} \subset \mathbb{R}^3 \setminus \overline{\Omega}_{\mathrm{aux}}. \tag{1}$$

Here, the overbar in $\overline{\Omega}_{\mathrm{aux}}$ denotes the topological closure, ensuring that no observation points fall on the domain boundary.

The method should therefore not be described merely as valid "outside the magnet." Observation points located outside the physical magnet but inside the auxiliary source domain may approach or cross the moving source locations as the quadrature order changes. In that region, convergence can become nonuniform, the identified coefficients can become unstable, and the resulting field should not be interpreted as a reliable approximation of the physical solution. The geometrical relationship among the physical magnet, the auxiliary source domain, the Gauss–Legendre equivalent-source nodes, and the intended observation region is illustrated schematically in Figure 2.

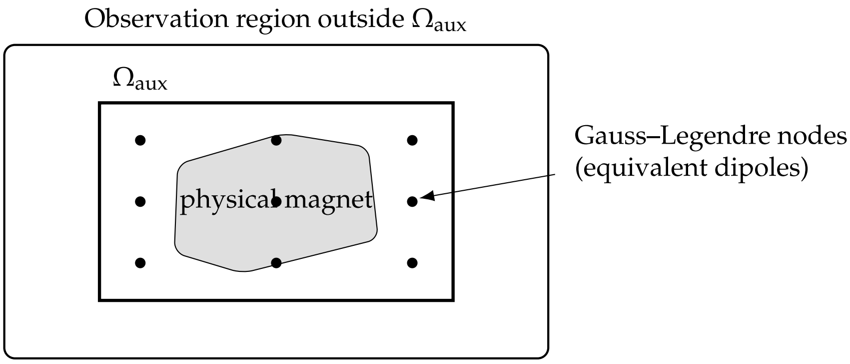

Figure 2. The Gauss–Legendre nodes are defined on an auxiliary parallelepiped enclosing the real magnet. The intended reconstruction region lies outside the auxiliary source domain, not merely outside the physical magnet.

The use of Gauss–Legendre nodes should accordingly be understood as a source-placement strategy rather than as a direct quadrature of the actual magnet geometry. Tensor-product nodes provide a regular family of source configurations of increasing order, preserve reflection symmetries, and provide a convenient non-optimized baseline for source placement. The quadrature weights can be absorbed into the unknown effective moments, which are subsequently identified from field data. This interpretation also avoids attributing a local physical meaning to the equivalent moments.

The reduced-order representation is reusable only while the magnetic source state remains unchanged. In a weakly coupled partition between a source region B and a region of interest A, the total field may be written conceptually as

$$\boldsymbol{B}_{\text{tot}} \simeq \boldsymbol{B}_B^{\text{ROM}} + \boldsymbol{B}_A. \tag{2}$$

Here, $\boldsymbol{B}_B^{\text{ROM}}$ is the calibrated background field generated by the fixed source, whereas $\boldsymbol{B}_A$ denotes the contribution associated with the region of interest. The decomposition is applicable provided that modifications introduced in A do not significantly perturb the current distribution or magnetic state represented in B. Substantial changes in the source state therefore require recalibration.

The SHiP spectrometer magnet at CERN [6] provides a demanding validation case. Its large dimensions, ferromagnetic yoke, coils, and extended stray-field region lead to a complex full-order model and a costly mesh-generation process. In this work, the reference magnetic field is obtained from a high-fidelity finite-element model, whereas the equivalent dipoles are placed inside a regular auxiliary source domain, represented by a parallelepiped in the present implementation. Model order, calibration-sample distribution, and calibration clearance are then assessed under a controlled protocol before the selected model is subjected to dense post-selection validation.

Within this integrated framework, the specific contributions of the present study can be summarized as follows:

- A unified reduced-order formulation for three-dimensional DC stray fields that combines prescribed equivalent-source placement, symmetry-constrained parameterization, and regularized data-driven identification;
- A clear distinction between the physical magnet, the auxiliary source domain, and the external observation region;
- A family of equivalent-source configurations based on tensor-product Gauss–Legendre nodes of the auxiliary source domain;
- Symmetry-constrained parameterization of the equivalent moments, reducing the number of independent unknowns;
- Regularized inverse identification from magnetic-flux-density data, followed by rapid mesh-free online evaluation;
- A controlled numerical assessment of model order, calibration clearance, and calibration-sample distribution, followed by dense post-selection validation against the high-fidelity SHiP finite-element solution.

The remainder of the paper is organized as follows. Section 2 introduces the equivalent-dipole formulation and the regularized inverse problem. Section 3 describes the auxiliary source domain, the symmetry-constrained source mapping, and the investigated model-order family. Section 4 presents the SHiP case study, the calibration and model-selection protocol, the numerical results, and the dense post-selection validation. Section 5 discusses the accuracy–complexity trade-off, the range of applicability, and possible extensions of the method. Finally, Section 6 summarizes the main findings and conclusions.

2. Mathematical Modeling

Let an ideal magnetic dipole with magnetic moment $\boldsymbol{m}$ be located at $\boldsymbol{r}'$. For an observation point $\boldsymbol{r}$, define

$$\boldsymbol{R} = \boldsymbol{r} - \boldsymbol{r}', \qquad R = \|\boldsymbol{R}\|, \qquad \widehat{\boldsymbol{R}} = \frac{\boldsymbol{R}}{R}. \tag{3}$$

In free space, considering μ_0 as the vacuum magnetic permeability, the magnetic vector potential generated by the dipole is [27,28]

$$\boldsymbol{A}(\boldsymbol{r}) = \frac{\mu_0}{4\pi}\frac{\boldsymbol{m}\times\boldsymbol{R}}{R^3}. \tag{4}$$

The corresponding magnetic flux density is obtained as

$$\boldsymbol{B}(\boldsymbol{r}) = \nabla\times\boldsymbol{A}(\boldsymbol{r}). \tag{5}$$

Using the identity

$$\nabla\times(\boldsymbol{m}\times\boldsymbol{w}) = \boldsymbol{m}\,\nabla\cdot\boldsymbol{w} - (\boldsymbol{m}\cdot\nabla)\boldsymbol{w}, \tag{6}$$

with constant $\boldsymbol{m}$, together with

$$\nabla\cdot\frac{\boldsymbol{R}}{R^3} = 0, \qquad \boldsymbol{R}\neq\boldsymbol{0}, \tag{7}$$

one obtains

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi}\frac{3(\boldsymbol{m}\cdot\widehat{\boldsymbol{R}})\widehat{\boldsymbol{R}} - \boldsymbol{m}}{R^3}. \tag{8}$$

Equation (8) is the exact field of an ideal point dipole for $\boldsymbol{r}\neq\boldsymbol{r}'$. The approximation involved in the present method does not concern the dipole kernel itself, but rather the replacement of the original magnetic device by a finite set of equivalent dipoles.

For a continuous distribution of dipole moment density $\boldsymbol{M}(\boldsymbol{r}')$ occupying a volume τ, the vector potential is

$$\boldsymbol{A}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int_\tau \boldsymbol{M}(\boldsymbol{r}') \times \frac{\boldsymbol{r} - \boldsymbol{r}'}{\|\boldsymbol{r} - \boldsymbol{r}'\|^3}\, \mathrm{d}\tau'. \tag{9}$$

Taking the curl with respect to the observation coordinates gives

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int_\tau \frac{3\big[\boldsymbol{M}(\boldsymbol{r}') \cdot \widehat{\boldsymbol{R}}\big]\widehat{\boldsymbol{R}} - \boldsymbol{M}(\boldsymbol{r}')}{R^3}\, \mathrm{d}\tau'. \tag{10}$$

The same continuous magnetization distribution may also be represented by the equivalent volume and surface current densities [28,29]

$$\boldsymbol{J}_{\mathrm{m}v} = \nabla' \times \boldsymbol{M}, \qquad \boldsymbol{J}_{\mathrm{ms}} = \boldsymbol{M} \times \boldsymbol{n}', \tag{11}$$

where $\boldsymbol{n}'$ is the outward unit normal to the boundary S of τ. Accordingly,

$$\boldsymbol{A}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \left[\int_\tau \frac{\boldsymbol{J}_{\mathrm{m}v}(\boldsymbol{r}')}{R}\, \mathrm{d}\tau' + \int_S \frac{\boldsymbol{J}_{\mathrm{ms}}(\boldsymbol{r}')}{R}\, \mathrm{d}S' \right], \tag{12}$$

and

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \left[\int_\tau \boldsymbol{J}_{\mathrm{m}v}(\boldsymbol{r}') \times \frac{\boldsymbol{R}}{R^3}\, \mathrm{d}\tau' + \int_S \boldsymbol{J}_{\mathrm{ms}}(\boldsymbol{r}') \times \frac{\boldsymbol{R}}{R^3}\, \mathrm{d}S' \right]. \tag{13}$$

Equation (13) has the form of the Biot–Savart law and shows the equivalence between the magnetization description and the associated bound-current description.

To obtain the finite-dimensional model used in this work, Equation (10) is approximated by a weighted sum over n equivalent dipoles located at the points $\boldsymbol{r}_i$:

$$\boldsymbol{B}(\boldsymbol{r}) \simeq \frac{\mu_0}{4\pi} \sum_{i=1}^{n} w_i \frac{3\big[\boldsymbol{m}_i \cdot \widehat{\boldsymbol{R}}_i\big]\widehat{\boldsymbol{R}}_i - \boldsymbol{m}_i}{R_i^3}, \tag{14}$$

where

$$\boldsymbol{R}_i = \boldsymbol{r} - \boldsymbol{r}_i, \qquad R_i = \|\boldsymbol{R}_i\|, \qquad \widehat{\boldsymbol{R}}_i = \frac{\boldsymbol{R}_i}{R_i}. \tag{15}$$

The coefficients w_i account for the adopted quadrature rule. If the quadrature weights are absorbed into the dipole moments by defining

$$\boldsymbol{p}_i = w_i \boldsymbol{m}_i, \tag{16}$$

Equation (14) becomes

$$\boldsymbol{B}(\boldsymbol{r}) \simeq \frac{\mu_0}{4\pi} \sum_{i=1}^{n} \frac{3\big[\boldsymbol{p}_i \cdot \widehat{\boldsymbol{R}}_i\big]\widehat{\boldsymbol{R}}_i - \boldsymbol{p}_i}{R_i^3}. \tag{17}$$

In this representation, $\boldsymbol{p}_i$ has units of $\mathrm{A} \cdot \mathrm{m}^2$ and represents the effective magnetic moment of the i–th equivalent dipole. Equation (17) is the form used in the direct field reconstruction.

For compactness, define the dipole interaction tensor

$$\boldsymbol{T}_i(\boldsymbol{r}) = \frac{1}{R_i^3}\Big(3\widehat{\boldsymbol{R}}_i\widehat{\boldsymbol{R}}_i^{\mathsf{T}} - \boldsymbol{I}\Big). \tag{18}$$

Then,

$$\boldsymbol{B}(\boldsymbol{r}) \simeq \frac{\mu_0}{4\pi} \sum_{i=1}^{n} \boldsymbol{T}_i(\boldsymbol{r})\boldsymbol{p}_i. \tag{19}$$

For P observation points, the three Cartesian field components are stacked in the vector

$$\boldsymbol{b} = \begin{bmatrix} B_x(\boldsymbol{r}_1) & B_y(\boldsymbol{r}_1) & B_z(\boldsymbol{r}_1) & \cdots & B_x(\boldsymbol{r}_P) & B_y(\boldsymbol{r}_P) & B_z(\boldsymbol{r}_P) \end{bmatrix}^{\mathsf{T}} \in \mathbb{R}^{3P}. \tag{20}$$

The effective moments of the n equivalent dipoles are collected as

$$\boldsymbol{p} = \begin{bmatrix} \boldsymbol{p}_1^{\mathsf{T}} & \boldsymbol{p}_2^{\mathsf{T}} & \cdots & \boldsymbol{p}_n^{\mathsf{T}} \end{bmatrix}^{\mathsf{T}} \in \mathbb{R}^{3n}. \tag{21}$$

Let $\boldsymbol{K} \in \mathbb{R}^{3P\times 3n}$ denote the matrix assembled from the interaction tensors in Equation (18). The general discrete forward problem can then be written as

$$\boldsymbol{b} = \frac{\mu_0}{4\pi}\boldsymbol{K}\boldsymbol{p} + \boldsymbol{\varepsilon} = \boldsymbol{A}_0\boldsymbol{p} + \boldsymbol{\varepsilon}, \tag{22}$$

where

$$\boldsymbol{A}_0 = \frac{\mu_0}{4\pi}\boldsymbol{K}, \tag{23}$$

and $\boldsymbol{\varepsilon}$ collects numerical, modeling, and measurement errors.

The least-squares estimate of the effective source moments is

$$\widehat{\boldsymbol{p}} = \arg\min_{\boldsymbol{p}} \|\boldsymbol{A}_0\boldsymbol{p} - \boldsymbol{b}\|_2^2. \tag{24}$$

If

$$\boldsymbol{A}_0 = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^{\mathsf{T}} \tag{25}$$

is the singular value decomposition of the general interaction matrix, a Truncated Singular Value Decomposition (TSVD) estimate may be written as [30]

$$\widehat{\boldsymbol{p}}_\tau = \sum_{\sigma_\ell/\sigma_1 \geq \tau} \frac{\boldsymbol{u}_\ell^{\mathsf{T}}\boldsymbol{b}}{\sigma_\ell}\boldsymbol{v}_\ell, \tag{26}$$

where τ is a relative truncation threshold.

Equation (22) represents the general unconstrained inverse problem. The problem-specific reduction associated with source grouping, reflection symmetries, and symmetry-plane constraints is introduced in Section 3.3, where the complete source vector is parameterized as $\boldsymbol{p} = \boldsymbol{S}\boldsymbol{q}$.

Finally, at distances much larger than the dimensions of the equivalent source distribution, $\boldsymbol{R}_i \simeq \boldsymbol{r}$ for every i, and the field reduces to that of the total equivalent magnetic moment

$$\boldsymbol{m}_{\text{tot}} = \sum_{i=1}^{n} w_i\boldsymbol{m}_i. \tag{27}$$

Hence,

$$\boldsymbol{B}(\boldsymbol{r}) \simeq \frac{\mu_0}{4\pi}\frac{3(\boldsymbol{m}_{\text{tot}}\cdot\widehat{\boldsymbol{r}})\widehat{\boldsymbol{r}} - \boldsymbol{m}_{\text{tot}}}{r^3}, \qquad r \gg \max_i \|\boldsymbol{r}_i\|. \tag{28}$$

This expression provides the asymptotic limit of the multi-dipolar model and confirms the expected r^{-3} decay of the reconstructed stray magnetic field. The adopted source-placement strategy provides a systematic and symmetry-compatible baseline for the equivalent-source representation. The specific tensor-product Gauss–Legendre configurations considered in this work are introduced in Section 3. However, optimization of both the shape of the auxiliary domain and the placement of associated sources remains possible, as discussed in Section 5. In the following section, tensor-product Gauss–Legendre nodes are used for the specific parallelepiped implementation adopted in the present SHiP study. The same equivalent-source concept can be extended to other regular auxiliary domains admitting a convenient quadrature or reference-domain mapping.

3. Symmetry-Constrained Equivalent-Dipole Configurations

The general formulation introduced in Section 2 is now specialized to a finite set of equivalent magnetic dipoles inside an auxiliary domain Ω_{aux} that encloses the active magnetic region to be replaced by the reduced-order model. The auxiliary domain is not a geometrical approximation of the real magnet and does not reproduce its material distribution. Its role is instead to provide a regular support for the singular elementary solutions, to permit a deterministic source placement, and to separate the equivalent-source representation from the detailed geometry and finite-element discretization of the original device.

In the implementation considered in this work, Ω_{aux} is a parallelepiped centered at $\boldsymbol{r}_{\mathrm{c}}$, with half-lengths h_x, h_y, and h_z. The complete physical magnetic assembly is contained inside this auxiliary source region, whereas the field reconstructed from the equivalent dipoles is intended for observation points outside it.

3.1. Gauss–Legendre Source Placement

For a tensor-product Gauss–Legendre rule of orders (n_x, n_y, n_z), the equivalent-source positions are

$$\boldsymbol{r}_{abc} = \boldsymbol{r}_{\mathrm{c}} + h_x \xi_a^{(n_x)} \widehat{\boldsymbol{x}} + h_y \xi_b^{(n_y)} \widehat{\boldsymbol{y}} + h_z \xi_c^{(n_z)} \widehat{\boldsymbol{z}}, \tag{29}$$

where

$$a = 1, \ldots, n_x, \qquad b = 1, \ldots, n_y, \qquad c = 1, \ldots, n_z, \tag{30}$$

and $\xi^{(n)}$ denotes a one-dimensional Gauss–Legendre node. The total number of physical dipoles is therefore

$$N = n_x n_y n_z. \tag{31}$$

As discussed in Section 2, the quadrature weights and the geometric Jacobian may be absorbed into the unknown effective moments. The inverse problem consequently identifies directly the effective magnetic moments $\boldsymbol{p}_i$, with units $\mathrm{A} \cdot \mathrm{m}^2$. In the present formulation, Gauss–Legendre nodes should therefore be interpreted primarily as a deterministic source-placement strategy rather than as pointwise samples of the real magnetization.

The use of a parallelepiped is a convenient implementation of the method, but it is not a fundamental requirement. More generally, Ω_{aux} may be any regular enclosing domain for which a stable Gaussian-type quadrature, or an equivalent mapping from a reference domain, can be constructed conveniently. If $\widehat{\Omega}$ denotes a regular reference domain, $\widehat{\boldsymbol{\xi}}_i$ its quadrature nodes, and $\mathcal{F} : \widehat{\Omega} \to \Omega_{\mathrm{aux}}$ a regular mapping, the equivalent-source locations may be written as

$$\boldsymbol{r}_i = \mathcal{F}(\widehat{\boldsymbol{\xi}}_i). \tag{32}$$

Accordingly, the shape and dimensions of the auxiliary source region may in principle be treated as design variables and optimized so as to follow more closely the active magnetic region that is to be replaced by the equivalent model. This possibility may reduce unnecessary empty source volume and may improve the representation of the field close to the active region. When symmetry reduction is employed, however, the chosen auxiliary domain and its node distribution must be compatible with the symmetry group of the source region, or the source-group mapping must be modified consistently.

For completeness, the tensor-product family considered in the development of the method is

$$\begin{aligned} (n_x, n_y, n_z) &= (2,2,2) &&\Longrightarrow && N = 8, \\ (n_x, n_y, n_z) &= (2,2,3) &&\Longrightarrow && N = 12, \\ (n_x, n_y, n_z) &= (2,3,3) &&\Longrightarrow && N = 18, \\ (n_x, n_y, n_z) &= (3,3,3) &&\Longrightarrow && N = 27, \\ (n_x, n_y, n_z) &= (5,5,5) &&\Longrightarrow && N = 125. \end{aligned} \tag{33}$$

Gauss–Legendre rules of different orders are generally non-nested. Equation (33) therefore defines a family of different reduced spaces rather than a conventional nested mesh-refinement sequence. No source order is assumed to be optimal at this stage. The configurations are compared under a common frozen identification protocol in Section 4, where the retained model order is selected from the resulting accuracy–complexity comparison within the prescribed discrete candidate family.

3.2. Reflection Symmetries

The reference SHiP source model exploits the three coordinate planes. With the coordinate convention adopted in this work, its magnetic flux density satisfies

$$\boldsymbol{B}(-x, y, z) = \boldsymbol{T}_x \boldsymbol{B}(x, y, z), \qquad \boldsymbol{T}_x = \mathrm{diag}(1, -1, -1), \tag{34a}$$

$$\boldsymbol{B}(x, -y, z) = \boldsymbol{T}_y \boldsymbol{B}(x, y, z), \qquad \boldsymbol{T}_y = \mathrm{diag}(1, -1, 1), \tag{34b}$$

$$\boldsymbol{B}(x, y, -z) = \boldsymbol{T}_z \boldsymbol{B}(x, y, z), \qquad \boldsymbol{T}_z = \mathrm{diag}(1, 1, -1). \tag{34c}$$

Equivalent sources located at symmetrically equivalent positions—i.e., sharing the same absolute values of their spatial coordinates—are assigned to the same representative vector group.

Let $g(i)$ denote the group associated with source i, and let $\boldsymbol{a}_g$ denote the representative moment in the non-negative octant. The effective moment of source i is generated as

$$\boldsymbol{p}_i = \boldsymbol{T}_x^{\eta_{xi}} \boldsymbol{T}_y^{\eta_{yi}} \boldsymbol{T}_z^{\eta_{zi}} \boldsymbol{P}_{g(i)} \boldsymbol{a}_{g(i)}, \tag{35}$$

where

$$\eta_{\alpha i} = \begin{cases} 0, & r_{i,\alpha} \geq 0, \\ 1, & r_{i,\alpha} < 0, \end{cases} \qquad \alpha \in \{x, y, z\}. \tag{36}$$

The projector $\boldsymbol{P}_g$ eliminates moment components that are incompatible with a source located on a symmetry plane. For the three coordinate planes,

$$\boldsymbol{P}_x = \mathrm{diag}(1, 0, 0), \qquad \boldsymbol{P}_y = \mathrm{diag}(1, 0, 1), \qquad \boldsymbol{P}_z = \mathrm{diag}(1, 1, 0). \tag{37}$$

This construction embeds the required parity relations directly into the reduced basis. The reconstructed background field therefore satisfies the prescribed reflection properties at arbitrary observation points and not merely at the samples used for calibration.

An important distinction is that the symmetries imposed in Equations (34) are the symmetries of the original source region represented by the reduced model, hereafter denoted as region B. They need not be symmetries of a second region A in which the reconstructed stray field is subsequently used. Thus, region A may contain geometrically asymmetric objects, asymmetric sources, or conductors with a non-symmetric material distribution. In that case, only the background contribution generated by B retains the parities imposed above, whereas the total field in A is not required to be symmetric. This differs from a conventional single-domain finite-element symmetry reduction, for which symmetry boundary conditions can be used to reduce the computational domain only when the complete coupled problem—including geometry, materials, sources, and boundary conditions—is invariant under the same transformation. The present background-field decomposition therefore allows the symmetry of the source region B to be exploited even when the region of interest A does not possess that symmetry.

3.3. Reduced Source Mapping

Let

$$G = \left\lceil \frac{n_x}{2} \right\rceil \left\lceil \frac{n_y}{2} \right\rceil \left\lceil \frac{n_z}{2} \right\rceil \tag{38}$$

be the number of representative vector groups. The number of independent scalar coefficients is

$$d = \sum_{g=1}^{G} \mathrm{rank}(\boldsymbol{P}_g). \tag{39}$$

After stacking the Cartesian components of all physical source moments, the complete moment vector is written as

$$\boldsymbol{p} = \boldsymbol{S}\boldsymbol{q}, \tag{40}$$

where

$$\boldsymbol{p} \in \mathbb{R}^{3N}, \qquad \boldsymbol{q} \in \mathbb{R}^{d}, \qquad \boldsymbol{S} \in \mathbb{R}^{3N\times d}. \tag{41}$$

The sparse matrix $\boldsymbol{S}$ contains the group relations, parity transformations, and symmetry-plane constraints.

For P observation points, the reduced forward model is

$$\boldsymbol{b} = \frac{\mu_0}{4\pi}\boldsymbol{K}\boldsymbol{S}\boldsymbol{q} + \boldsymbol{\varepsilon} = \boldsymbol{A}\boldsymbol{q} + \boldsymbol{\varepsilon}, \tag{42}$$

where $\boldsymbol{K}$ collects the closed-form dipole interaction tensors and

$$\boldsymbol{A} = \frac{\mu_0}{4\pi}\boldsymbol{K}\boldsymbol{S}. \tag{43}$$

For the highest-order candidate in Equation (33),

$$N = 125, \qquad G = 27, \qquad d = 51. \tag{44}$$

The distinction among these quantities is important: N is the number of physical point dipoles, G is the number of representative vector groups, and d is the number of independent scalar coefficients. Thus, although the 27 representative vectors nominally contain 81 Cartesian components, the symmetry-plane constraints reduce the number of independent coefficients to 51. Analogous reductions apply to the lower-order configurations.

The source order is selected numerically rather than prescribed a priori. To isolate the effect of model complexity, all configurations in Equation (33) are compared using the same auxiliary domain, calibration data, calibration clearance, frozen spatial split, symmetry constraints, column scaling, and validation-based TSVD procedure. Only the source order and the corresponding reduced dimension are changed. The resulting model-order study is reported in Section 4.3.

4. Numerical Validation for the SHiP Spectrometer Magnet

The proposed reduced-order representation has been assessed using the spectrometer magnet developed for the Scattering and Neutrino Detector of the SHiP experiment at CERN [6]. The reference field is obtained from a three-dimensional COMSOL Multiphysics 6.4 model. The reference finite-element model contains 322,840 elements, with a reported mesh-quality value of 0.92, and exploits the geometric symmetries of the SHiP magnet. The nonlinear magnetic behavior of the ferromagnetic yoke is represented through its material magnetization curve. The geometry, material properties, and three-dimensional finite-element formulation of the SHiP spectrometer magnet are described in detail in Ref. [6]. The FEM solution is therefore regarded here as the numerical reference rather than as an exact ground truth. Consequently, the errors reported below quantify the discrepancy of the reduced model with respect to this reference solution. The present benchmark considers a single magnetostatic source state. Therefore, all calibration and evaluation data originate from the same underlying finite-element solution, although they are sampled on distinct spatial sets. Accordingly, the reported errors quantify the spatial reconstruction accuracy of the reduced model for this fixed source state and should not be interpreted as evidence of transferability across different excitation or magnetization states.

The numerical study is organized so as to distinguish explicitly between parameter identification, model selection, internal holdout testing, and post-selection assessment. The successive stages comprise the selection of the calibration protocol, the comparison of the equivalent-source order, and the post-selection evaluation of the frozen model on spatial datasets not used in the preceding stages.

4.1. Numerical Protocol and Datasets

For the SHiP application, the auxiliary source domain is the parallelepiped centered at the origin with half-lengths

$$h_x = 1.2\ \mathrm{m}, \qquad h_y = 2.0\ \mathrm{m}, \qquad h_z = 3.6\ \mathrm{m}. \tag{45}$$

For an observation point $\boldsymbol{r}$, its Euclidean distance from this domain is

$$d_{\mathrm{aux}}(\boldsymbol{r}) = \left[\sum_{\alpha\in\{x,y,z\}} \max(|r_\alpha - r_{\mathrm{c},\alpha}| - h_\alpha, 0)^2\right]^{1/2}. \tag{46}$$

Two distinct distances are used below. The calibration clearance d_{cal} is the minimum d_{aux} admitted during inverse identification, whereas d_{valid} is a post-calibration empirical boundary defining the external region in which the selected model meets the adopted accuracy criterion.

Two first-octant calibration grids with the same nominal size ($35 \times 35 \times 11$ = 13,475 nodes) were compared. The exponential grid contains a higher density of points near the magnetic source, while the uniform grid is evenly spaced. Both grids span $0 \le x, y \le L_{xy}$, with $L_{xy} = 7.379$ m, and contain 35 coordinates along each of these directions. For $j = 0, \ldots, 34$, the uniform grid uses $x_j = y_j = jL_{xy}/34$, whereas the exponential grid uses $x_j = y_j = L_0[\exp\{(j/34)\ln(1 + L_{xy}/L_0)\} - 1]$, with $L_0 = 1$ m. Both grids use the same 11 coordinates along z, namely $z_k = 0.001 + 0.720k$ m, with $k = 0, \ldots, 10$. After removing non-finite and non-external samples, they provide 9750 and 8111 external observations, respectively. A third shifted regular grid, containing 17,104 finite external points over $0.325 \le d_{\mathrm{aux}} \le 5.0$ m, is used only as a common sensitivity/model-selection grid. Because this grid was used to compare calibration clearances and source orders, it is treated explicitly as a sensitivity/model-selection dataset and is not regarded as an independent final validation set. All reported calibration and common-grid errors are computed from the finite strictly external first-octant samples. Reflection to the other octants leaves the relative vector L^2 error unchanged because both fields satisfy the same verified parities.

For each calibration grid, all finite observations strictly external to the auxiliary domain were partitioned through a deterministic, balanced blockwise spatial train/validation/test split. The occupied coordinate ranges were divided into five, five, and four intervals along the three Cartesian directions, respectively, defining up to $5 \times 5 \times 4 = 100$ rectangular blocks. In the present datasets, 83 of these blocks are occupied. Each occupied block was assigned in its entirety to one subset, thereby preventing spatial leakage between training, validation, and testing.

The assignment targeted nominal point fractions of 60%, 20%, and 20% and was optimized to balance the total number of observations, the distribution of d_{aux}, the distribution of the reference-field magnitude, and their joint distribution. A fixed random seed 260, 805 was used for 6000 multistart assignments, followed by up to 12 local-improvement passes based on single-block moves and pairwise block swaps.

The block assignment was generated once using all finite strictly external calibration points and was then kept frozen throughout the clearance and model-order studies. For each value of d_{cal}, points not satisfying $d_{\mathrm{aux}} \geq d_{\mathrm{cal}}$ were removed from their preassigned subsets without any reassignment. The design matrix was column-scaled before TSVD, and the selected rank was the smallest rank whose validation error lay within 1% of the minimum validation error. After rank selection, the coefficients were re-estimated using the union of the training and validation subsets. The test subset was not used either for coefficient fitting or for TSVD-rank selection and therefore constitutes a spatially blockwise holdout. For the selected exponential-grid configuration at $d_{\mathrm{cal}} = 0.40$ m and $N = 125$, Figure 3 shows the normalized singular values of the column-scaled training matrix and the corresponding validation error as a function of the retained TSVD rank. The prescribed 1% validation rule selects $r = 51$.

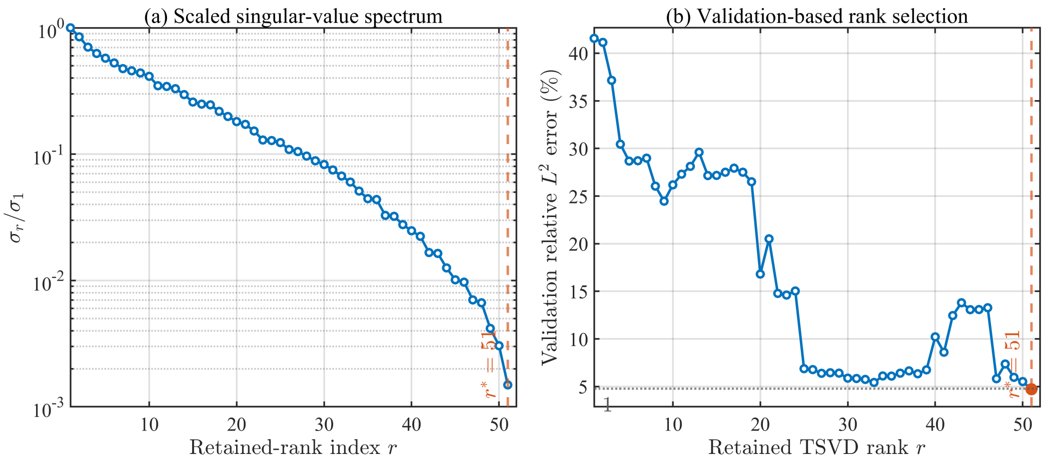

Figure 3. TSVD-rank selection for the final exponential-grid configuration with $d_{\mathrm{cal}} = 0.40$ m and $N = 125$ equivalent dipoles. (**a**) Singular values of the column-scaled training matrix, normalized by the largest singular value. (**b**) Relative validation error as a function of the retained rank. The vertical dotted line identifies the selected rank $r = 51$, determined as the smallest rank whose validation error lies within 1% of its minimum value.

For the subsequent model-order comparison, the exact frozen train/validation/test indices associated with the selected exponential calibration case at $d_{\mathrm{cal}} = 0.40$ m were reused for every investigated source order. Consequently, differences among the model-order results cannot originate from changes in the observations assigned to the identification subsets. The resulting test subset therefore constitutes a spatially blockwise holdout, since entire spatial blocks are excluded from both coefficient fitting and TSVD-rank selection.

The principal global accuracy measure is the relative vector error

$$e_{\mathrm{rel},2} = \frac{\|\boldsymbol{B}_{\mathrm{EDM}} - \boldsymbol{B}_{\mathrm{FEM}}\|_{\mathrm{F}}}{\|\boldsymbol{B}_{\mathrm{FEM}}\|_{\mathrm{F}}}. \tag{47}$$

where $\boldsymbol{B}_{\text{EDM}}$ and $\boldsymbol{B}_{\text{FEM}}$ denote the fields reconstructed by the equivalent-dipole model (EDM) and computed by the reference finite-element model (FEM), respectively. Furthermore, the subscript "F" denotes the Frobenius norm, which computes the square root of the sum of the absolute squares of its elements, whereas the subscript "2" denotes the standard Euclidean norm (or L^2 norm) in $\mathbb{R}^3$, applied to the local magnetic flux density vector evaluated at a specific spatial point $\mathbf{r}$. Here, $\boldsymbol{B}_{\text{EDM}}, \boldsymbol{B}_{\text{FEM}} \in \mathbb{R}^{P\times 3}$ are point-by-component field matrices. Their Frobenius norm is equivalent to the Euclidean norm of the corresponding stacked field vector defined in Equation (20).

For the error maps, the pointwise relative vector error is

$$\varepsilon_{\text{rel}}(\boldsymbol{r}) = 100\frac{\|\boldsymbol{B}_{\text{EDM}}(\boldsymbol{r}) - \boldsymbol{B}_{\text{FEM}}(\boldsymbol{r})\|_2}{\|\boldsymbol{B}_{\text{FEM}}(\boldsymbol{r})\|_2}. \tag{48}$$

Equation (48) is evaluated only where $\|\boldsymbol{B}_{\text{FEM}}(\boldsymbol{r})\|_2 > 0$. Since this pointwise normalization may amplify the apparent error where the reference field is very small, the global vector norm $e_{\text{rel,2}}$ remains the primary quantitative metric. Distance-resolved errors are also evaluated in 0.05-m shells. The quantity $d_{5\%}$ denotes the lower boundary of the first shell containing at least 50 finite external points for which the shell relative vector error is below 5% and remains below 5% for all subsequent shells. The 5% threshold is adopted here as a pragmatic accuracy criterion for defining the empirical validity boundary and should not be interpreted as a regulatory or safety limit.

4.2. Calibration Clearance and Sampling Strategy

Table 1 reports representative results of the clearance sweep for the 125-dipole candidate. The parameter r_{TSVD} denotes the truncation rank, i.e., the number of retained singular values. Using all external points ($d_{\text{cal}} = 0$) strongly destabilizes the identification: validation retains only one singular direction for the exponential grid and four for the uniform grid, with common-grid errors of 61.447% and 49.544%, respectively. Increasing the clearance first improves stability and predictive accuracy, whereas excessive clearance removes the near-source information required to reconstruct the near/intermediate field.

Table 1. Representative calibration-clearance results for the 125-dipole model. Errors are evaluated on the common three-dimensional sensitivity/model-selection grid.

	Exponential Calibration Grid			**Uniform Calibration Grid**		
d_{cal} **[m]**	r_{TSVD}	$e_{\text{rel,2}}$ **[%]**	$d_{5\%}$ **[m]**	r_{TSVD}	$e_{\text{rel,2}}$ **[%]**	$d_{5\%}$ **[m]**
0.00	1	61.4	–	4	49.5	–
0.20	51	5.09	0.55	50	5.95	0.55
0.40	51	4.59	0.45	50	5.61	0.50
0.60	49	6.88	0.55	48	6.32	0.55
1.00	50	8.90	0.55	50	11.0	0.60

The best overall compromise for the exponential calibration is

$$d_{\text{cal}} = 0.40 \text{ m}, \tag{49}$$

for which all 51 singular directions are retained and the scaled condition number is 6.71×10^2. At the same clearance, the exponential grid gives a lower global common-grid error than the uniform grid (4.59% versus 5.61%) and improves the near/intermediate-field shells. Farther from the source, the two strategies become very similar, with the uniform grid slightly better in some far-field metrics. The result is therefore interpreted as a sampling-distribution effect, not as a universal superiority of the exponential grid. For the selected exponential-grid calibration at $d_{\mathrm{cal}} = 0.40$ m, the frozen blockwise partition retains 5567, 1762, and 1677 observations in the training, validation, and test subsets, respectively. The test subset, which is not used for either coefficient identification or TSVD-rank selection, yields a relative vector error of $e_{\mathrm{rel},2} = 4.20\%$.

Sensitivity to imperfect calibration data was assessed by adding independent Gaussian perturbations with prescribed relative L^2 amplitudes of 0.1%, 0.5%, 1%, 2%, and 5% to the training and validation field values. For each nonzero noise level, 50 realizations were considered, while the source configuration, calibration clearance, spatial split, column scaling, and TSVD selection rule were kept unchanged. Up to 2% noise, the selected rank remained $r = 51$ in every realization and the median error on the unperturbed blockwise test subset remained between 4.19% and 4.20%; at 2% noise, its 5th–95th percentile interval was 4.10–4.32%. At 5% noise, the median test error was 4.26%, although 9 of the 50 realizations selected $r = 33$, increasing the 95th-percentile test error to 9.87%.

4.3. Model-Order Selection

After fixing the calibration protocol above, the five source configurations of Section 3 were compared while keeping the auxiliary domain, exponential calibration grid, $d_{\mathrm{cal}} = 0.40$ m, frozen spatial split, column scaling, symmetry constraints, and TSVD selection rule unchanged. Thus only the reduced-model order is varied, as summarized in Table 2.

Table 2. Model-order sensitivity under the frozen identification protocol. A dash indicates that the persistent 5% shell criterion is not reached within the investigated external domain.

(n_x, n_y, n_z)	N	G	d	r_{TSVD}	$\kappa_2(\overline{A})$	$e_{\mathrm{rel},2}$ [%]	$d_{5\%}$ [m]
$(2,2,2)$	8	1	3	1	1.10	85.3	–
$(2,2,3)$	12	2	5	3	1.83	66.9	–
$(2,3,3)$	18	4	8	5	3.81	55.0	–
$(3,3,3)$	27	8	12	11	8.24	21.2	2.00
$(5,5,5)$	125	27	51	51	671	4.59	0.45

The transition from $N = 27$ to $N = 125$ is decisive: the global common-grid error decreases from 21.2% to 4.59%, while $d_{5\%}$ decreases from 2.00 m to 0.45 m. For $d_{\mathrm{aux}} \geq 0.50$ m, the cumulative relative vector error falls from 13.7% to 1.71%. Although the highest-order system is more ill-conditioned, the validation rule selects $r = 51$ for the present unperturbed calibration dataset, so that no singular direction is truncated in this particular fit.

Within the investigated discrete candidate family, these results justify retaining

$$N = 125 \tag{50}$$

as the best-performing tested configuration. This selection should not be interpreted as establishing a global accuracy–complexity optimum; a denser assessment of intermediate and higher tensor-product orders would be required for that purpose.

4.4. Final Validity Domain and Post-Selection Validation

For the selected exponential 125-dipole model, the common-grid shell analysis gives $d_{5\%} = 0.45$ m. A conservative operational value

$$d_{\mathrm{valid}} = 0.50 \text{ m} \tag{51}$$

is adopted. The imposed reflection parities are satisfied to floating-point accuracy, with relative defects of order 10^{-16}.

After the calibration distribution, calibration clearance, source order, regularization procedure, and validity boundary had been fixed, an additional dense uniform three-dimensional COMSOL dataset was generated exclusively for post-selection assessment. The grid has a spacing of 0.10 m and contains 232,560 nominal points; among these, 177,743 finite external observations satisfy $d_{\text{aux}} \geq 0.50$ m. No coefficient fitting or model-selection decision was performed using these data. On this three-dimensional post-selection dataset, the frozen $N = 125$ model gives a relative vector L^2 error of 1.77%, a vector MAE of 5.52 μT, and a 95th-percentile absolute vector error of 21.4 μT.

The two dense COMSOL planes reported below were likewise generated after the model-development choices had been frozen and are retained as complementary spatially resolved assessments of the same reduced model. On the physically representative mid-plane $z = 0$, the 43,635 finite points satisfying $d_{\text{aux}} \geq 0.50$ m give

$$e_{\text{rel},2} = 0.95\%. \tag{52}$$

The vector MAE and 95th-percentile absolute vector error are 5.40 μT and 21.6 μT, respectively. The distance-resolved result independently confirms the adopted validity boundary: the shell $0.40 \leq d_{\text{aux}} < 0.50$ m gives 5.01%, whereas the next shell, $0.50 \leq d_{\text{aux}} < 0.60$ m, gives 2.09% and the error continues to decrease with distance. The corresponding field and error maps are shown in Figure 4.

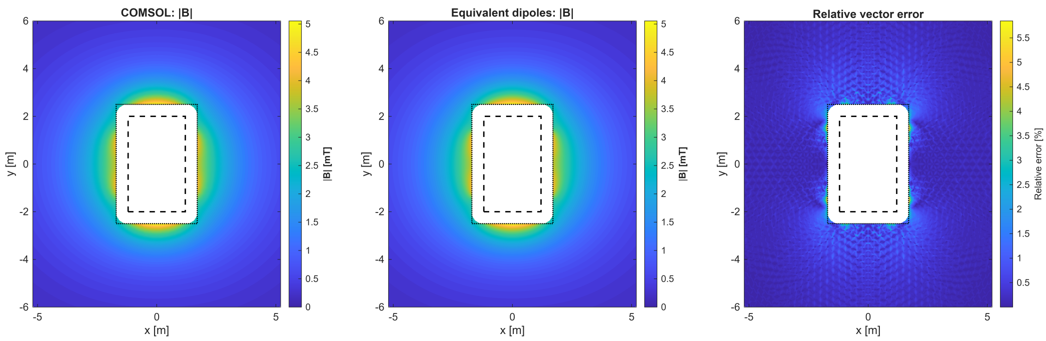

Figure 4. Dense post-selection comparison on the mid-plane $z = 0$. The first two panels show the COMSOL and equivalent-dipole field magnitudes; the third shows the pointwise relative vector error defined in Equation (48). Only the region $d_{\text{aux}} \geq 0.50$ m is displayed for quantitative comparison. The complete map is obtained by reflection of the simulated quarter-plane using the verified field parities.

A second dense plane at $z = 4.125$ m, shown in Figure 5, is entirely outside the conservative validity boundary and provides a complementary check over 47,879 finite external points, for which

$$e_{\text{rel},2} = 2.74\%. \tag{53}$$

Unlike the mid-plane, this plane does not intersect the auxiliary source region because its minimum clearance is 0.525 m. It therefore provides a complementary fully external validation of the same frozen model.

4.5. Computational Cost and Online Evaluation

The computational cost of the proposed approach was assessed by separating the one-time offline stage from the subsequent online field-evaluation stage. All timings were obtained on the same computational platform, i.e., a workstation equipped with an 11th Gen Intel® Core™ i7-1165G7 @2.80 GHz and NVIDIA GeForce MX450, using COMSOL Multiphysics® 6.4 and MATLAB® R2025b. Unless otherwise stated, figure rendering and file-writing operations were excluded from the timing measurements.

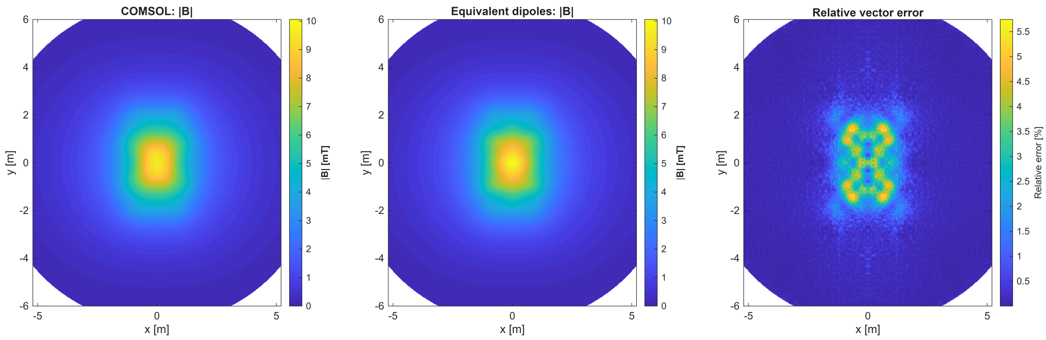

Figure 5. Dense post-selection comparison on the plane $z = 4.125$ m. The plane is entirely external to the auxiliary source region and has a minimum clearance of 0.525 m. The first two panels show the COMSOL and equivalent-dipole field magnitudes; the third shows the pointwise relative vector error.

The full-order finite-element solution required 1936 s. Evaluation and export of the field data used for calibration required an additional time of about 2 s. Once the reference data were available, assembly of the reduced inverse problem for the selected $N = 125$ configuration required 1.30 s, while column scaling, singular-value decomposition, validation-based rank selection, and final refitting required 0.061 s. The EDM-specific offline identification therefore required 1.36 s, whereas the complete one-time offline workflow, including generation and extraction of the finite-element reference data, required approximately 1939 s, as summarized in Table 3.

Table 3. Computational cost of the full-order finite-element model and the selected equivalent-dipole model.

Operation	Wall Time
FEM stationary solution	1936 s
FEM calibration-data evaluation/export	≈2 s
EDM matrix assembly	1.30 s
EDM scaling, SVD, rank selection, and final fit	0.061 s
Dense 3-D EDM evaluation (177,743 points)	0.65 s
Dense 3-D FEM evaluation/export (177,743 points)	≈2 s

For the dense three-dimensional set of 177,743 observation points, online evaluation of the calibrated equivalent-dipole model required 0.65 s, whereas evaluation and export at the same points from the already computed full-order finite-element solution required about 2 s. The corresponding online speed-up was therefore approximately 3.10.

These results clarify the computational role of the proposed model. In the present FEM-calibrated implementation, the equivalent-dipole approach does not eliminate the one-time computational cost required to establish the reference magnetic state. Rather, this cost is incurred during the offline stage and is subsequently amortized over repeated field evaluations. Once calibrated, the magnetic state is represented by a compact set of equivalent-dipole coefficients, so that subsequent evaluations do not require re-solving the full-order problem or retaining and interrogating its complete discretized representation.

5. Discussion

The results confirm that the proposed model is a data-calibrated reduced-order representation of the external field of a source whose magnetic state is known during the offline stage. The equivalent dipoles are not a discretization of the real magnetization; rather, they provide a compact analytical basis for reproducing the field outside the auxiliary source support.

5.1. Identification Strategy and Model Complexity

The numerical study shows that calibration clearance, sampling distribution, and model order are coupled design choices. Including points arbitrarily close to Ω_{aux} does not improve the reduced model: the rapidly varying near-boundary field and the singular elementary basis make the inverse problem more difficult, and validation then suppresses most singular directions. Conversely, excessively large clearances remove observations that constrain the near/intermediate field. The selected value $d_{\mathrm{cal}} = 0.40$ m represents the best compromise for the present configuration, whereas $d_{\mathrm{valid}} = 0.50$ m is an independent empirical statement about the predictive domain.

The calibration-grid comparison likewise shows that sampling should be designed for the region of interest. The exponential grid improves the near/intermediate-field reconstruction, while the uniform grid becomes slightly more accurate in some far-field metrics; therefore, no universal superiority of one distribution is claimed.

Finally, the number of dipoles is properly interpreted as a reduced-model order, not as a finite-element mesh density. The $N = 125$ model is retained because the improvement over $N = 27$ is substantial: the common-grid error decreases from 21.2% to 4.59% and $d_{5\%}$ from 2.00 m to 0.45 m. Although the scaled condition number increases, the validation rule selects $r = 51$ for the present dataset. The richer source space is therefore required to represent spatial content that is not captured by the lower-order models, particularly close to the intended validity boundary.

5.2. Generality of the Equivalent-Source Formulation

This section summarizes the broader applicability of the proposed formulation, generalizing the conceptual points introduced earlier in the manuscript beyond the specific case study.

While the parallelepiped geometry used in the SHiP study is highly convenient—as it supports a straightforward tensor-product Gauss–Legendre construction and facilitates reflection-group mappings—it is merely one possible implementation. As noted previously, this bounding box must not be identified with the physical geometry of the magnet. The theoretical framework can be extended to any suitable auxiliary enclosing solid that admits a robust quadrature rule or a convenient mapping from a reference domain. Therefore, its shape and dimensions may in principle be treated as design variables and optimized to follow the active source region more closely. Such an optimization should balance enclosure of the physical sources, conditioning, distance from the intended observation region, and compatibility between the node distribution and the symmetry group used in the reduced mapping.

The imposed symmetry constraints belong to the source region B represented by the reduced model and need not be shared by the local region A. Hence, asymmetric geometry, materials, conductors, or local sources may be treated in A while retaining the symmetry-reduced background field generated by B. This differs from a monolithic FEM symmetry reduction, which requires the complete coupled problem to be invariant under the same transformation.

With reference to the weakly coupled decomposition introduced in Equation (2), the present DC model provides the background contribution generated by region B, provided that the reaction field associated with region A does not appreciably alter the source state used for calibration.

5.3. Limitations and Outlook

The present coefficients are calibrated from a high-fidelity finite-element solution. Identification from measurements is possible in principle, but sensor uncertainty, positioning errors, incomplete coverage, and systematic offsets would then need explicit treatment. Moreover, the reported d_{valid} is configuration-dependent: changes in the auxiliary support, source order, operating state, sampling distribution, or regularization require renewed distance-resolved verification.

Gauss–Legendre placement also does not by itself regularize the inverse problem. Instead, column scaling and TSVD remain important. Future developments may therefore address optimized auxiliary domains, adaptive or multiple source layers, mixed dipole–multipole bases, measurement-based calibration, uncertainty propagation, and parametric or time-dependent coefficient models.

A possible extension of the present formulation concerns weakly coupled magneto-quasistatic problems with a time-varying background field. In this case, the full-order model may be solved at a sequence of selected time instants, and the corresponding reduced coefficients may be identified from the total magnetic field at each sampled state. The equivalent-source representation does not require the individual contributions to the total field to be separated; therefore, imposed currents, ferromagnetic magnetization, and eddy-current contributions within source region B can all be implicitly incorporated into the time-dependent coefficient vector $\boldsymbol{q}(t)$. The resulting coefficient snapshots may then be interpolated to provide the background excitation for a local time-dependent problem in region A. The temporal sampling must, however, be sufficiently fine to resolve the fastest field variations relevant to the intended application. In particular, transient field peaks associated with strong induced currents must be adequately sampled to avoid their loss during interpolation. As in the static formulation, this weakly coupled use requires the electromagnetic back-reaction of region A on the source state represented in region B to remain negligible. This time-dependent extension is not validated in the present work and is therefore left for future investigation.

6. Conclusions

A symmetry-constrained equivalent-dipole reduced-order model has been developed for the rapid reconstruction of three-dimensional DC stray magnetic fields. The method represents the external field of a fixed magnetic source through a finite superposition of closed-form point-dipole fields whose effective moments are identified by a one-time regularized inverse problem.

For the SHiP spectrometer case study, five symmetry-constrained model orders were compared under the same identification protocol. The final configuration contains $N = 125$ physical dipoles arranged into $G = 27$ representative vector groups and requires $d = 51$ independent scalar coefficients. Compared with the $N = 27$ configuration, the common three-dimensional evaluation-grid error decreases from 21.2% to 4.59%, providing a quantitative justification for the selected model order.

After the calibration distribution, clearance, regularization, model order, and validity boundary had been fixed, an additional dense three-dimensional COMSOL dataset was generated exclusively for post-selection assessment. Over approximately 1.78×10^5 finite external points satisfying $d_{\text{aux}} \geq 0.50$ m, the selected model gives a relative vector L^2 error of 1.77%, with a vector MAE of 5.52 μT and a 95th-percentile absolute vector error of 21.4 μT. The complementary dense mid-plane assessment gives a relative vector L^2 error of 0.95% over 4.4×10^4 points in the same validity region. Together with the distance-resolved analysis, these results support the conservative empirical validity distance $d_{\text{valid}} = 0.50$ m adopted for the final model.

The proposed representation complements rather than replaces the full-order magnetostatic analysis. Once the magnetic state of the source has been established, the calibrated model provides a compact analytical representation for mesh-free online evaluation of the external field provided that the source state remains unchanged. By contrast, substantial changes in geometry, excitation, or magnetization generally require a new calibration.

Author Contributions: Conceptualization, A.Q. and R.F.; methodology, R.F., V.P.L., A.Q. and C.S.; software, V.P.L. and C.S.; validation, R.F., V.P.L., A.Q. and C.S.; formal analysis, R.F., V.P.L., A.Q. and C.S.; investigation, R.F., V.P.L., A.Q. and C.S.; data curation, V.P.L. and C.S.; writing—original draft preparation, R.F., V.P.L., A.Q. and C.S.; writing—review and editing, R.F., V.P.L., A.Q. and C.S.; visualization, V.P.L. and C.S.; supervision, A.Q., R.F. and V.P.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data and MATLAB routines supporting the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (OpenAI, GPT-5.6 Sol) for language editing, consistency checks, and assistance in reviewing MATLAB routines. All scientific decisions, numerical results, interpretations, and final text were reviewed and verified by the authors, who take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Jagielski, M.; Chernyshova, M.; Malinowski, K.; Quercia, A.; Fornal, T. Magnetic field influence on electron transport in planned GEM-based radiated power diagnostic for DEMO. *Phys. Plasmas* **2024**, *31*, 113905. [CrossRef]
2. CERN HSE Unit. Safety Guideline SG-NIR-1-0-1: Exposure to Static Magnetic Fields. Available online: https://edms.cern.ch/ui/file/2975978/LAST_RELEASED/SG-NIR-1-0-1_EN.pdf (accessed on 28 July 2026).
3. European Parliament and Council of the European Union. Directive 2013/35/EU of the European Parliament and of the Council of 26 June 2013 on the minimum health and safety requirements regarding the exposure of workers to the risks arising from physical agents (electromagnetic fields). *Off. J. Eur. Union* **2013**, *L 179*, 1–21.
4. International Commission on Non-Ionizing Radiation Protection (ICNIRP). Guidelines on limits of exposure to static magnetic fields. *Health Phys.* **2009**, *96*, 504–514. [CrossRef] [PubMed]
5. Feychting, M. Health effects of static magnetic fields—A review of the epidemiological evidence. *Prog. Biophys. Mol. Biol.* **2005**, *87*, 241–246. [CrossRef] [PubMed]
6. Ahdida, C.; Albanese, R.; Alexandrov, A.; Anokhina, A.; Aoki, S.; Arduini, G.; Atkin, E.; Azorskiy, N.; Back, J.J.; Bagulya, A.; et al. The magnet of the scattering and neutrino detector for the SHiP experiment at CERN. *J. Instrum.* **2020**, *15*, P01027. [CrossRef]
7. Lobry, J. A FEM-Green Approach for Magnetic Field Problems with Open Boundaries. *Mathematics* **2021**, *9*, 1662. [CrossRef]
8. Schröder, J.; Reichel, M.; Birk, C. An Efficient Numerical Scheme for the FE-Approximation of Magnetic Stray Fields in Infinite Domains. *Comput. Mech.* **2022**, *70*, 141–153. [CrossRef]
9. Biccioli, L.; Di Rienzo, L.; Petrone, C.; Russenschuck, S. Boundary Element Method Reconstruction of Two-Dimensional Magnetic-Field Maps from Measured Boundary Data in Accelerator Magnets. *IET Sci. Meas. Technol.* **2019**, *13*, 60–66. [CrossRef]
10. Ion, I.G.; Liebsch, M.; Simona, A.; Loukrezis, D.; Petrone, C.; Russenschuck, S.; De Gersem, H.; Schöps, S. Local Field Reconstruction from Rotating Coil Measurements in Particle Accelerator Magnets. *Nucl. Instrum. Methods Phys. Res. Sect. A* **2021**, *1011*, 165580. [CrossRef]
11. Papakanellos, P.J.; Tsitsas, N.L.; Anastassiu, H.T. The Method of Auxiliary Sources (MAS) in Computational Electromagnetics: A Comprehensive Review of Advancements over the Past Two Decades. *Electronics* **2024**, *13*, 3520. [CrossRef]
12. Kaklamani, D.; Anastassiu, H. Aspects of the method of auxiliary sources (MAS) in computational electromagnetics. *IEEE Antennas Propag. Mag.* **2002**, *44*, 48–64. [CrossRef]
13. Cheng, A.H.D.; Hong, Y. An Overview of the Method of Fundamental Solutions—Solvability, Uniqueness, Convergence, and Stability. *Eng. Anal. Bound. Elem.* **2020**, *120*, 118–152. [CrossRef]
14. Stratton, J.A. *Electromagnetic Theory*; John Wiley & Sons: Hoboken, NJ, USA, 2007.
15. Alves, C.J.S. On the Choice of Source Points in the Method of Fundamental Solutions. *Eng. Anal. Bound. Elem.* **2009**, *33*, 1348–1361. [CrossRef]

16. Antunes, P.R.S.; Santos, V.; Serranho, P. The MFS-SVD Method for the Laplace Equation in Three Dimensions. *SIAM J. Sci. Comput.* **2025**, *47*, A454–A471. [CrossRef]
17. Formisano, A.; Martone, R. Different Regularization Methods for an Inverse Magnetostatic Problem. *Int. J. Appl. Electromagn. Mech.* **2019**, *60*, S49–S62. [CrossRef]
18. Chiariello, A.G.; Formisano, A.; Fresa, R.; Ledda, F.; Martone, R.; Pizzo, F. Effective magnetic field computation in Tokamaks in presence of magnetic materials. *Fusion Eng. Des.* **2015**, *96*, 641–644. [CrossRef]
19. Haus, H.A.; Melcher, J.R. *Electromagnetic Fields and Energy*; Prentice Hall: Englewood Cliffs, NJ, USA, 1989.
20. McGreivy, N.; Zhu, C.; Gunderson, L.M.; Hudson, S.R. Computation of the Biot–Savart line integral with higher-order convergence using straight segments. *Phys. Plasmas* **2021**, *28*, 082111. [CrossRef]
21. Sokolov, A.; Kennedy, M.P.; Blokhina, E. Method of Equivalent Currents for the Calculation of Magnetic Fields in Inductors and Magnets with Application to Electronics. In *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*; IEEE: Piscataway, NJ, USA, 2019; pp. 1–4.
22. Braams, B.J. The interpretation of tokamak magnetic diagnostics. *Plasma Phys. Control. Fusion* **1991**, *33*, 715–748. [CrossRef]
23. von Frese, R.R.B.; Hinze, W.J.; Braile, L.W.; Luca, A.J. Spherical-Earth Gravity and Magnetic Anomaly Modeling by Gauss–Legendre Quadrature Integration. *J. Geophys.* **1981**, *49*, 234–242.
24. Olver, F.W.J.; Olde Daalhuis, A.B.; Lozier, D.W.; Schneider, B.I.; Boisvert, R.F.; Clark, C.W.; Miller, B.R.; Saunders, B.V.; Cohl, H.S.; McClain, M.A. *NIST Digital Library of Mathematical Functions*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2025.
25. Trefethen, L.N.; Bau, D. *Numerical Linear Algebra*; Society for Industrial and Applied Mathematics (SIAM): Philadelphia, PA, USA, 2022.
26. Stroud, A.H. *Approximate Calculation of Multiple Integrals*; Prentice-Hall: Englewood Cliffs, NJ, USA, 1971.
27. Chow, T.L. *Introduction to Electromagnetic Theory: A Modern Perspective*; Jones & Bartlett Learning: Englewood Cliffs, NJ, USA, 2006.
28. Jackson, J.D. *Classical Electrodynamics*; John Wiley & Sons: Hoboken, NJ, USA, 1998.
29. Van Bladel, J.G. *Electromagnetic Fields*; Wiley-IEEE Press: Hoboken, NJ, USA, 2007.
30. Brunton, S.L.; Kutz, J.N. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*; Cambridge University Press: Cambridge, UK, 2019.