Politecnico di Torino

DENERG Department of Energy "Galileo Ferraris"

Master's Degree Course in Electrical Engineering

Master's Thesis

Design of a Permanent Magnet Array for Head Imaging in a Low-Field Portable MRI System

Student
Fausto Rubicondo

Supervisor
Prof. Fabio Freschi
Co-Supervisor
Prof. Feng Liu

Academic Year 2024/2025

Abstract

In this thesis the problem of optimization of a permanent magnetic structure for low-field imaging applications is addressed, with particular attention to the homogeneity of the B_z field within a spherical region of interest defined as FOV (Field of View).
The whole work is based on the use of a numerical code in MATLAB environment, based on a Coulomb model for magnetic field calculation. In the initial phase, the code was subjected to an accurate validation by comparison with established software for magnetostatic analysis commonly used in academic and industrial environments. The good agreement between the results confirmed the reliability of the numerical approach adopted.
Subsequently, several magnetic configurations from the literature have been analyzed to assess their performance in terms of intensity and homogeneity of the magnetic field generated in the FOV. Each configuration was simulated and compared systematically, with the aim of identifying the most promising geometry to be optimized.
Among the options studied, one promising initial configuration was selected, which was then subjected to an extensive process of global optimization using a genetic algorithm (ga) to explore the space of possible geometries in depth. During the course of the work, the configuration has been progressively modified and enhanced, through the introduction of new structural elements and geometric variants to further improve its homogeneity and efficiency.
With the proposed magnetic arrays, it was possible to obtain fields characterized by homogeneity of less than 25,000 ppm within a spherical FOV of 20 cm in diameter, with magnetic field intensity varying between 0.16 T and 0.19 T, depending on the specific changes made to the configuration.
Although the results obtained represent a good goal in terms of homogeneity and intensity of the field,the current structure is not yet suitable for imaging applications since it lacks of gradient fields, essential for spatial encoding in MRI. Nevertheless, the configurations developed in this thesis provides provide a solid foundation for future developments towards more advanced 3D imaging solutions and integration with portable diagnostic systems.

Contents

List of Figures

List of Tables

Introduction

Magnetic resonance imaging (MRI) is one of the most advanced and versatile imaging techniques in medicine, as it provides detailed information about soft tissue without the use of ionising radiation. Historically, the development of MRI has been driven by the search for increasingly sharp and precise images, promoting the use of stronger magnetic fields. This has led to the spread of high and ultra-high field systems capable of exceeding 3 Tesla and providing high diagnostic performance, especially in neurology and oncology.

However, this technological direction also has significant structural and operational limitations. High-field MR systems are based on superconducting coils, that require cryogenic cooling and have high acquisition and maintenance costs. In addition, these devices require controlled, shielded and fixed environments, making them inaccessible in decentralised environments or where logistical resources are limited.

At the same time, an alternative line of research is emerging oriented towards the use of weaker magnetic fields and permanent magnets with the aim of developing portable, affordable and sustainable magnetic resonance systems. These systems, known as low-field MRI, are gaining interest due to their potential applicability in point-of-care (PoC) environments such as emergency rooms, ambulances, territorial clinics or areas with low technology availability. Their use could revolutionise the very concept of imaging, bringing MRI directly to the patient's bedside.

However, the reduction of the magnetic field also brings with it new technical challenges, in particular the difficulty of generating sufficiently homogeneous and stable fields in clinically interesting volumes such as a head or an extremity. The development of magnetic sources based on permanent magnet arrays capable of offering a good compromise between compactness, intensity and field homogeneity, is an extremely current and open research topic.

This is the context of the work carried out in this thesis, which aims to contribute to the definition of new geometries and optimization strategies for low-field systems, The aim is to provide a solid basis for future developments in the field of compact and portable magnetic resonance imaging.

The thesis is organized as follows:

- **Chapter 1** provides the fundamental physical principles of magnetic resonance and discusses the rationale behind the growing interest in low-field magnetic resonance;
- **Chapter 2** describes the MATLAB-based code developed for magnetic field

computation and presents its validation through comparison with well-established simulation codes;

- **Chapter 3** presents the analysis and simulation of various magnetic configurations reported in the literature, evaluating their performance with the goal of identifying the most promising initial geometry for further optimization;
- **Chapter 4** focuses on the selection of the initial magnetic configuration and its first-stage optimization through the addition of supplementary magnetic components, aimed at identifying the most effective direction for further structural development;
- **Chapter 5** carries out a more refined optimization of the magnetic structure by exploring different configuration options and comparing them to identify the most effective one;
- **Chapter 6** compares the optimized structure with a reference design from the literature and proposes future developments aimed at addressing the integration of gradient systems required for imaging applications.

1 Basics of Magnetic Resonance

1.1 Physical Principles of MRI

Magnetic Resonance Imaging (MRI) belongs to a broader class of techniques that exploit the fundamental physical phenomenon known as Nuclear Magnetic Resonance (NMR). NMR refers to the ability of certain atomic nuclei to interact with external magnetic fields in such a way that they can absorb and later re-emit electromagnetic energy under specific conditions. The discovery of NMR in the mid-20th century marked a major milestone in physics and chemistry. Initially developed as a tool for analyzing molecular structures, providing insight into atomic-scale interactions and chemical composition, NMR later evolved into the foundation for non-invasive imaging techniques, including MRI.

Nuclear Magnetic Resonance lies on a fundamental quantum mechanical property of certain atomic nuclei: the *nuclear spin* [1]. Since the mass of the proton is not zero, the rotation of the proton leads to angular momentum. Since the electric charge of the proton is distributed in a small spherical volume, this rotation constitutes a circulating charge around its axis. This motion generates a microscopic current, which in turn generates a magnetic field called the *Magnetic moment* $\boldsymbol{\mu}$.

The link between the angular momentum $\boldsymbol{J}$ and the magnetic moment $\boldsymbol{\mu}$ of a nucleus is given by

$$\boldsymbol{\mu} = \gamma \times \mathbf{J}$$

where γ is a constant characteristic of a given nucleus called *gyromagnetic ratio*. For nuclei containing an odd number of protons or neutrons, it is not possible to achieve complete cancellation of individual spins, resulting in a non-zero net angular momentum. Among the various NMR-active nuclei, hydrogen-1 (^{1}H) plays a particularly significant role in magnetic resonance imaging. Hydrogen atoms consist of a single proton and, consequently, a non-zero magnetic moment. This simple nuclear structure makes hydrogen very sensitive to magnetic fields and ideal for NMR-based techniques. An additional advantage is the abundance of hydrogen

nuclei in the human body, primarily due to the high water content and the presence of organic molecules rich in hydrogen atoms. When placed in a static external magnetic field B_0 the magnetic moments of the hydrogen nuclei tend to align either parallel (low-energy state) or antiparallel (high-energy state) to the field direction. Although the difference in population between these two states is very small, it is sufficient to produce a macroscopic net magnetization aligned with B_0. However, due to their intrinsic magnetic properties, the magnetic moments do not remain perfectly aligned with the external field. Instead, they tend to rotate around the direction of the magnetic field, following a motion known as *precession* (Fig.1.1).

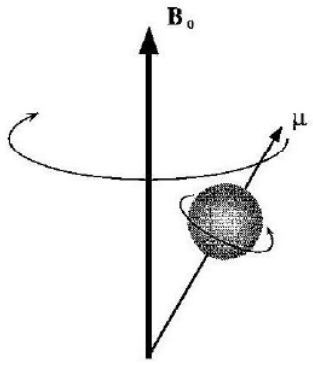

Figure 1.1: Precession motion around B_0 axis, adapted from [1]

The speed at which this precession occurs is not arbitrary, but is determined by a simple and fundamental relationship: the *Larmor frequency*, which expresses how fast the magnetic moments rotate around the field axis. This frequency ω_0 is directly proportional to the strength of the external magnetic field and is given by:

$$\omega_0 = \gamma \times B_0$$

At equilibrium, the net magnetization vector resulting from the collective behavior of the hydrogen nuclei is aligned along the direction of the static magnetic field B_0. However, in order to detect a signal, this equilibrium must be disturbed by a dynamic change in the orientation of the magnetization. This is achieved by using an RF coil (*radio frequency*) that generates an oscillating magnetic field B_1 perpendicular to the static magnetic field B_0. The RF field is generated by passing an alternating current through the coil at a frequency that corresponds exactly to the Larmor frequency of the nuclei. When the RF pulse is applied, the magnetic moments of the nuclei absorb energy and are deflected from their alignment with B_0. After the RF pulse is turned off, the hydrogen nuclei maintain the excess energy absorbed during the excitation phase while the external field B_0 once again

dominates the behavior of the nuclear spins. As $\boldsymbol{B}_1$ is shut down the transverse magnetization does not stay still but keeps rotating around the direction of $\boldsymbol{B}_0$ at the Larmor frequency. During this rotation, the nuclei gradually release the absorbed energy, returning to their equilibrium condition where the net magnetization is realigned along $\boldsymbol{B}_0$. This relaxation process is associated with the emission of a weak electromagnetic signal, *Free Induction Decay* (*FID*). This FID signal, detected by a receiver coil, represents the total behavior of the nuclei in the sample. However, on his own, the signal does not have any clear spatial information. Spatial details are extracted by applying controlled magnetic field variations called *gradients*. These changes alter the resonance conditions of the nuclei according to their position within the sample. This means that each point of the sample is subjected to a slightly different magnetic field and, for this reason, precesses at a different frequency. By exploiting this position-frequency correspondence it is possible to associate each detected signal with its point of origin enabling, through proper acquisition and reconstruction techniques, the composition of magnetic resonance images.

1.2 High-Field MRI vs Low-Field MRI

Since the growing interest in enhancing the accessibility and flexibility of MRI systems, it is fundamental to consider the difference and the trade-offs between high-field and low-field technologies. Although both approaches rely on the same fundamental physical principles, they differ significantly in terms of technical implementation, performance, cost, and clinical applicability.
The following section highlights the key advantages and limitations of each system across several critical aspects of MRI operation:

- **Signal-to-Noise Ratio (*SNR*):** is a crucial factor in determing MRI image quality, a higher SNR enhances image resolution, improves tissue contrast, and increases diagnostic accuracy. In MRI, the strength of the detected signal is directly linked to the sample's net magnetization, the more $\boldsymbol{B}_0$ increases, the more nuclear spins align with it, meaning that the SNR improves with increasing magnetic field strength. For this reason, high-field MRI systems inherently provide significantly higher SNR compared to low-field systems;

- **Spatial Encoding:** as anticipated in Sec 1.1, it refers to the process of determining the spatial origin of the MR signal in the sample. In conventional high-field MRI systems, spatial encoding is performed using gradient coils that

produce magnetic fields varying linearly along the orthogonal axes, ensuring a precise localization of the MR signal within the imaging volume and the reconstruction of high-resolution images.
Instead many low-field MRI systems adopt alternative encoding strategies like the use of *Spatial Encoding Magnetic Fields* (**SEMs**). SEMs are specially designed magnetic field shapes, often non-linear, that inherently contain spatial information. Unlike gradient coils, SEMs are not separate active components, they are indeed integrated into the main magnetic structure and typically generated by permanent magnets or specially designed static field configurations. Although SEMs reduce hardware complexity by removing the need for powered gradient systems, they often require more advanced reconstruction methods and may be less adaptable in defining the imaging volume.
Although SEMs represent a promising approach for portable, low-field MRI systems, high-field systems with conventional gradient coils still offer superior performances since their capability of generating highly linear and controlled magentic gradients;

- **Portability and accessibility:** Due to their lower magnetic field strength, these systems do not require the extensive infrastructure required by high-field MRI machines. As a result, low-field MRI systems are lighter, more compact, and require significantly less power to operate. These characteristics make them particularly suitable for bedside diagnostics, rural or under-resourced healthcare environments, emergency rooms, and even in-field use in military or disaster scenarios;

- **Cost and infrastructures:** High-field MRI machines involve significant costs, not only for the scanner itself but also for, installation, operation and long-term maintenance. For example a new $1.5\,\boldsymbol{T}$ or $3\,\boldsymbol{T}$ MRI system can cost anywhere from approximately $1 million to over $3 million, depending on the specific model and features, with a general estimate of about $1 million per Tesla of magnetic field strength. These systems typically require complex infrastructure, including radiofrequency shielding, controlled ventilation, and cryogenic cooling systems for the superconducting magnets. Moreover, their operation and maintenance demand highly trained personnel, restricting their use to specialized facilities. These requirements also lead to higher energy consumption and impose strict constraints on site selection and environmental conditions (like reinforced flooring, temperature and humidity control, and protection from electromagnetic interference).

On the other end low-field MRI units offer a more cost-effective and logistically flexible alternative. Their simpler design, often based on permanent magnets or resistive electromagnets, eliminates the need for cryogenic cooling and allows installation in standard clinical settings with minimal modifications. With costs generally in the range of tens of thousands of dollars, these systems offer a more affordable solution for healthcare environments with limited budgets or decentralized infrastructures.

Regardless of system architecture or clinical context, the ability to generate a strong and highly homogeneous magnetic field remains essential for producing high-quality MR images. This becomes even more critical in low-field systems, where signal levels are intrinsically lower and every improvement in field performance can significantly enhance diagnostic value. Moreover, in low-field configurations, the relative strength of spatial encoding fields or other interfering sources can approach that of the main $\boldsymbol{B}_0$ field. This increases the risk of signal distortion and reduces the accuracy of spatial encoding. For this reason, improving both the absolute strength and the homogeneity of $\boldsymbol{B}_0$ is not just desirable, but essential to maintain image quality and ensure reliable system performance.
However, achieving both high field strength and high homogeneity typically necessitates the use of large and heavy magnets. In order to combine these two competing requirements, one possible strategy is to miniaturize the hardware and restrict the region of interest to a smaller imaging volumes. This trade-off between imaging volume and field quality is a key aspect of system design and forms the basis of the present work, which focuses on optimizing magnetic configurations for low-field MRI systems operating under strict spatial constraints.

2 Magnetic Field Computation

2.1 Coulombian Approach

The computation of the magnetic field generated by permanent magnets is a fundamental step in the design and optimization of magnetic systems.
In most studies found in the literature, this field is calculated using the current model, where the magnetization of the material is replaced by an equivalent distribution of electric currents. These equivalent 'bound' currents, circulating on the magnet's surface or within its volume, serve as the sources of the resulting magnetic field. Despite its physical clarity, this method requires evaluating surface or volume integrals over the full current distribution, which can be challenging in the presence of complex geometries. Consequently, the current model becomes computationally intensive, particularly in three-dimensional cases.
In this work, a different method is adopted: the *Coulombian approach* or *charge model* [2].
This method replaces the magnetization with a distribution of fictitious magnetic charges on the surfaces of the volume.
The magnet is considered ideal, with a magnetization vector $\boldsymbol{J}$ that remains constant in both magnitude and direction throughout the entire volume. Another crucial assumption is that the relative permeability of the magnet is equal to 1, i.e. $\mu_m = \mu_0$ (or $\mu_r = 1$ valid for NdFeB and SmCo magnets). This condition is essential for several reasons reasons. First, when the magnet's permeability equals that of free space, there are no discontinuities in the magnetic field inside the material and the entire magnetic field can be attributed to surface effects (see Figure 2.1).

Figure 2.1: Example of polyhedron-shaped magnet. (A) Constant and uniform magnetization J [T]. (B) Equivalent Coulombian representation [2]

Second, this assumption guarantees linearity between the magnetic field H and the magnetic flux density B, through the relation $B = \mu_0 H$.. Finally, this assumption preserves the formal analogy with electrostatics, in which the magnetic scalar potential behaves similarly to the electrostatic potential generated by charge distributions.
The magnetic field H is derived from the scalar magnetic potential

$$\vec{H} = -\vec{\nabla}\phi_m = \frac{\sigma}{4\pi\varepsilon_0} \iint_S \frac{\vec{r}}{r^3}\, dr \tag{2.1}$$

where the geometrical parameter r is the distance between the point where the field is to be calculated and the magnetic charges.
The analytical calculation of the magnetostatic scalar potential (or the magnetic field) generated by a uniformly charged surface of arbitrary polygonal shape is a complex task. In general, the integral in 2.1 does not admit a closed-form solution. Nonetheless, an exact analytical result can be obtained in a specific case: when the surface is a right triangle. This is the reason why this calculation method is based on the decomposition of any polygon in a series of right triangles.
In the context of this thesis, this approach offers an effective compromise between analytical precision and computational efficiency, and is extensively employed to evaluate the magnetic fields produced by complex magnet configurations.

2.2 Code Implementation

The magnetic field simulations presented in this work are based on a computational tool called `PmLab`, originally developed by the thesis supervisor and provided as a foundation for this project. This code implements the Coulombian approach described in Sec.2.1 and it has been employed throughout the thesis to study the magnetic field produced by various magnet configurations.
Fundamentally, the code computes the magnetic field of a uniformly magnetized

polyhedral magnet by decomposing its surfaces into a mesh of right triangular elements and, from this general formulation different common geometries (like cubes, bricks, cylinders etc...) are extracted, providing flexibility and adaptability to both simple and complex magnet arrangements.
Given the size of the full implementation, only the most relevant parts are shown here, with the goal of illustrating the core principles behind the simulation process. The code below shows the main routine for computing the magnetic field produced by a uniformly magnetized polyhedral magnet.

2.2.1 *PmPolyhedron* Initialization

```
% construct empty pmPolyhedron object
            p = pmPolyhedron;
            % assemble different objects
            for i = 1:nargin
                p.faces = [p.faces; varargin{i}.faces+size(p.
                    vertices,1)];
                p.vertices = [p.vertices; varargin{i}.vertices];
                p.charge = [p.charge; varargin{i}.charge];
            end
            % assign the generic polyhedron class
            p.shape = 'polyhedron';
            % assign the object of the first PM
            p.object = varargin{1}.object;
        end
```

The first section of the code handles the initialization of a new `pmPolyhedron` object. The use of `varargin` allows the function to accept multiple input components, making it possible to assemble complex geometries from simpler predefined shapes. Each component adds its own vertices, faces, and magnetic charges to the resulting polyhedral structure.

2.2.2 Magnetic Charge Computation

```
        % magnetic charge
        function p = pmCharge(p)
            % check arguments
            arguments
                p {mustBeA(p,'pmPolyhedron')}
            end
            V = bsxfun(@minus,p.vertices,p.origin)*p.rotation;
```

```
            % outward unit vectors
            uA = versor(triArea(V,p.faces));
            % charge density
            p.charge = dotProduct(p.BrFun,uA);
            % relative tolerance to avoid useless contributions
            RELTOL = 1e-6;
            % charge threshold
            CHARGETOL = RELTOL*max(abs(p.charge));
            % remove "small" contribuitons
            p.charge(abs(p.charge) < CHARGETOL) = 0;
        end
```

The `pmCharge` function computes the surface magnetic charge for each triangular face of the polyhedron, taking into account the local magnetization direction and the geometric characteristics of the element.
At first, all vertex coordinates are converted from the global frame to the magnet's local coordinate system, according to the definition of `BrFun`. This is achieved by applying a translation and rotation based on the magnet's local origin and rotation matrix. Then the `TriArea` function is used to calculate the area vectors of triangle starting from its vertices, these vectors are then normalized using `versor` to return the outward-pointing unit normals.
The surface charge is then calculated through the dot product between the magnetization vector and the corresponding unit normal. A relative tolerance is then introduced in order to eliminate insignificant contributions to the field.

2.2.3 Magnetic Charge Calculation

```
% evaluate magnetic field
  function H = pmField(p,Q,options)
      % empty space magnetic permeability
      Mu0 = 4e-7*pi;
      % preallocation
      H = zeros(size(Q));
      % field calculation
      for i = 1:size(p.faces,1)
          % delete small contributions
          if p.charge(i) ~= 0
          [~,I1Grad] = intGreen2d(p.vertices(p.faces(i,:),:),Q)
          H = H-p.charge(i)/(4*pi*Mu0)*I1Grad;
          end
      end
      if options.inside
```

```
16              % points inside the permanent magnet
                idIn = inSolid(p,Q);
                % correction for points inside the PM
                H = H+bsxfun(@times,idIn,(p.BrFun*p.rotation.')/Mu0);
            end
21      end
```

The `pmField` function calculates the magnetic field in a specified set of observation points in space `Q`. The function begins by initializing the output matrix H, which holds the magnetic field vectors at each point in `Q`, and proceeds, through a `for` cycle, to all triangular surface elements. For each triangle with a nonzero surface charge, the `intGreen2d`[1] function is called, which provides the exact analytical expression for the gradient of the scalar magnetic potential generated by that triangle at all specified observation points. The resulting gradient vector is then scaled and subtracted from the total field, in accordance with the definition $H = -\nabla\sigma_m$. When the optional inside flag is active, the function identifies observation points located inside the magnet and adds a constant correction term (equal to the magnetization vector divided by μ_0) to the computed field. This reflects the uniform internal field expected in a magnetized body with relative $\mu_r = 1$.

2.2.4 Basic Elements for Practical Implementation

Although the general polyhedral implementation offers great versatility, practically only a limited set of specific shapes has been extensively used in the simulations. To simplify the construction of the configurations that will be shown in the following chapters, two fundamental geometric shapes have been derived from the general formulation: the rectangular prism (`pmBrick`) and the circular sector (`pmTile`), as presented in Chapter 2 of [7].

Other basic geometries, including cylinders, spheres, and toroids, have been implemented within the same framework. However, with the exception of the cylindrical element, the remaining shapes have not been used in the present work.
Alongside the geometric components, a set of specific functions has been introduced to enable common operations like:

- `Scale`: scales the magnet by a factor k;
- `Translate`: Moves the magnet in space by applying a translation defined by a vector V;

[1] See M. Fabbri, "Magnetic Flux Density and Vector Potential of Uniform Polyhedral Sources", IEEE Transactions on Magnetics, vol. 44, no. 1, Jan 2008, pp. 32–36

- `Revolve`: performs a spatial rotation of the magnet using a given rotation matrix R;
- `pmAssemble`: combines two or more magnets into a single object of `pmPolyhedron` type for joint processing and field evaluation;
- `plot`: generates a 3D plot of the magnet geometry, including the polarization vector for visual reference.

2.3 Code Validation

In order to ensure the accuracy of the results obtained with the `PmLab` code provided for this work, its performances have been validated by comparing them with other well-established computational tools for magnetic field analysis: `MagPyLib` and `MagTetris`. These tools were selected because of their proven accuracy and widespread use in both academic and industrial contexts.
Magpylib is an open-source Python library for modeling and simulating static magnetic fields generated by sources such as permanent magnets, current loops, and magnetic dipoles. Its computational approach is based on the concept of equivalent surface-bound currents. Essentially, the magnetization is modeled through surface tangential currents whose entity is proportional to the magnetization vector and the normal to the surface itself. The MagPy library supports basic geometric shapes such as cuboids, cylinders, and spheres, and is particularly suitable for the design and optimization of magnetic systems in both scientific and engineering contexts.
MagTetris, instead, is a MATLAB-based analytical tool specifically designed to compute the magnetic field produced by configurations of rectangular permanent magnets. As PmLab its computational core is based on the Coulombian model. Unlike Magpylib, which supports full 3D field evaluation, MagTetris is specifically designed to analyze the field on two-dimensional observation planes, surely offering less computational costs but at the cost of reduced versatility.
Validation is carried out on three benchmark cases, offering a comprehensive assessment of the accuracy and efficiency of PmLab.

2.3.1 Case 1: Magnetic Brick

The first benchmark case consists of a single magnetic brick with uniform magnetization. (see Figure 2.2)

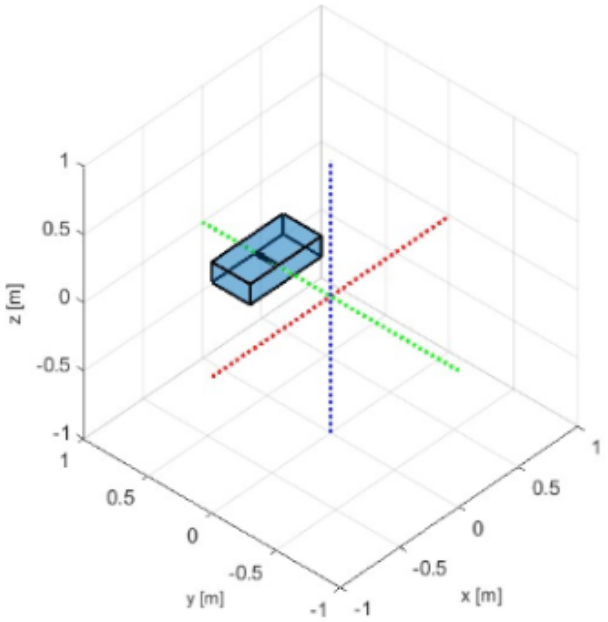

Figure 2.2: Magnetic Brick geometry and orientation

Parameter	Value
Dimensions	$D_x = 0.6$ m, $D_y = 0.3$ m, $D_z = 0.15$ m
Position	$\mathbf{P}_0 = [0,\ 0.5,\ 0]$ m
Magnetization	$B_r = 1.4$ T, directed along the y-axis
Observation points (x-axis)	From $[-1,\ 0,\ 0]$ to $[1,\ 0,\ 0]$
Observation points (y-axis)	From $[-1,\ 0,\ 0]$ to $[1,\ 0,\ 0]$
Observation points (z-axis)	From $[-1,\ 0,\ 0]$ to $[1,\ 0,\ 0]$
Evaluated points per axis	10,000

PmLab - MagPyLib Comparison

The comparison was carried out along the three principal axes (x, y, z), evaluating the magnitude of the magnetic flux density $|B|$.
In Figures 2.3–2.5, the magnetic field profiles computed by PmLab and MagPyLib are presented. The curves show a perfect matching in both magnitude and spatial behavior, validating the consistency and the accuracy of the PmLab formulation for this configuration.

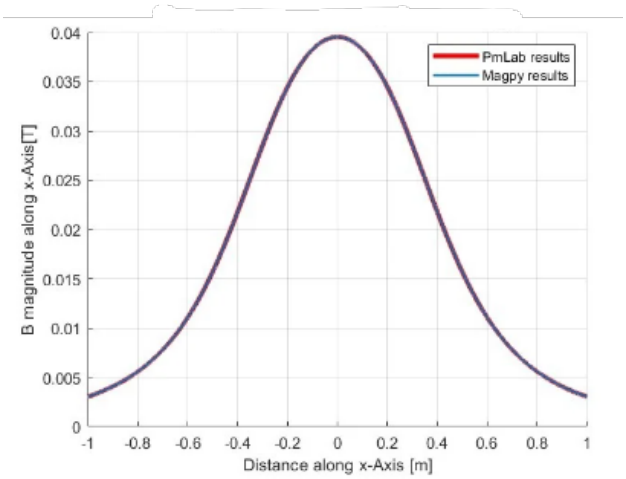

Figure 2.3: Magnetic flux density magnitude |**B**| along the y-axis

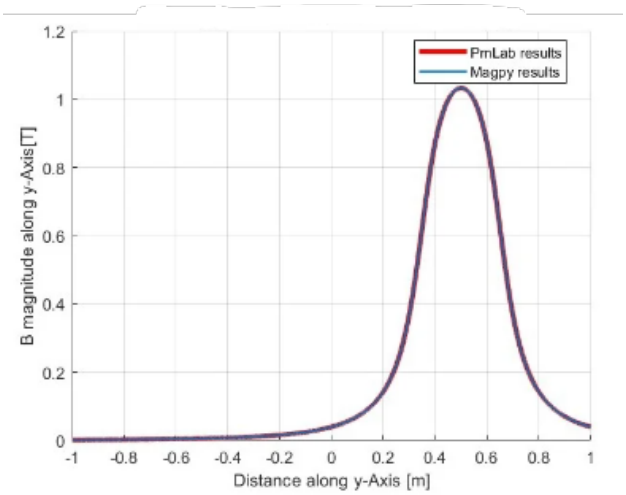

Figure 2.4: Magnetic flux density magnitude |**B**| along the x-axis

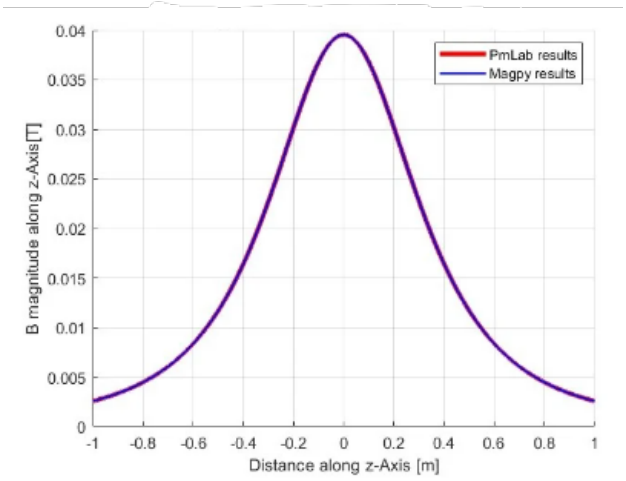

Figure 2.5: Magnetic flux density magnitude |**B**| along the z-axis

PmLab - MagTetris Comparison

Since MagTetris field analysis is restricted to two-dimensional geometries, the comparison was limited to field evaluations along the midplanes of the magnet, specifically along the $x-$ and $y-$ axes at $z = 0$.
Also in this case, as shown in Figures 2.6 and 2.7, the curves produced by PmLab and MagTetris are in excellent agreement, confirming the correctness of the field computation in a planar setting.

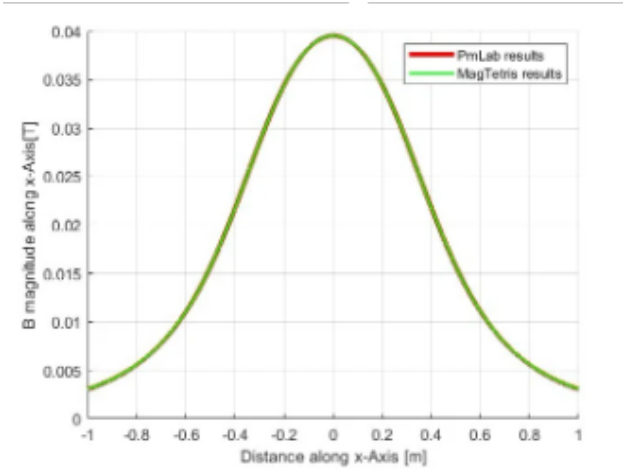

Figure 2.6: Magnetic flux density magnitude |**B**| along the x-axis

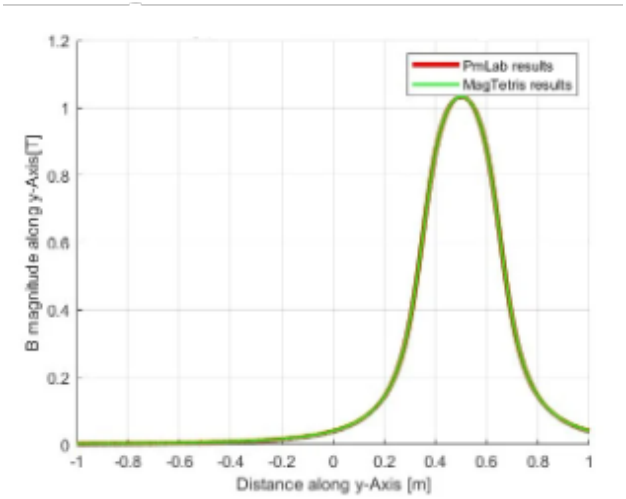

Figure 2.7: Magnetic flux density magnitude |**B**| along the y-axis

Before moving on to the next benchmark case, it is important to examine a numerical detail related to how the magnetic field is computed inside the magnetized volume. Figure 2.8 shows the magnetic field computed along the y-axis using PmLab, with and without the correction that excludes observation points located

inside the magnet volume. In the black curve, no internal-point correction is applied. The result is that $|B|$ appears to be underestimated in the region where the observation point falls inside the magnet. This derives directly from the definition of the Coulombian model, which sees the magnetization as a superficial charge distribution and becomes undefined within the source region. To overcome this issue a conditional check is included (see Section 2.2.3) and the resulting curve (in green) represents the expected field trend for a uniformly magnetized brick.

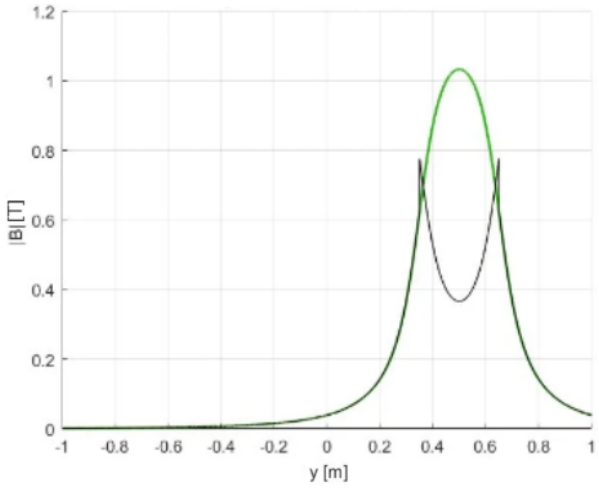

Figure 2.8: Effect of internal-point correction in PmLab

Code	**Calculation Time [s]**
PmLab	1.101
MagPy	5.971
MagTetris	0.583[1]

Table 2.1: Case 1 - Computation time required by each code

Table 2.1 reports the computation time required by each code to evaluate the magnetic field produced by a single brick magnet. As expected, MagTetris exhibits the shortest execution time, owing to its simplified 2D formulation.
In contrast, PmLab achieves an effective compromise between computational speed and geometric versatility. Compared to MagPyLib, it ensures markedly shorter execution times, without sacrificing 3D capabilities.

[1]Note: MagTetris does not allow field evaluation along the z-axis and was therefore tested under a reduced configuration.

2.3.2 Case 2: Halbach Configuration with Bricks

The second benchmark case is about a circular Halbach array composed of magnetized bricks. A Halbach array, as it will be better explained in Section 3.2, is a particular arrangement of permanent magnets where the magnetization of the elements varies progressively along the structure in order to intensify the field on one side of the structure while reducing it on the other. This is done by assigning to each magnet a magnetization vector that rotates with a fixed angular step, as shown in Figure 2.9.

Parameter	**Value**
Number of rings	1
Magnets per ring	24
Magnet dimensions	$D_x = 0.02$ m, $D_y = 0.02$ m, $D_z = 0.1$ m
Ring radius	$R = 0.16$ m
Magnetization	Halbach pattern
Observation points for B_x	100 points from $[-0.1,\ 0,\ 0]$ to $[0.1,\ 0,\ 0]$
Evaluation grid (colormap)	$R_{\text{base}} = 0.1$ m 100 points along x from $-R_{\text{base}}$ to R_{base} 100 points along y from $-R_{\text{base}}$ to R_{base} Total number of points: 100,000

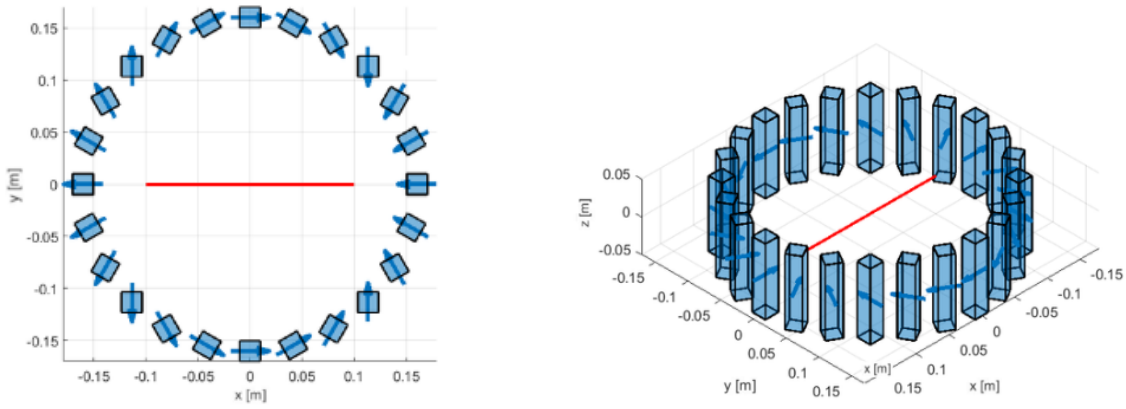

Figure 2.9: Case 3 - Halbach disposition and magnetization with bricks, in red the observation points for B_x

PmLab - MagPyLib Comparison

In this case the comparison was perfromed along two distinct evaluation domains:

1. a one-dimensional line across the center of the array;

2. a two-dimensional plane parallel to the array cross-section.

The analysis is restricted to the B_x component of the magnetic flux since it represents the dominant contribution in this specific Halbach array. Figure 2.16 shows the one-dimenional B_x comparison. The profiles align very closely throughout the entire domain

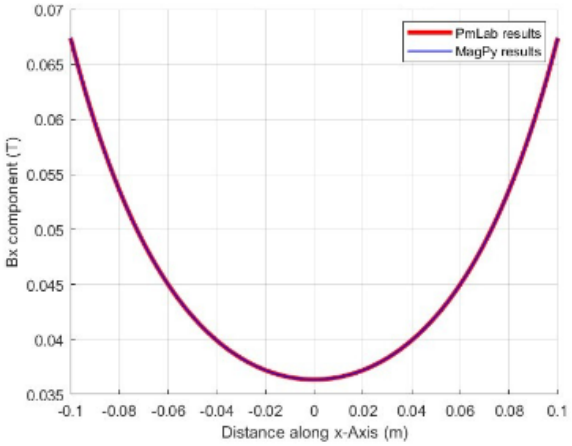

Figure 2.10: Case 2 - PmLab-MagPy comparison of the B_x component along the observation points

Figures 2.11 and 2.12 show the 2D comparison of the B_x component of the field. Despite the different graphical style, both figure exhibits the same features: at the center of the array the fields reach a minimum while moving outward the fields show elliptical patterns that increase smoothly reaching local maxima near the edges of the magnet ring.

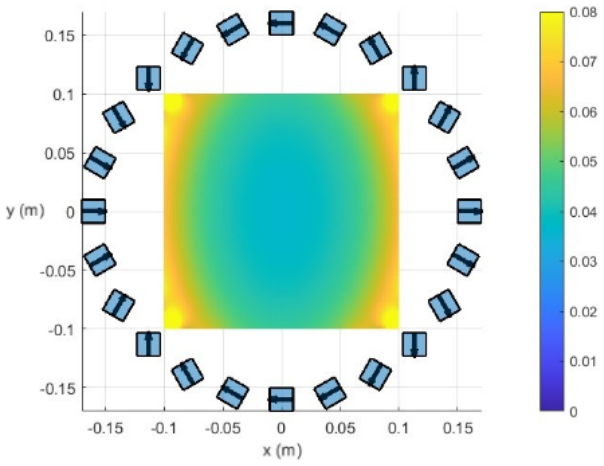

Figure 2.11: Case 2 - Two-dimensional PmLab B_x colormap

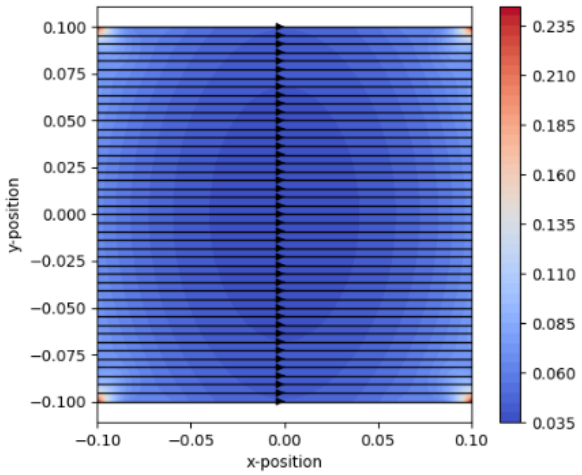

Figure 2.12: Case 2 - Two-dimensional MagPy B_x colormap

PmLab - MagTetris Comparison

Figure 2.13 shows the one-dimensional comparison of the B_x component between PmLab and MagTetris, evaluated along the same line as in the previous test. The two curves are virtually indistinguishable.

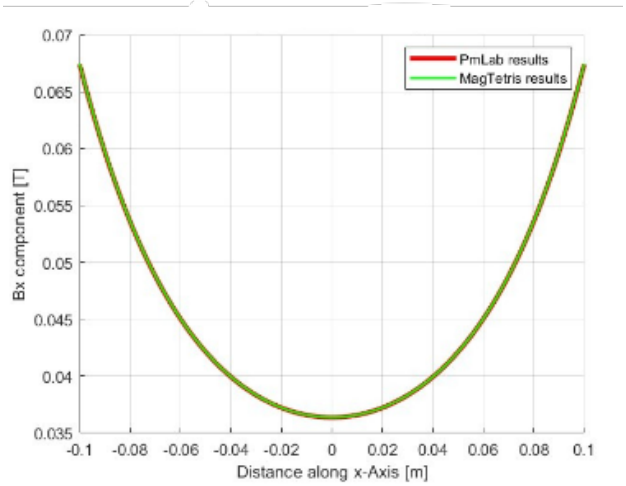

Figure 2.13: Case 2 - PmLab-MagTetris comparison of the B_x component along the observation points

In Figure 2.14, instead, the two-dimensional B_x map is displayed. As it can be seen, the field distribution matches exactly the one made through the PmLab code, both in terms of spatial symmetry and magnitude. Any visual discrepancies between the field maps—such as differences in color scales, vector overlays, or grid resolution are merely the result of distinct plotting styles and default configurations of the codes.

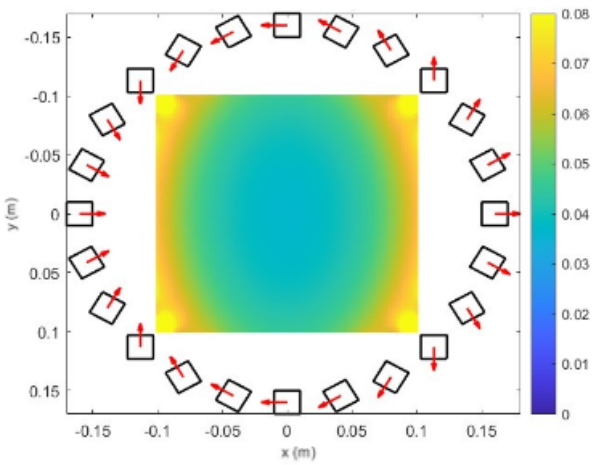

Figure 2.14: Two-dimensional *MagTetris* B_x colormap, case 2

Code	**Calculation Time [s]**
PmLab	4.91
MagPy	17.39
MagTetris	3.164

Table 2.2: Case 2 - Computation time required by each code

The Halbach ring configuration, composed of multiple brick magnets, naturally leads to higher computational load due to the accumulation of contributions from several elements. Although MagTetris is still the fastest tool, its use is limited by the absence of full 3D analysis. PmLab maintains robust performance even in this more demanding configuration, offering a notable advantage in computation time compared to MagPy, whose execution remains substantially slower. This difference in performance becomes more and more relevant as the number of magnetic elements increases. In summary, the results demonstrate that PmLab handles increasing system complexity with excellent scalability.

2.3.3 Case 3: Halbach Configuration with Tiles

The third example considers an Halbach array composed of segmented magnetic tiles, in order to follow the curvature of the ring geometry. The main parameters of the configuration are summarized in the table below:

Parameter	Value
Number of rings	1
Magnets per ring	22
Magnet dimensions	$R_{int} = 0.15$ m, $R_{ext} = 0.17$ m, $D_z = 0.1$ m
Ring radius	$R = 0.16$ m
Magnetization	Halbach pattern
Observation points for B_x	100 points from $[-0.1,\ 0,\ 0]$ to $[0.1,\ 0,\ 0]$
Evaluation grid (colormap)	$R_{\text{base}} = 0.1$ m 100 points along x from $-R_{\text{base}}$ to R_{base} 100 points along y from $-R_{\text{base}}$ to R_{base} Total number of points: 100,000

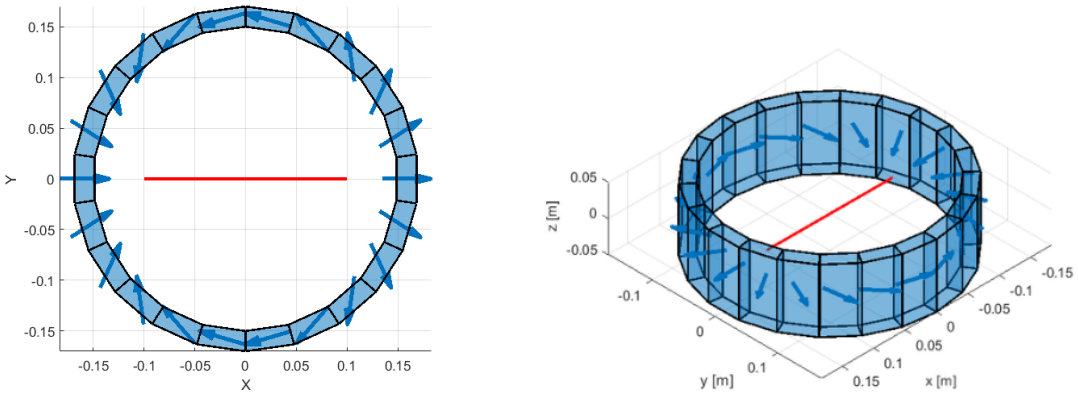

Figure 2.15: Case 3 - Halbach disposition and magnetization with tiles, in red the observation points for B_x

PmLab - MagPyLib Comparison

Before presenting the validation results for this configuration it is important to underline that the comparison for this case was made only with MagPyLib code. This is due to the fact that MagTetris code only supports parallelepipedal geometries and therefore cannot be used to model curved structures like the tiled Halbach array.

As in the previous case, the results from both codes are in excellent agreement, with nearly overlapping curves across the entire domain.
From Figures 2.17 and 2.18, again, it is possible to notice the coherent spatial distribution of the field, with a clear peak near the inner radius of the array and a symmetric decay toward the center. Unlike the previous Halbach configuration with brick-shaped magnets, the field strength is clearly higher. This is due to the greater magnetic filling factor: the curved tiles occupy a larger fraction of the ring's circumference, reducing the amount of non-magnetic gaps and improving the overall

flux concentration.

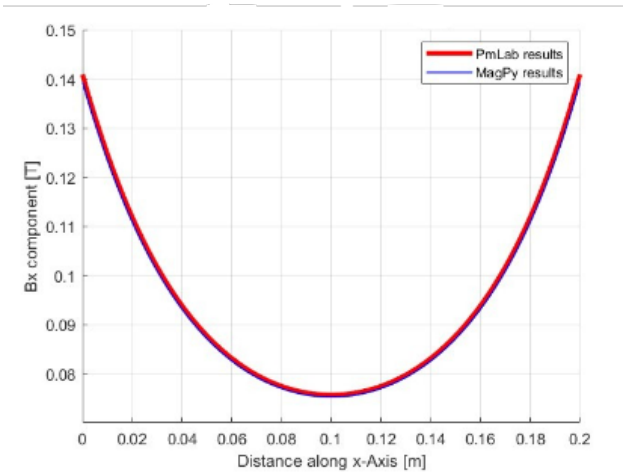

Figure 2.16: Case 3 - PmLab-MagPy comparison of the B_x component along the observation points

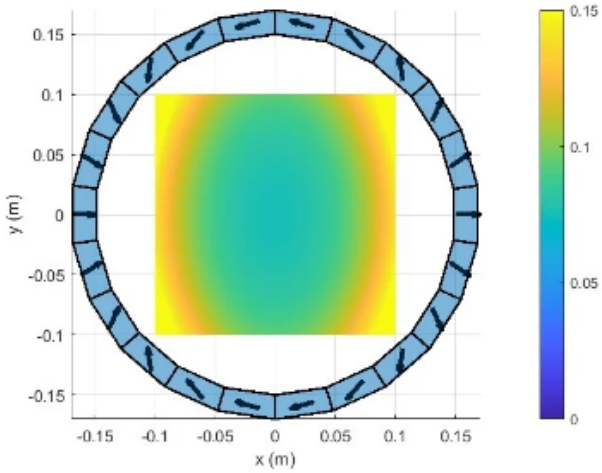

Figure 2.17: Case 3 - Two-dimensional PmLab B_x colormap

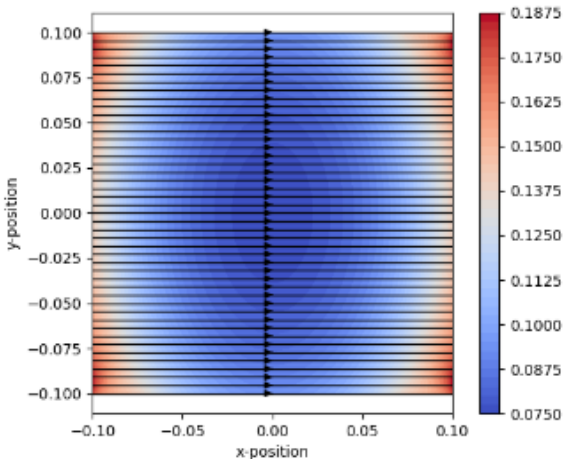

Figure 2.18: Case 3 - Two-dimensional MagPy B_x colormap

Conclusion

Through the three benchmark cases discussed in this chapter, the PmLab code has proven to be both reliable and accurate, indeed the comparisons with the well-established MagPyLib and MagTetris have shown excellent agreement in all scenarios. These results not only provide evidence of the effectiveness of the code but also of the Coulombian formulation itself. Although conceptually different from traditional current-based models, the Coulombian approach proves to be accurate, flexible and fast, even for intricate configurations.

This validation establishes a solid foundation for the numerical simulations presented in the following chapters. With this awareness, the analysis moves forward searching for the best magnetic configuration, aiming to find the most effective solution for the target application.

3 Permanent Magnet Arrays

This new chapter is dedicated to the exploration and evaluation of different permanent magnet configurations whose purpose is to generate a static magnetic field suitable for low-field magnetic resonance requirements. The goal is to identify a specific layout that provides a good compromise between field strength and homogeneity and that can be used as a baseline for the following optimization stage.
To this end a targeted literature review was carried out in order to find which magnet arrangements are commonly used in this specific field and, even if the range of available solutions is quite limited, most of the configurations can be classified in two main categories: **Halbach-type** and **Aubert-type** designs, both adopted for their ability to generate magnetic fields in compact and open geometries.
The following sections present a selection of representative configurations, inspired by existing studies and grouped according to the two main categories. To ensure coherence in the analysis, all the configurations have been simulated using the simulation tool adopted in this work. This approach eliminates potential differences between different modeling strategies that may have been used, allowing for a direct and fair comparison of results.

3.1 Permanent Magnets

Before presenting the various configurations selected for the analysis, it is important to define the material and the magnetization assumptions made in this work. Although in the literature a variety of magnets, with different properties, are employed, the comparisons in this chapter have been performed under uniform conditions.
In particular, all the configuration are simulated under the assumption of uniform remanent magnetization B_r for all magnets.
But what type of permanent magnet should be used?
Among the various permanent magnet materials available, only Neodymium-Iron-Boron (NdFeB) and Samarium-Cobalt (SmCo) are considered suitable for the pur-

poses of this study. This is because the configurations analyzed do not rely ferromagnetic yokes to channel the magnetic field. Instead, they rely only on permanent magnets in order to generate a sufficiently strong magnetic field in air, where, of course, the field of view (FOV) is located.
NdFeB magnets are commonly preferred due to their high remanence and energy density, though they are more sensitive to temperature variations. SmCo magnets, while slightly less powerful, provide greater thermal stability and resistance to demagnetization.

Magnet	$\boldsymbol{B_r}$ (T)	$\boldsymbol{H_c}$ (kA/m)	$\boldsymbol{T_c}$ (°C)
SmCo (sintered)[2]	0.8–1.1	600–2000	720
NdFeB (sintered)	1.0–1.4	750–2000	70–200

Table 3.1: Magnetic properties of NdFeB and SmCo

For all the simulations a fixed $B_r = 1.4.\ T$ has been assumed. This value corresponds to a high-performance permanent magnet and is intended to represent a best-case scenario from a purely magnetic perspective. No considerations such as manufacturing cost, material availability, and thermal stability have been taken into account.
The selected value of $\boldsymbol{B_r}$ is representative of commercial NdFeB magnets in the N48 to N52 grade range, which are among the most powerful rare-earth permanent magnets available. Table 3.2 summarizes the typical remanent flux densities for common NdFeB grades.

NdFeB Grade	$\boldsymbol{B_r}$ (T)
N35	1.17–1.21
N42	1.29–1.32
N48	1.38–1.40
N52	1.42–1.44

Table 3.2: Typical remanent flux density values for commercial NdFeB magnet grades.

[2]The term sintered refers to a manufacturing process in which magnetic powders are compacted and heated below their melting point to produce dense, high-performance permanent magnets. Sintered magnets generally offer superior magnetic properties compared to bonded types.

3.2 Halbach Arrays

A Halbach array, as briefly mentioned in Section 2.3.2, is a particular configuration of permanents magnets placed in such a way to concentrate the field on one side of the assembly while minimizing it on the other side. This asymmetric field distribution is achieved through a specific orientation of the magnetization vectors within the array. Halbach configurations can be implemented in different geometries: one-dimensional (linear array), two-dimensional (cylindrical structure), or three-dimensional (spherical arrangement). For magnetic resonance imaging, the cylindrical (2D) geometry is the one of practical interest due to its ability to generate a strong and relatively homogeneous magnetic field within a central bore.
An ideal cylindrical Halbach array can be described mathematically by a continuous rotation of the magnetization vector around the azimuthal[3] coordinate of the cylinder. This rotation is characterized by the relation:

$$\phi = (k + 1)\theta \tag{3.1}$$

where:

- θ is the azimuthal position of the magnet element (i.e., its angular position around the ring);
- ϕ is the angle of the magnetization vector within that element;
- k is the multipolarity index that defines the type of Halbach field being generated.

This expression indicates that the direction of magnetization rotates more rapidly than the physical placement of the magnets. The value of k directly controls how the magnetization rotates along the ring, and thus how the field is shaped, as presented in Figure 3.1:

- for $k = +1$, the magnetization rotates at twice the rate of the angular position, generating a dipolar field concentrated inside the cylinder;
- for $k = 2$, the configuration corresponds to a quadrupolar field, with a more spatially varying distribution within the structure.

[3]The azimuthal coordinate θ in cylindrical coordinates defines the angular position around the central axis of the ring, measured from a fixed reference direction in the plane perpendicular to the axis.

As k increases, the resulting field patterns become progressively more complex, while negative values of k lead to field configurations that are concentrated outside the cylindrical assembly rather than in it.

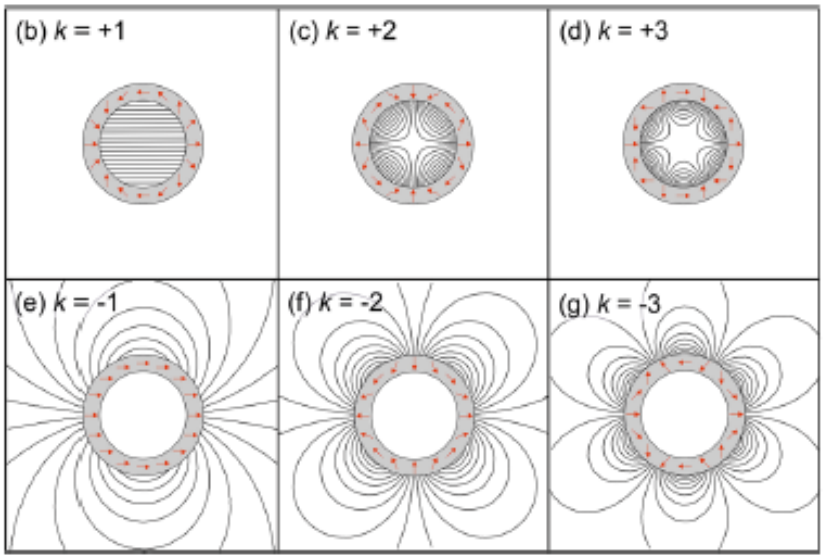

Figure 3.1: Conceptual construction of cylindrical ideal Halbach multipoles (adapted from [3])

Among all possible Halbach configurations, only the case with $k = 1$ is relevant for MRI. This specific arrangement produces a dipolar field confined inside the cylinder, with the magnetic field oriented perpendicularly to the cylinder's axis (Figure 3.2). Such geometry is well-suited for generating the static magnetic field required in MRI systems, especially in open-access and low-field designs.

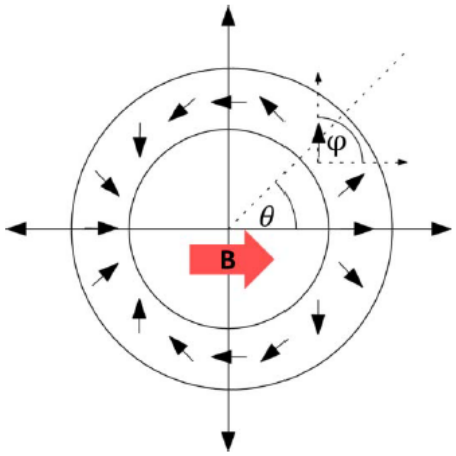

Figure 3.2: Magnet polarization for a dipolar Halbach cylinder, adapted from [4]

The Halbach configuration described so far represents an idealized case, valid under the assumption of a continuously varying magnetization along the magnetic

material and infinitely long cylindrical geometry. However, such ideal conditions are not practically achievable, the fabrication of magnets with constantly rotating magnetization is technologically unfeasible, at least at the scale required for MRI systems[4].

The closest realizable alternative consists in discretizing the magnetization profile of the cylinder by dividing it into individual segments - typically cuboid, cylindrical, or trapezoidal bars - each uniformly magnetized in a fixed direction. While this segmented approach is significantly more economical and manufacturable, it inherently compromises both the strength and uniformity of the magnetic field.

Moreover, when the array is constructed using bars additional field distortions arise near the axial ends of the structure. These so-called fringe effects can significantly degrade the field quality in the central region.

Having discussed the theoretical principles and practical limitations of Halbach arrays, we now turn our attention to the specific configurations considered in this work. These geometries, inspired by existing designs in the literature, represent actual implementations of segmented Halbach structures.

3.2.1 Configuration 1

The first configuration analyzed in this study is derived from the portable Halbach-based MRI system described in [8]. This design, referred to as **Mandhala** (Magnet Arrangements for Novel Discrete Halbach Layout), represents a compact and practical implementation of a discretized cylindrical Halbach array specifically optimized for low-field imaging.

The magnet structure consists of a $36\,cm$ diameter cylinder composed of twenty square-section rungs, each containing stacked NdFeB bar magnets arranged to approximate a dipolar Halbach configuration. The magnets dimensions are $1x1x14''$ ($1'' = 2.54\,cm$ so the dimensions are $2.54 \times 2.54 \times 35.5\,cm$), magnetized along the $1''$ thickness. To reduce axial field decay due to the limited length of the array, two additional Halbach end-rings were added, each composed of twenty cubic NdFeB magnets with 1" sides (Figure 3.3).

[4]Small bonded continuous magnetization Halbach rings can be produced using injection molding, this is done for small size brushlss motors [4]

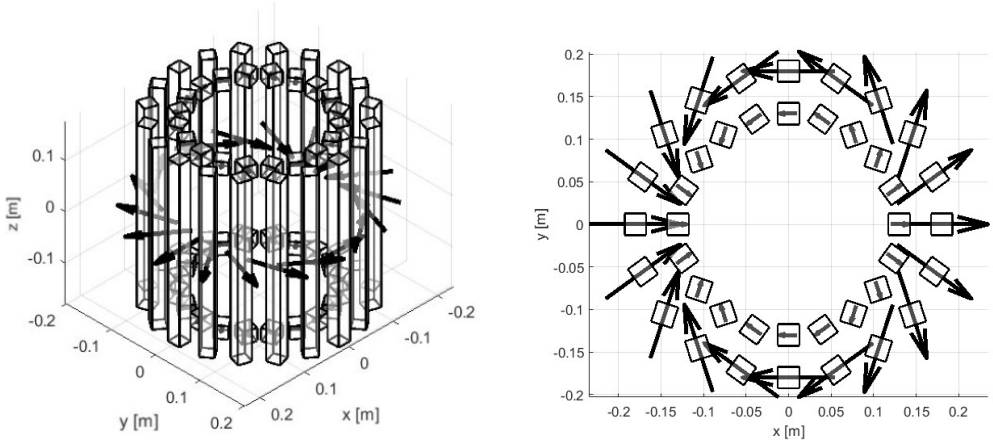

Figure 3.3: Configuration 1 - Schematics of the array

As in [8], the field was evaluated within a spherical field of view (FOV) of $0.1\,m$ radius, centered at the geometric center of the array. This choice ensures consistency with the original analysis and allows for a direct comparison of field behavior.
The analysis focuses exclusively on the B_x component of the magnetic field, as the angular coordinate $\theta = 0$ is aligned with the x-axis (as in Fig.3.2). In this orientation, B_x is the dominant component; B_y cancels out by symmetry, and the B_z component is practically negligible.

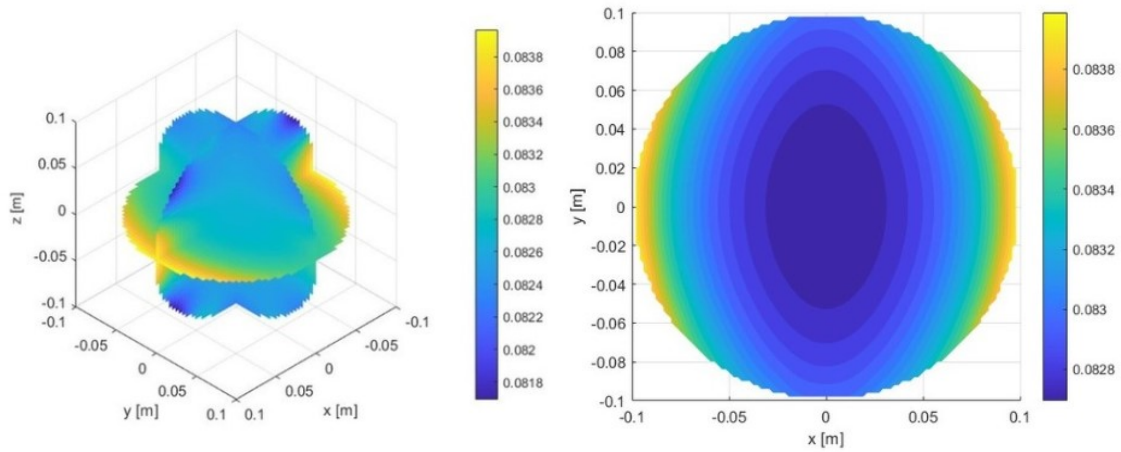

(a) (b)

Figure 3.4: Configuration 1 – Evaluation of the B_x component in the FOV (a) and on the xy-plane (b)

Center	[0, 0, 0] m
Obs. points	30976
$B_{x,\max}$	0.0840 T
$B_{x,\min}$	0.0817 T
$B_{x,\text{mean}}$	0.0827 T
Homogeneity	27507 ppm

Center	[0, 0, 0] m
Obs. points	7668
$B_{x,\max}$	0.0840 T
$B_{x,\min}$	0.0817 T
$B_{x,\text{mean}}$	0.0830 T
Homogeneity	15568 ppm

Table 3.3: Configuration 1 - Results of PmLab computation on the spherical FOV (left) and on the xy-plane (right)

The results of the magnetic field simulations are summarized in Tables 3.3. For both domains, the maximum, minimum, and average values of the B_x component are reported, along with the number of observation points.
In both cases, the field homogeneity is quantified using the standard metric expressed in parts per million (ppm), calculated according to the following formula:

$$\text{Homogeneity}_{\text{ppm}} = \frac{B_{x,\max} - B_{x,\min}}{B_{x,\text{mean}}} \times 10^6 \tag{3.2}$$

This metric provides a relative measure of the field variation within the region of interest. As expected, the homogeneity value is lower when computed over the central plane, where the field is more stable, than over the entire spherical FOV, which includes fringe regions near the ends of the magnet array.
While this method is widely adopted for quantifying field homogeneity due to its simplicity, it provides only a coarse estimate as it considers only the maximum deviation from the mean field, offering no view of the spatial distribution of the field within the FOV. As a result, two different field distributions may produce the same homogeneity value under this definition.
For a more complete characterization, a more advanced method based on **spherical harmonics** can be applied [9]. This technique is based on the fact that, in source-free regions, the magnetic field can be expressed as a sum of spherical harmonic components. Each component represents a distinct spatial pattern, and the relative amplitudes of these terms provide quantitative insight into how the field deviates from an ideal, perfectly uniform distribution.
Harmonic analysis enables both the quantification and characterization of field inhomogeneities, making it especially useful in MRI, where certain types of distortion can be directly associated with specific spherical harmonic terms.
Each spherical harmonic component is identified by two indices: degree n and order m. The degree n determines the overall spatial complexity of the field component: as n increases, the magnetic field exhibits more spatially intricate patterns, with a

larger number of regions where the field direction or magnitude changes. The order m, which ranges from -n to +n, controls the azimuthal variation of the harmonic component, i.e., how the field changes when rotating around the different axis.

The table 3.4 reports the relative contribution of each harmonic degree for configuration 1 mentioned above.

Degree n	**Relative Contribution**	**Dominant Orders m**
0	9.997×10^{-1}	(uniform field)
1	5.039×10^{-11}	negligible
2	3.296×10^{-7}	$m = -2$ (XY), $m = 1$ (ZX)
3	3.371×10^{-10}	negligible

Table 3.4: Configuration 1 - Relative contribution of the spherical harmonic degrees

As expected, the degree $n = 0$ component, corresponding to a perfectly uniform field, dominates the spectrum. Its normalized weight is close to 1, indicating that most of the energy is concentrated in the uniform mode. All higher-order terms contribute by several orders of magnitude less. The components $n = 1$ and $n = 3$ are essentially negligible, while a small but noticeable contribution is observed for $n = 2$, which corresponds to quadrupolar field distortions. Among the five possible harmonic components of degree $n = 2$, the analysis reveals that the most influential terms are those of order $m = -2$ and $m = 1$. These correspond respectively to field distortions in the xy- and zx-planes, as illustrated in Figures 3.5 and 3.6.

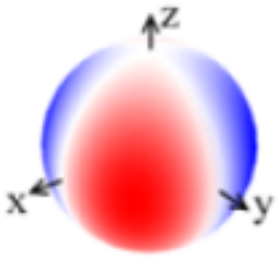

Figure 3.5: $n = 2, m = -2$

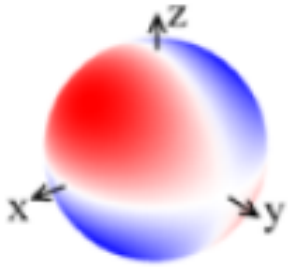

Figure 3.6: $n = 2, m = 1$

3.2.2 Configuration 2

The second configuration analyzed in this work is based on the sparse Halbach magnet array described in [4] and [10]. The array consists of a cylindrical frame with 2 layers of 24 rectangular rods (internal and external diameters are respectively 41 cm and 50 cm), each one composed of 18 NdFeB cubes of 1″ side (Figure 3.7

). A key feature of this design is its asymmetry along the axial direction: to accommodate anatomical constraints, such as shoulder width, the array adopts an asymmetric design: only 7 rows are positioned below the isocenter, while 11 are placed above it. For this reason, in order to compensate for the magnetic field drop-off near the lower end of the imaging volume, an additional ring of cubic magnets ($36\,cm$ in diameter) is placed at the bottom of the array, while no corresponding top-end ring is present since the longer structure provides a smoother field decay

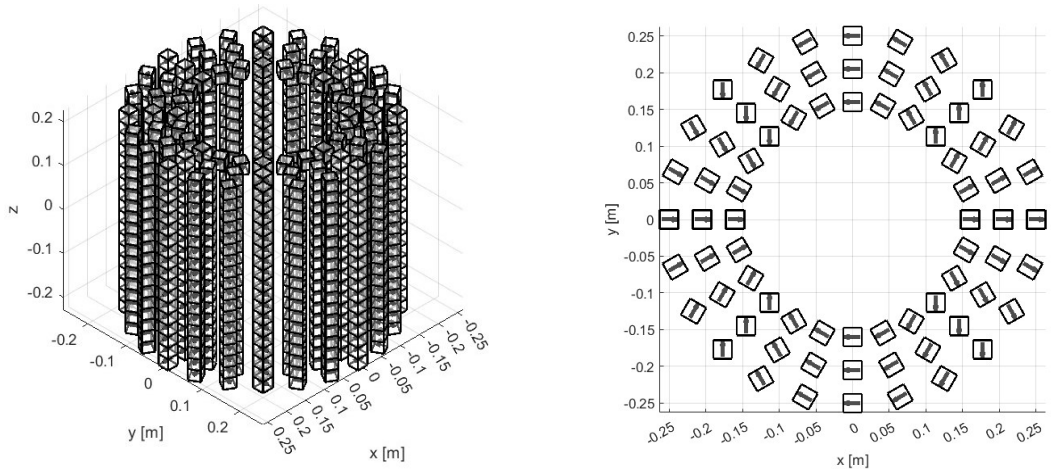

Figure 3.7: Configuration 2 - Schematics of the array

As in the previous case, the magnetic field is evaluated within a spherical FOV of radius $0.1\,m$. However, due to the asymmetry of the structure, the FOV is slightly offset in the axial direction and centered in $[0, 0, 0.05]\,m$ in the current reference frame (Figure 3.8).

In both [4] and [10] the structure was discretized into a large number of $1''$ cubes in order to allow the genetic optimization process. Indeed this segmentation enabled each cube to be independently defined as non-magnetic or composed of either N42 or N52 grade material. In the reproduction carried out in this work, however, all magnetized elements were assigned the same B_r, thereby removing any material variability. Even though the cube-based design was kept for consistency, the system behaves like a set of identical magnetized bars.

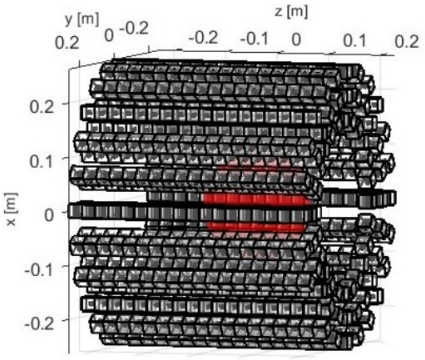

Figure 3.8: Configuration 2 - Array side view (FOV in red)

The results of the simulation performed with *PmLab*, always relatively to the B_x component of the magnetic field, are reported in Figure 3.14

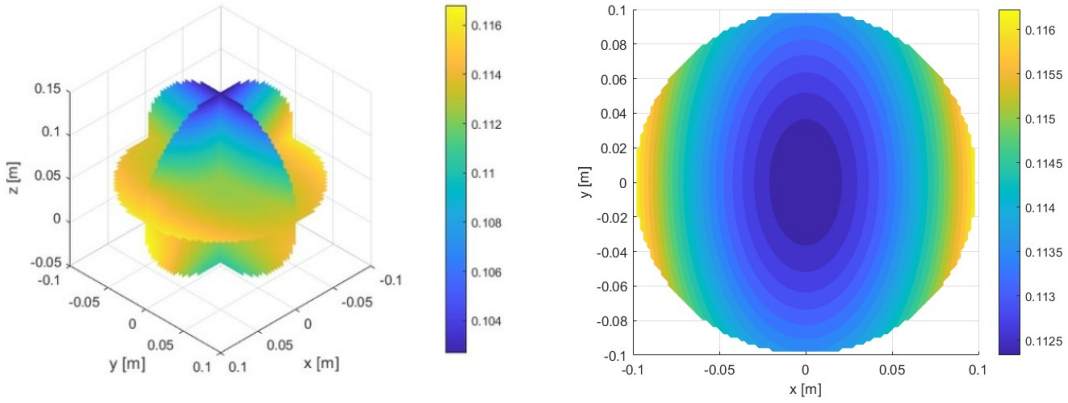

(a) (b)

Figure 3.9: Configuration 2 – Evaluation of the B_x component in the FOV (a) and on the xy-plane (b)

Center	[0, 0, 0.05] m
Obs. points	30976
$B_{x,\text{max}}$	0.1168 T
$B_{x,\text{min}}$	0.1027 T
$B_{x,\text{mean}}$	0.1123 T
Homogeneity	125582 ppm

Center	[0, 0, 0.05] m
Obs. points	7668
$B_{x,\text{max}}$	0.0162 T
$B_{x,\text{min}}$	0.1123 T
$B_{x,\text{mean}}$	0.1137 T
Homogeneity	34242 ppm

Table 3.5: Configuration 2 - Results of PmLab computation on the spherical FOV (left) and on the xy-plane (right)

Compared to sec.3.2.1, relevant differences are noticeable in both strenght and homogeneity. In this case the field intensity is significantly higher (0.1123 T against 0.0827 T in mean values). This is mostly due to the greater amount of magnetic

used in the second structure, which presents a more densely packed arrangement of bars. Despite the higher peak field in Configuration 2, it exhibits a significantly lower homogeneity, particularly in the spherical volume. This illustrates a key, and sometimes counterintuitive, concept: adding more magnetic material does not automatically result in better field uniformity. Therefore designing an effective permanent magnet system isn't just a matter of adding more material, it also requires careful considerations of geometry, magnetization direction and symmetry, especially in open structures where magnetic flux cannot be easily guided or confined.

Degree n	**Relative Contribution**	**Dominant Orders m**
0	9.989×10^{-1}	(uniform field)
1	6.370×10^{-8}	negligible
2	7.616×10^{-4}	$m = -2$ (XY), $m = 1$ (ZX)
3	3.239×10^{-4}	negligible

Table 3.6: Configuration 2 - Relative contribution of the spherical harmonic degrees

The harmonic decomposition of the magnetic field (Table 3.6) shows that the order $n = 0$ is dominant, reflecting the presence of a strong and nearly uniform field across the FOV. Similarly to the first configuration, the relative magnitude of the second order is the strongest one and, among the degree-two components, the most prominent distortions are those associated with the orders $m = -2$ and $m = 1$, already presented in Figures 3.5 and 3.6. However a non-negligible contribution also comes from the third-order harmonics, among which the most relevant presence is the one associated with the order $m = 0$, as in Figure 3.10.

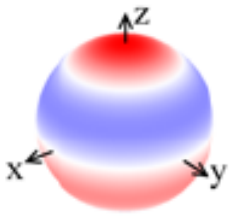

Figure 3.10: $n = 3$, $m = 0$

This behavior directly reflects the asymmetry of the physical structure, in particular the presence of the compensating lower ring and the absence of a mirrored upper ring. As a result, the magnetic field tends to be stronger in the lower portion of the FOV, as also confirmed by the field distribution calculated within the spherical region.

3.2.3 Configuration 3

The third configuration explored in this study is inspired by the elliptical Halbach structure proposed in [11]. Unlike traditional cylindrical Halbach arrays, this design uses an elliptical bore, which could offer several practical benefits. These include enhanced patient accessibility, easier integration of RF coils, and improved anatomical conformity, particularly for applications focused on the head or limbs.
Since the original study does not provide detailed geometric parameters, the implementation adopted here preserves the same magnet arrangement as in the previous cylindrical configurations, modifying only the inner profile to follow an elliptical shape. The purpose of this analysis is to evaluate whether introducing an elliptical deformation leads to improvements in field homogeneity. A decrease in field strength is already expected since the magnets (especially those near the major semiaxis) are displaced further from the FOV. However, one could hypothesize that such a deformation may attenuate the field in directions where it is stronger (at the boundary of the FOV along the x-axis for example) contributing to a more homogeneous distribution.
The main geometric parameters of the elliptical Halbach structure and its representation, derived from Condiguration 1 and adapted to an elliptical profile, are reported in Table 3.7 and Figure 3.11.

Parameter	**Value**
Number of vertical bars	20
Bars dimensions	$0.025 \times 0.025 \times 0.355\,m$
Bars semi-axes (X, Y)	$0.22\,m$, $0.18\,m$
Number of shielding rings	2
Cubes per ring	20
Cubes dimensions	$0.025 \times 0.025 \times 0.025\,m$
Shielding ring semi-axes (X, Y)	$0.18\,m$, $0.14\,m$
FOV center	$[0,\ 0,\ 0]\,m$
FOV radius	$0.1\,m$

Table 3.7: Main geometric parameters, Configuration 3*b*

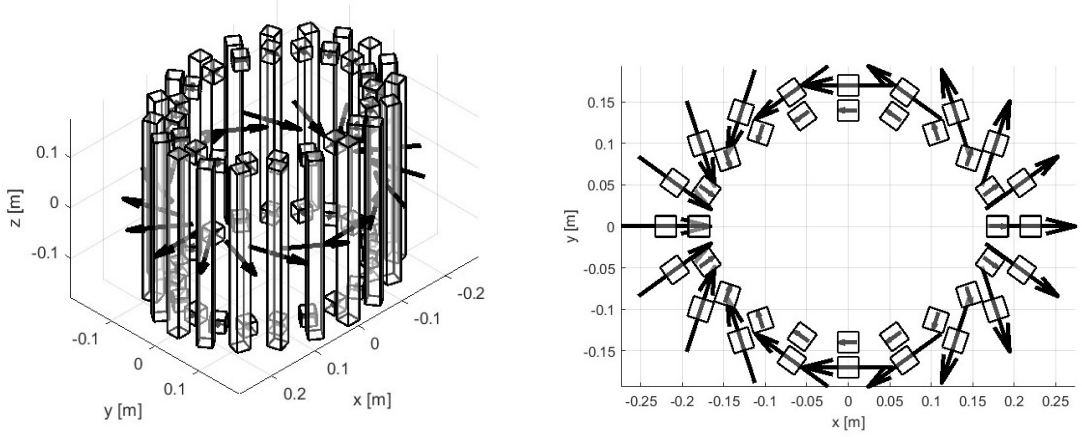

Figure 3.11: Schematics of the array, Configuration 3*a*

In Figure 3.12 it is presented the second elliptical configuration, this time derived from Configuration 2. The components of the structure and their disposition remain the same, the only difference is in the geometry (Table 3.8).

Parameter	**Value**
Number of Layers	2
Number of vertical rungs	24
Cubes per rungs	18
Cubes dimensions	$0.025 \times 0.025 \times 0.025\,m$
Rungs semi-axes (X, Y)	$0.32\,m$, $0.25\,m$ (ext. layer)
	$0.27\,m$, $0.20\,m$ (int. layer)
Number of shielding rings	1
Cubes per ring	24
Cubes dimensions	$0.025 \times 0.025 \times 0.025\,m$
Shielding ring semi-axes (X, Y)	$0.22\,m$, $0.15\,m$
FOV center	$[0,\ 0,\ 0.05]\,m$
FOV radius	$0.1\,m$

Table 3.8: Main geometric parameters, Configuration 3*b*

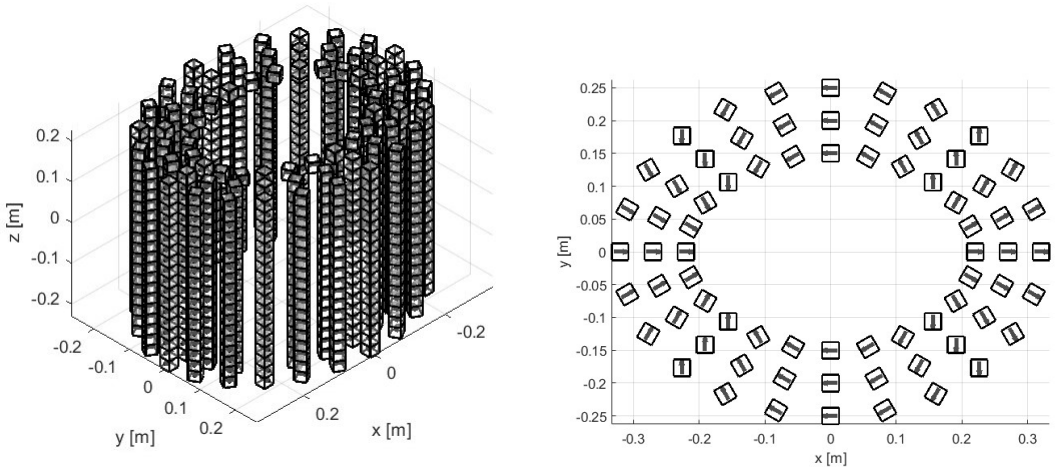

Figure 3.12: Schematics of the array, Configuration 3*b*

The magnetic field distributions produced by both elliptical configurations (Cases 3*a* and 3*b*) were evaluated using the same methodology adopted for the previous structures. The simulation outcomes for both cases are presented in the following figures and tables.

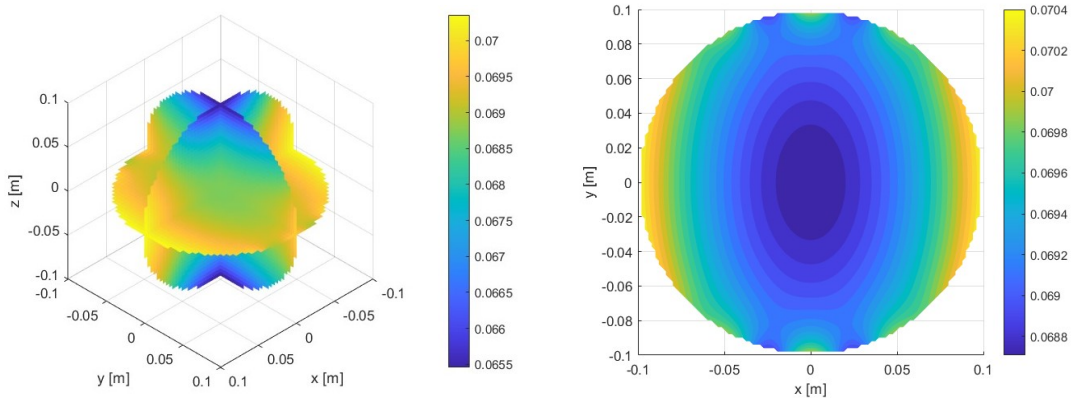

(a) (b)

Figure 3.13: Configuration 3*a* – Evaluation of the B_x component in the FOV (a) and on the *xy*-plane (b)

Center	[0, 0, 0] *m*
Obs. points	30976
$B_{x,\text{max}}$	0.0704 T
$B_{x,\text{min}}$	0.0655 T
$B_{x,\text{mean}}$	0.0687 T
Homogeneity	71327 ppm

Center	[0, 0, 0] *m*
Obs. points	7668
$B_{x,\text{max}}$	0.0704 T
$B_{x,\text{min}}$	0.0687 T
$B_{x,\text{mean}}$	0.0693 T
Homogeneity	24396 ppm

Table 3.9: Configuration 3*a* - Results of PmLab computation on the spherical FOV (left) and on the *xy*-plane (right)

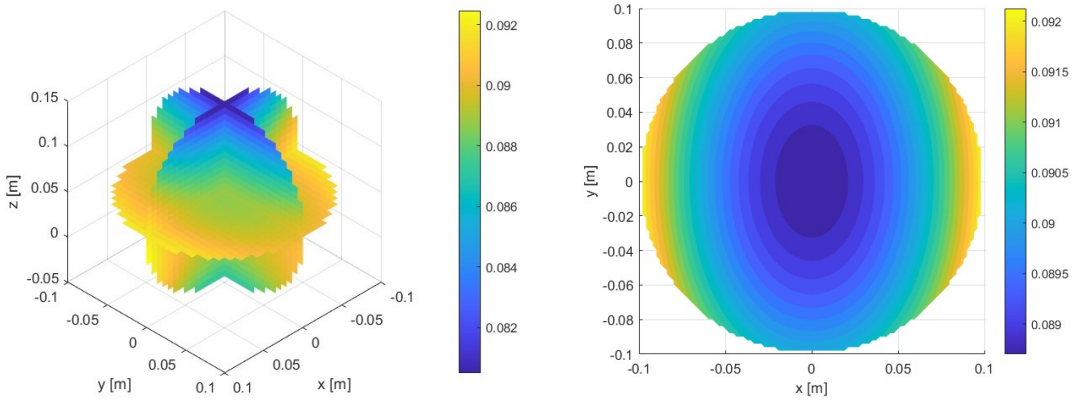

(a) (b)

Figure 3.14: Configuration 3*b* – Evaluation of the B_x component in the FOV (a) and on the *xy*-plane (b)

Center	[0, 0, 0.05] *m*
Obs. points	30976
$B_{x,\text{max}}$	0.0925 T
$B_{x,\text{min}}$	0.0805 T
$B_{x,\text{mean}}$	0.0827 T
Homogeneity	134957 ppm

Center	[0, 0, 0.05] *m*
Obs. points	7668
$B_{x,\text{max}}$	0.0921 T
$B_{x,\text{min}}$	0.0887 T
$B_{x,\text{mean}}$	0.0900 T
Homogeneity	38088 ppm

Table 3.10: Configuration 3*b* - Results of PmLab computation on the spherical FOV (left) and on the *xy*-plane (right)

The results obtained for the elliptical configurations show that deforming the bore profile does not lead to improvements in field uniformity. In fact, both elliptical variants exhibit lower field intensity and worse or comparable homogeneity relative to their cylindrical counterparts. Although the elliptical deformation was introduced to reduce field excess along the *x*-direction by increasing the distance between the magnets and the boundaries of the FOV, the results show that this adjustment alone does not effectively eliminate the inherent inhomogeneities. By altering the arrangement of the magnets, it is not possible to selectively modify the magnetic field in only one portion of the FOV without inevitably affecting the rest of the region. Any attempt to reduce field intensity at specific locations, such as the outer edges of the FOV, inevitably leads to a reduction of the field in its central area and the homogeneity results either preserved or degraded. This highlights a fundamental limitation: trying to correct the field in a specific area by changing the geometry of the magnet arrangement tends to affect the entire field distribution.

3.3 Aubert Rings

An effective and technically valid alternative for generating the $\boldsymbol{B}_0$ field in portable MRI systems is provided by a specific class of magnetic structures known as Aubert ring pairs or I/O ring pairs.
Unlike the more common Halbach-based configurations, which produce a transverse magnetic field with respect to the system axis, Aubert rings are designed to generate a longitudinal field, aligned with the axis of the scanner (typically the z-axis). Thanks to this feature, the Aubert configuration is highly advantageous for system integration, as it allows the direct reuse of conventional RF coil design from conventional superconducting MRI. The radiofrequency coils, as briefly anticipated in Section 1.1, are essential components in MRI systems. They are both used to transmit the excitation pulses that flip the nuclear spins and to receive the signal emitted by the precessing magnetization and, for better performances these coils have to be aligned with $\boldsymbol{B}_0$, as their sensitivity is maximized when the oscillating magnetic field they produce is perpendicular to $\boldsymbol{B}_0$[5].
An Aubert ring pair consists of two coaxial radial magnet rings, one magnetized inward and the other outward. The opposing radial magnetizations generate magnetic field lines that converge and reinforce each other along the longitudinal axis, resulting in a predominantly axial magnetic field within the central region of the structure. This effect is achieved without the need for ferromagnetic yokes, as illustrated in Figure 3.15.

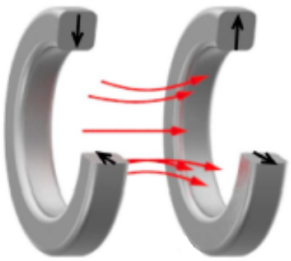

Figure 3.15: Sectional view of an Aubert ring pair, adopted from [5]

Unlike Halbach configurations, specifically designed to confine the field on one side while suppressing it on the opposite, the I/O ring pair geometry lacks of a field-shielding mechanism. For this reason, although the field is directed towards

[5]This condition holds in conventional imaging systems that do not use Spatial Encoding Magnetic fields (SEM). When SEMs are employed, the encoding is embedded directly in the spatial structure of the static field, and the orientation constraints on the RF coils can be relaxed, as their operation no longer strictly depends on the alignment with a uniform $\boldsymbol{B}_0$ field.

the bore center, it is not strongly confined and non-negligible fring field may be present beyond the system's boundaries.
Similarly to Halbach arrays, the ideal Aubert ring, with a smoothly varying radial magnetization pointing inward or outward around the annular piece, is very difficult to implement in practice. To overcome this limitation the structure is usually discretized into magnetic segments (like the usual bricks and tiles), each with a uniform magnetization that approximates the radial pattern.
The following sections present two discretized implementations of the Aubert ring concept. These configurations provide realistic and practically achievable models for assessing the magnetic performance of Aubert-based assemblies in low-field MRI applications.

3.3.1 Configuration 4

The fourth configuration is derived from the optimized design proposed in [12]. In the original study, the structure was subject to a multi-stage optimization process, including additional magnetic components specifically introduced to generate gradient fields. In the present analysis, however, only the base structure, consisting of radially magnetized brick magnets arranged in an IO ring pair, is considered, in order to isolate and evaluate the fundamental properties of the Aubert configuration. Each ring of the presented configuration is composed of 22 identical NdFeB bricks, arranged uniformly along a circular path. The blocks are, of course, magnetized radially, with opposing directions in the two rings.
Each magnet block has dimensions of $120 \times 30 \times 20\,mm$, and, as in the previous configurations, is assumed to have a remanent magnetization $B_r = 1.4\,T$. The magnets are placed with a uniform angular spacing of $360°/22 \approx 16.36°$ around each ring, forming a nearly continuous magnetic surface. Each of the two magnetic assemblies composing the Aubert ring pair consists of five stacked rings of brick-shaped magnets, aligned along the axial direction. As it can be seen in the sectional view in Figure 3.16, within each assembly, the rings present a slightly offset in the radial direction relative to one another. Starting from the outermost ring, the radii of the stacked layers follow the sequence: $[0.188, 0.180, 0.172, 0.172, 0.180]\,m$, forming a symmetric radial profile designed to enhance field homogeneity[6]. As a result of this arrangement, the effective bore diameter, defined as the minimum clear aperture through the magnet structure, is approximately $340\,mm$, while the distance between the two stacks has been set to $280\,mm$.

[6]The radial misalignment between stacked rings was not arbitrary, but resulted from a genetic optimization process described in the original study.

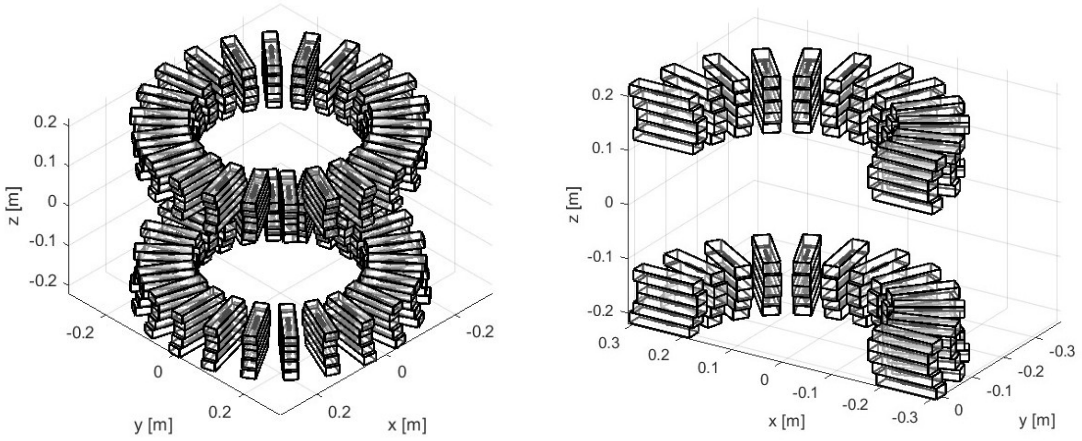

Figure 3.16: Configuration 4 – Schematics and sectional view of the array

Unlike the previous Halbach configurations that employed spherical regions of interest, the current and the following I/O ring based design adopt a cylindrical FOV, as defined in the reference studies In both cases, the FOV is modeled as a cylinder with a diameter of 20 cm and a height of 5 cm, centered at the origin of the coordinate system.

The dominant contribution to the field, the one that will be explored, is the B_z component, as shown in 3.17.

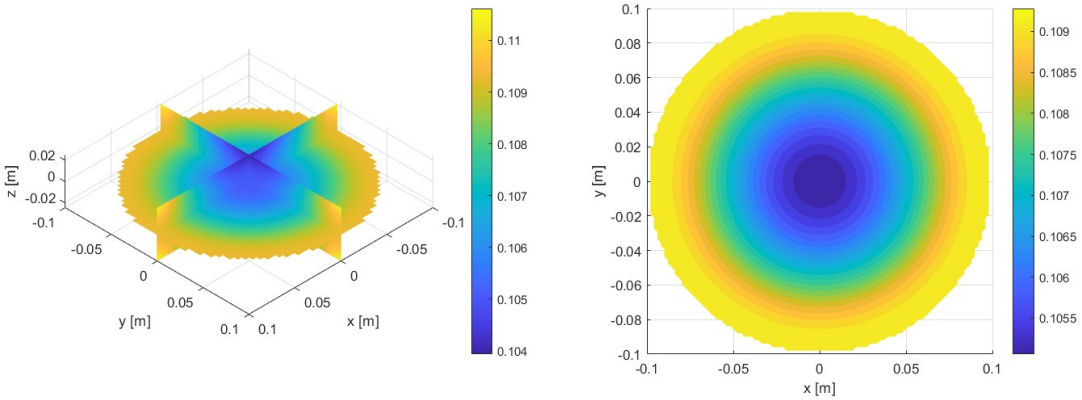

(a) (b)

Figure 3.17: Configuration 4 – Evaluation of the B_z component in the FOV (a) and on the xy-plane (b)

Center	[0, 0, 0] m
Obs. points	19940
$B_{z,\max}$	0.1106 T
$B_{z,\min}$	0.1039 T
$B_{z,\text{mean}}$	0.1080 T
Homogeneity	61874 ppm

Center	[0, 0, 0] m
Obs. points	7668
$B_{z,\max}$	0.1093 T
$B_{z,\min}$	0.1051 T
$B_{z,\text{mean}}$	0.1079 T
Homogeneity	39090 ppm

Table 3.11: Configuration 4 - Results of PmLab computation on the cylindrical FOV (left) and on the xy-plane (right)

Configuration 4 produces magnetic field strengths comparable to those obtained with Configuration 2 (Halbach array with vertical bars). This result is particularly promising given the low magnetic filling factor of the current setup. This indicates that this field strength is achieved without fully exploiting the available volume. As a result, the configuration demonstrates strong potential for generating even higher magnetic fields if the magnet density is increased.
The dominant B_z component exhibits a concentric and radially symmetric distribution that gradually tends to decrease from the periphery to the center of the FOV. Another noticeable consideration is that the maximum magnetic field is not located on the xy-plane at $z = 0$, but in adjacent axial planes, closer to the magnet rings. Both these aspects are a direct consequence of the disposition of the magnets and the distance between them and the central xy-plane.

Degree n	**Relative Contribution**	**Dominant Orders m**
0	9.904×10^{-1}	(uniform field)
1	1.107×10^{-9}	negligible
2	9.531×10^{-3}	$m = -2$ (XY), $m = 1$ (ZX)
3	2.275×10^{-8}	negligible

Table 3.12: Configuration 4 - Relative contribution of the spherical harmonic degrees

The results of the spherical harmonic decomposition (Table 3.12) show an harmonic spectrum close to the one obtained for Halbach structures. Second-order harmonics are, indeed, the main responsible for the inhomogeneities in the FOV.

3.3.2 Configuration 5

The last configuration under investigation is derived from [5], where it is realized using magnetized tiles, arranged in such a way to reproduce the radial magnetization pattern.
Unlike the original study, which involves extensive optimization procedures, the

version here analyzed adopts a simplified initial geometry, intended to provide a clear reference for comparative analysis.

The structure analyzed in Configuration 5 is composed of two symmetric assemblies, each consisting of nine concentric magnetic rings. All rings share a common outer diameter of $0.25\,m$ and a uniform axial thickness of $1\,cm$. Within each assembly, the inner diameter of the rings increases progressively from $0.15\,m$ to $0.23\,m$, with a step of $1\,cm$ between successive layers. This results in a set of nested annular tiles with decreasing radial width toward the center. The axial distance between the inner rings of the two assemblies is $0.20\,m$, defining the clear gap across which the magnetic field is concentrated, as illustrated in Figure 3.18.

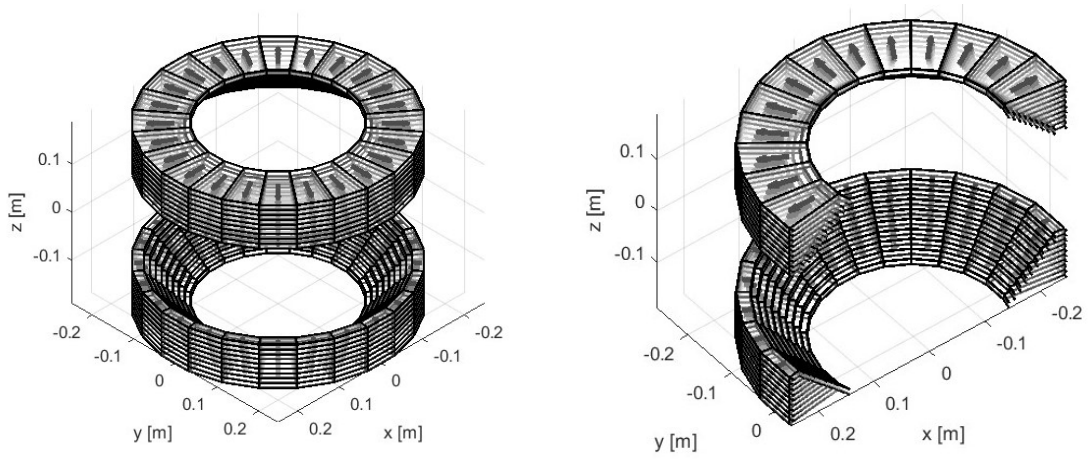

Figure 3.18: Configuration 5 – Schematics and sectional view of the array

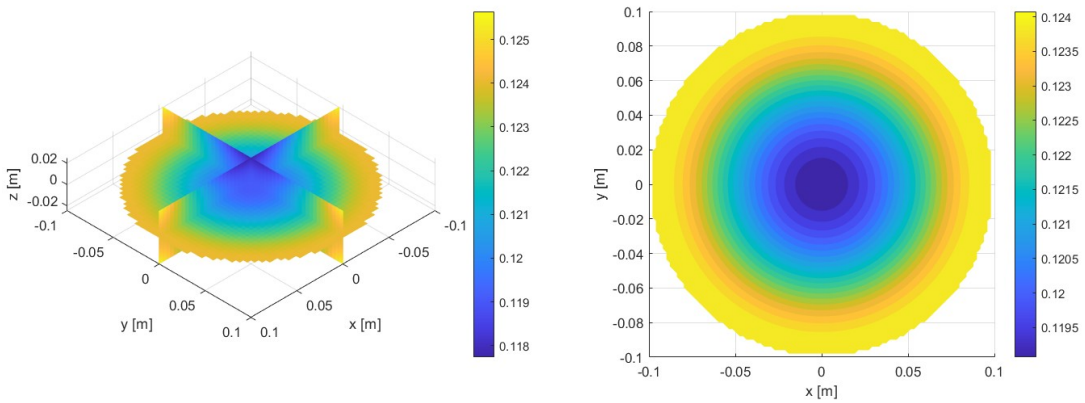

(a) (b)

Figure 3.19: Configuration 5 – Evaluation of the B_z component in the FOV (a) and on the xy-plane (b)

Center	[0, 0, 0] m
Obs. points	19940
$B_{z,\max}$	0.1256 T
$B_{z,\min}$	0.1177 T
$B_{z,\text{mean}}$	0.1225 T
Homogeneity	64531 ppm

Center	[0, 0, 0] m
Obs. points	7668
$B_{z,\max}$	0.1241 T
$B_{z,\min}$	0.1191 T
$B_{z,\text{mean}}$	0.1225 T
Homogeneity	40858 ppm

Table 3.13: Results of *PmLab* computation on the cylindrical FOV (left) and on the xy-plane (right), Configuration 5

Although a direct comparison is not entirely feasible, due to the conceptual similarity but geometric and volumetric differences in the magnetic structure between the two designs, the results for Configuration 5 show a clear increase in magnetic field strength compared to Configuration 4. Indeed, the use of a denser and more compact magnet arrangement results in an approximate increase of $0.1\,T$ in the average field intensity. This enhancement is achieved without a substantial deterioration in field homogeneity. However, it is important to note that the analysis for these configurations was conducted over a smaller cylindrical FOV, which does not extend significantly along the z-axis the direction in which the magnetic field typically decreases more sharply. As a result, the reported homogeneity values may be underestimated, as they do not fully capture the fringe effects and axial variations that would emerge in a larger or more elongated field of view.

Degree n	**Relative Contribution**	**Dominant Orders m**
0	9.985×10^{-1}	(uniform field)
1	1.213×10^{-9}	negligible
2	1.045×10^{-2}	$m = -2$ (XY), $m = 1$ (ZX)
3	2.503×10^{-8}	negligible

Table 3.14: Relative contribution of the main spherical harmonic degrees, Configuration 5

The analysis of spherical harmonics in Table 3.14 does not show significant differences compared to the previous ones. Only a larger contribution of the second degree component can be noted, among which the components of order $m = -2$ and $m = 1$ always have a predominant role.

4 Selection and Optimization

This chapter marks the beginning of the design refinement phase, whose purpose is trying to identify and improve the most suitable magnet configuration for low-field application.
Following the comparative analysis in Chapter 3, in the first part of this chapter, the choice of initial configuration will be made among the candidates previously analyzed. The choice is based on a combination of factors, including magnetic field intensity, spatial homogeneity, structural compactness, and practical feasibility.
Once the reference configuration has been selected, the chapter proceeds with its progressive optimization. The goal is to improve magnetic performances, in terms of strength of the field and homogeneity, by adjusting specific geometric and material parameters or adding new components. o guide this process, a genetic algorithm is employed, a type of evolutionary optimization particularly well suited to complex, non-linear problems. Its ability to efficiently explore a wide solution space makes it particularly effective in identifying configurations with superior magnetic performance.

4.1 Choice of the Reference Geometry

The configuration selected as the starting point for the optimization process is the tile-based Aubert ring structure described in Sec.3.3.2. Among the various geometries, this design provides a promising foundation thanks to its balanced combination of magnetic performance, manufacturability, and architectural flexibility. The choice is supported by several key considerations:

- **Higher field strength** in the region of interest, a key factor for improving the signal-to-noise ratio. A stronger field leads to a greater net magnetization of the spin system, which is essential for image contrast and sensitivity in low-field MRI;
- **Manufacturing simplicity**: although the structure consists of multiple

curved tiles arranged in concentric rings, each tile is identical in shape, size, and magnetization direction. This uniformity significantly simplifies the manufacturing process as the tiles can be mass-produced using a single mold and magnetized in bulk with the same direction;

- **Structural openness**: the configuration leaves a significant amount of free space between the two magnet layers, offering the opportunity to integrate supplementary components, either magnetic or non-magnetic, for improving field homogeneity;
- **Axial orientation of the magnetic field**, which facilitates seamless integration with standard RF coils, eliminating the need for custom redesign.

Despite its advantages, the selected configuration exhibits key limitations that must be resolved to satisfy the performance demands. In the previous analysis, the evaluation of field homogeneity was performed over a cylindrical field of view with a height of only $5\,cm$ and, even in this conservative situation, the configuration showed significant inhomogeneities, especially along the z-axis (as visible in Tables 3.14).
For the purposes of this study, however, the region of interest will be extended to a spherical FOV with a radius of $10\,cm$, which better reflects the dimensions of a realistic imaging volume.

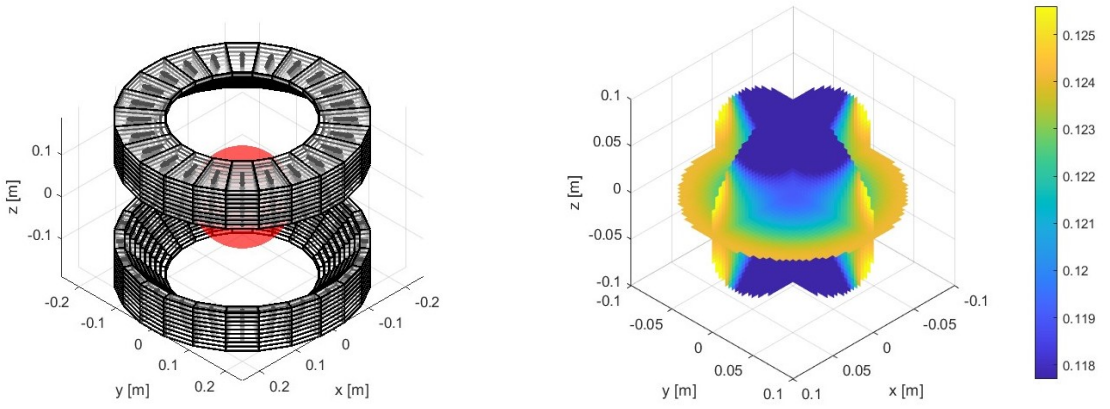

(a) (b)

Figure 4.1: Reference structure with the new spherical FOV (a) and computation of the B_z component on the new spherical FOV (b)[7]

[7]ATT: To emphasize the decrease in the field when moving on the z-axis, the color map has been adjusted to have the same chromatic scale as in figure 3.19

Center	[0, 0, 0] m
Obs. points	30976
$B_{z,\max}$	0.1264 T
$B_{z,\min}$	0.0909 T
$B_{z,\text{mean}}$	0.1191 T
Homogeneity	298467 ppm

Table 4.1: Results of PmLab computation of the reference structure on a spherical FOV

The increased volume makes even more difficult to maintain field uniformity within the FOV, especially near the two magnetic assemblies, where the axial field rapidly drops, as visible in Figure 4.1. This trend is confirmed by the results obtained in the simulation, shown in the Table 4.1, where we can observe that both the values of $B_{z,\max}$ and $B_{z,\min}$ are different from the ones in Table 3.13. In particular, $B_{z,\max}$ is slightly higher in the latter case, implying that the maximum field value is reached at a higher altitude, closer to the magnetic structures, whereas the $B_{z,\min}$ value is significantly lower (more than $0.2\,T$) to indicate how strong this drop is as we move towards the two magnetic layers along the z-axis. This increase in the difference between the two values leads to a substantial increase in the inhomogeneity within the FOV, which reaches a value almost 5 times higher than that seen in Configuration 5. Despite its limitations this configuration is adopted as the baseline for further analysis and optimization, thanks to its simplicity, magnetic efficiency, and promising potential for refinement.

4.2 Preliminary Adjustment: Axial Expansion

Before proceeding with the optimization steps, a preliminary adjustment to the baseline configuration is required. Specifically, the two magnetic layers are axially displaced to increase the gap between them. The motivation for this adjustment is illustrated in Figure 4.2, based on the `FEMM` simulation of Configuration 5. In the figure, the magnetic field intensity and flux lines are visualized across a vertical cross-section of the system. The central semicircular region represents the target FOV, while the two main magnet assemblies, located above and below, generate the static magnetic field.

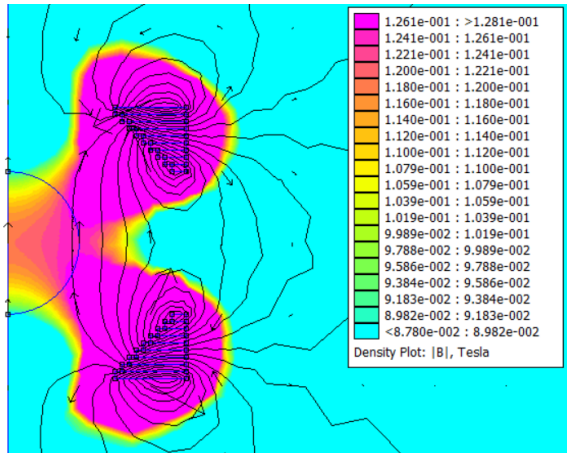

Figure 4.2: Magnetic field distribution in reference setup (FEMM)

Due to the original axial positioning of the magnetic layers, it can be observed that their influence does not sufficiently extend into the upper and lower extremities of the FOV. These peripheral regions, which lie closer to the z-axis boundaries of the FOV, are characterized by visibly weaker field magnitudes. This uneven coverage leads to significant spatial variation in field strength and thus deteriorates the global field homogeneity.

To address this issue, a preliminary structural adjustment is introduced: the axial gap between the two magnet layers is increased from $20\,cm$ to $25\,cm$. The purpose of this adjustment is to extend the region of effective magnetic coverage, allowing the high-field areas (in purple) to better encompass the FOV and produce a more uniform starting field distribution, thus offering a stronger foundation for further optimization.

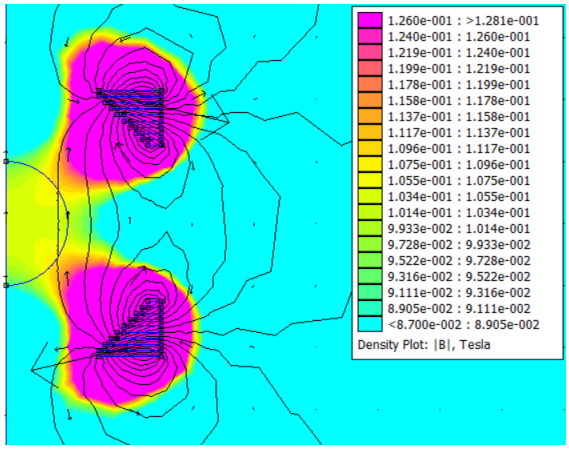

Figure 4.3: Magnetic field distribution after the adjustement (FEMM)[8]

[8]ATT: the same color scale of Figure 4.2 has been used.

The effect of increased gap is illustrated in Figure 4.3. Compared to the previous case, two key observations can be made. First, the overall field intensity within the FOV decreases slightly due to the increased distance between the magnetic sources and the central region. Second, this reduction is accompanied by a noticeable improvement in field uniformity, as visible from the more homogeneous color distribution. his qualitative observation is also supported by quantitative results obtained through PmLab simulations, as summarized in Table 4.2.

Center	[0, 0, 0] *m*
Obs. points	45686
$B_{z,\max}$	0.1044 T
$B_{z,\min}$	0.0948 T
$B_{z,\text{mean}}$	0.1005 T
Homogeneity	96189 ppm

Table 4.2: Results of PmLab computation after the axial adjustement

Starting from this stage of the analysis, all results will be systematically evaluated using both the FEMM finite element solver and the PmLab tool. To justify and support the combined use of both tools, a preliminary comparison is presented in this section. The aim is to evaluate the degree of agreement between the two simulation methods and to identify any systematic discrepancies that may influence subsequent analyses. Both solvers are applied to the same baseline configuration (the last one seen in the present section) and the magnetic field is computed on identical observation domains.

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.1044	0.1064	+1.91%
$B_{z,\min}$ [T]	0.0948	0.0922	-2.90%
$B_{z,\text{mean}}$ [T]	0.1005	0.1015	+0.59%
Homogeneity [ppm]	96189	140437	+44%
Points in FOV	45586	4736	—

Table 4.3: Comparison between PmLab and FEMM results for the reference configuration (25 *cm* gap)

The maximum value is slightly higher in FEMM compared to PmLab. This can be attributed to the fact that FEMM treats the problem as axisymmetric, assuming that each ring is perfectly continuous along its circumference. On the other end, PmLab models the rings as discretized arrays of consecutive tiles, each with individual orientation and position. This segmentation breaks the continuity of the

magnetic surface, resulting in a slight reduction of the peak field value.
In contrast the minimum value is lower in FEMM. This is likely due to the higher resolution of the FEMM simulation in detecting localized field drops within the observation domain (FEMM computes the field over a dense 2D semicircular cross-section, whereas PmLab samples the entire 3D spherical FOV, resulting in a lower point density). The increased field excursion observed in FEMM naturally results in a higher homogeneity value.

It is important to clarify that the discrepancy in homogeneity values observed between FEMM and PmLab is inherent and will persist throughout all subsequent analyses. This is primarily due to the different strategies the solvers use to sample the magnetic field within the region of interest:

- PmLab evaluates the magnetic field over a 3D spherical volume by generating a regular grid of observation points using MATLAB's `meshgrid` function. This ensures that the sampling points are evenly spaced throughout the domain;
- In FEMM, on the other hand, the field is computed using an adaptive triangular mesh, which subdivides the simulation domain into numerous interconnected triangles, adjusting the density of this mesh based according the local field behavior.

As a consequence, the set of observation points differs between the two tools, making it impractical to impose identical spatial sampling across simulations. In addition, since FEMM employs a finer and more adaptive mesh, especially in regions with steeper field gradients, it tends to produce consistently higher (and thus more conservative) estimates of field inhomogeneity.
For this reason, in our study, FEMM serves as the bottleneck: although PmLab provides reliable estimates, its sampling resolution is slightly lower due to the uniform but coarser 3D grid. Consequently, any configuration that meets the homogeneity constraints in FEMM can be considered robust, even if PmLab yields a more optimistic evaluation.

4.3 Step 1 – Inner Radii Optimization

As a first step in the optimization process of the magnetic structure, the attention was given to the inner radii of the magnetic layers. This phase represents the initial attempt to systematically refine the geometry of the system in order to explore its

potential to meet the design requirements. To this purpose, a Genetic Algorithm (GA) was utilized, a stochastic, population-based optimization method inspired by principles of natural selection and evolutionary theory.

A Genetic Algorithm explores the solution space by evolving a set of candidate configurations, where each individual is defined by a vector of decision variables, in this case the inner radii of the magnetic layers. The optimization starts from an initial population generated randomly. Each candidate solution is assessed using a fitness function that measures how effective a solution is in terms of optimization goals. The best-performing candidates are then chosen to serve as parents for the next generation. Offspring solutions are created by applying crossover, which blends characteristics from two selected parents, and mutation, which introduces stochastic alterations to promote diversity. This evolutionary process is iterated over successive generations, leading to a progressive improvement in solution quality.
In this specific case, the optimization was formulated as a multi-objective problem, aiming to balance two competing goals: the maximization of the magnetic field strength within the FOV and the minimization of field inhomogeneity across the same region. Unlike single-objective optimization, this approach leads a set of optimal trade-off solutions, collectively referred to as the Pareto front. A solution is part of the Pareto front if no other solution can improve one objective without compromising the other. As such, the front represents the full spectrum of compromises between field objectives, providing a comprehensive map of viable design alternatives. From this set of non-dominated solutions, the final configuration can be selected based according to the specific demands of the use case.

The optimization was applied to the internal radii of the magnetic rings, each one constrained between the $[0.15\,m,\ 0.25\,m]$ range. The algorithm was set to run for 50 generations with a population size of 100 individuals, ensuring both a broad exploration of the solution space and a stable convergence toward high-quality solutions. This choice, validated through multiple tests, proved to be effective in consistently reaching well-distributed Pareto fronts and robust solutions.
It is important to highlight that, unlike the previous analyses, where the FOV was divided using a fine mesh for accurate field mapping, in the optimization processes a coarser sampling was used to reduce the computational time. Although this approach results in a slight reduction in spatial resolution, it still provides a reliable estimate of the field to guide the optimization.

The output of the process is a Pareto front comprising non-dominated solutions, each representing a different trade-off between the two conflicting objectives: magnetic field strength and field homogeneity.

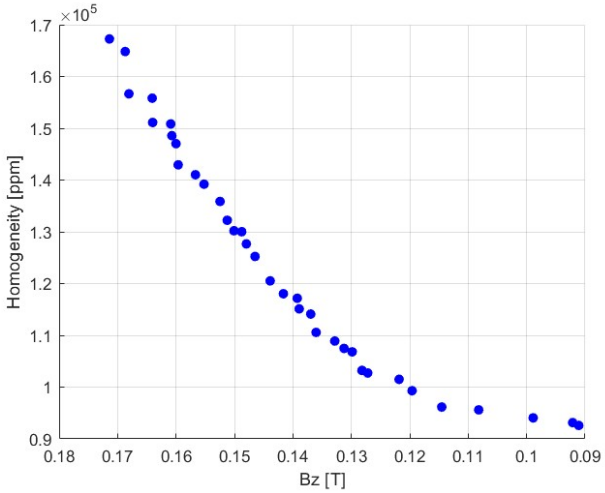

Figure 4.4: Pareto Front - Inner radii optimization

As clearly visible in the Pareto front (Figure 4.4), the two objectives are inherently conflicting, improving one typically results in the deterioration of the other. This observation highlights the necessity of adopting a multi-objective approach, as single-objective methods may fail to represent the complicate trade-offs present in the design space.

Among the available solutions, the one corresponding to the maximum magnetic field intensity was selected as the reference configuration. This choice reflects the two-step design strategy adopted in this work: first, prioritize the maximization of the magnetic field, then refine the geometry in a subsequent optimization phase to improve homogeneity.

One may wonder why, given the initial emphasis on maximizing the magnetic field intensity, the trivial solution with all internal radii set to $0.15\,m$ was not chosen as a starting point. The problem is that this configuration generates significant spatial variations in the magnetic field that are extremely difficult to mitigate in the next optimization steps.

The internal radii associated with this solution, ordered from the innermost to the outermost ring, are:

$$R_{\text{in}} = [0.164,\ 0.170,\ 0.159,\ 0.156,\ 0.166,\ 0.163,\ 0.171,\ 0.170,\ 0.162]\,m$$

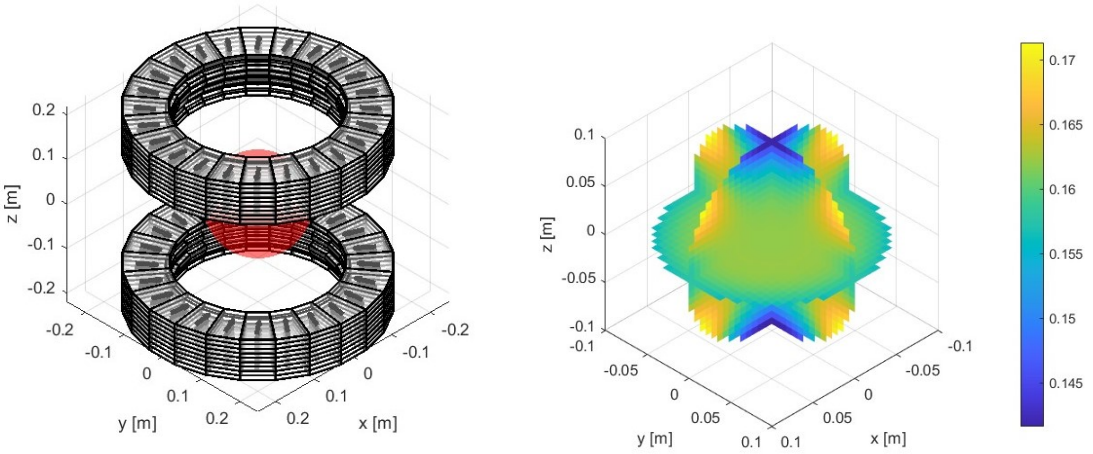

(a) (b)

Figure 4.5: Step 1 - 3D view of the optimized magnetic structure (a) and PmLab B_z field distribution in the FOV (b)

Figures 4.5 and 4.6 show the practical implementation of the optimized structure and the results of the two computational methods that are summarized in Table 4.4.

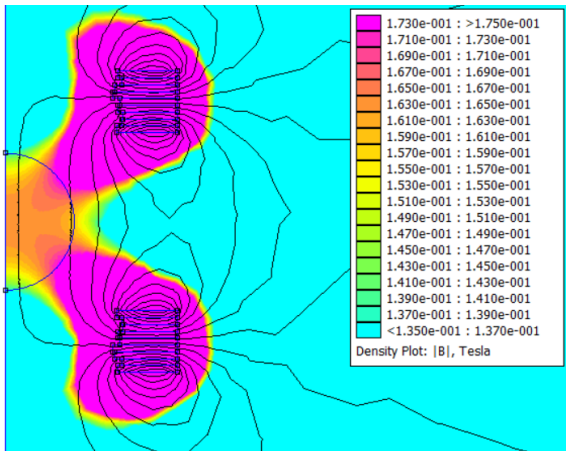

Figure 4.6: Step 1 - B_z field distribution (FEMM)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.1713	0.1739	+1.5%
$B_{z,\min}$ [T]	0.1443	0.1367	-5.2%
$B_{z,\text{mean}}$ [T]	0.1621	0.1621	=
Homogeneity [ppm]	166777	229731	+37%
Points in FOV	3544	4736	—

Table 4.4: Step 1 - Comparison between PmLab and FEMM results

Table 4.4 shows that the discrepancy between PmLab and FEMM's results is still similar to the previous case in Sec.4.2, with a general alignment on the field values and a much more significant difference on the homogeneity values, for the already mentioned reasons.
Compared to Table 4.3 (before the optimization of the internal radii) is it noticeable how the magnetic performances have significantly improved, with an overall increase of approximately 60%, but at expenses of homogeneity which sees a considerable increase (around 65% for both PmLab and FEMM).

4.4 Step 2 - External Ring Magnets

Following the optimization of the internal radii of the magnetic layers, this next step aims to further enhance the axial magnetic field in the region of interest by introducing a set of external annular permanent magnets.
These cylinders are positioned concentrically around the main magnet assembly, without entering the central area between the layers or interfering the FOV. Their job is to boost the total magnetomotive force, giving extra push to the magnetic flux and making the field inside the FOV stronger.

From a physical perspective, the influence of the outer cylinders can be interpreted through an analogy with electrical circuits: the magnetomotive force (MMF) acts as a voltage source, the magnetic flux corresponds to electric current, and magnetic reluctance is analogous to electrical resistance. Just like adding batteries in series raises the voltage in an electric circuit, adding external magnetized cylinders increases the total MMF in the magnetic system. This makes the magnetic flux in the central area stronger, even without changing the internal layout.

The external cylindrical magnets adopted in this phase of the study were assigned fixed dimensions: a radial thickness of $4\,cm$ and an axial height of $5\,cm$. While the optimization focused only on the internal radii and axial positions of the cylinders, their section was kept constant.
This design decision was motivated by both manufacturing considerations and material integrity constraints. From a practical point of view, these dimensions enable the fabrication of the cylinders as monolithic magnetic components, eliminating the need for segmentation. This is also made possible by the fact that these external cylinders are designed with axial magnetization (not a radial one, as the magnetic layers).

Moreover these dimensions reflect a careful balance between performance and practicality. They allow each cylinder to produce enough magnetomotive force to boost the field, while keeping the size small enough to avoid adding too much weight and compromise the portability of the system.
The algorithm parameters remained unchanged from the previous stage. By allowing both the position and radial extent of each cylinder to vary independently, and without imposing symmetry constraints (at least for the moment), the search space was kept sufficiently broad to explore non-trivial configurations capable of enhancing the field.
To guide the search, the following constraints were applied:

- The internal radii of the cylinders were allowed to vary between $0.26\,m$ and $0.40\,m$;
- The axial positions were constrained between $-0.12\,m$ and $0.12\,m$,, to cover the space between the layers without overlapping them axially.

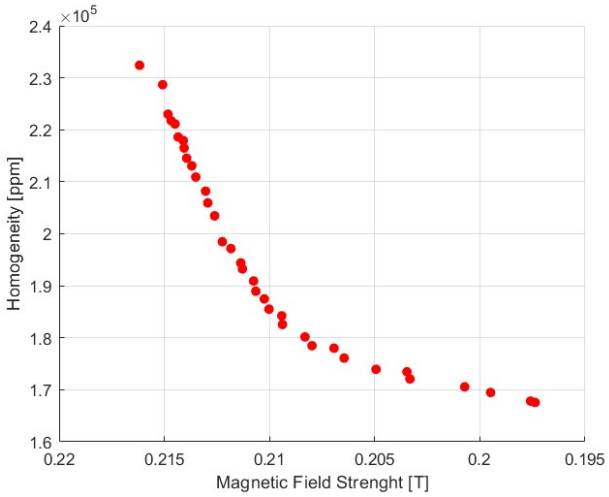

Figure 4.7: Pareto Front - External rings optimization

The results of the optimization are illustrated in Figure 4.7. The plot shows the usual trade-off between the two goals: setups with stronger magnetic fields tend to be less homogeneous, while those with better homogeneity usually have weaker fields.
The comparison shown in Figure 4.8 clearly highlights the significant increase in magnetic field strength made possible by the addition of the external elements. Moreover, the shape and slope of the red front provide additional insight: the curve appears to tilt rapidly, suggesting that the current configuration is approaching a

local performance maximum. In other words, further improvements in field intensity, within the limits of the current design and material constraints, would likely come at the cost of disproportionate losses in homogeneity.

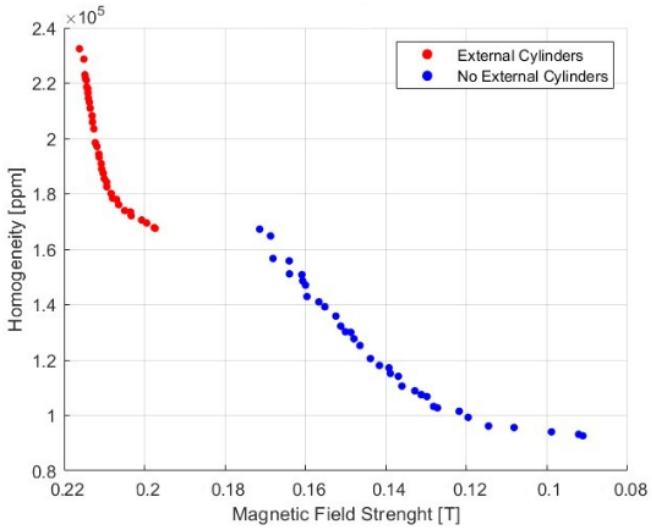

Figure 4.8: Pareto Front - Before and after the insertion of the external rings

Following with the strategy adopted in the previous optimization stage, the configuration selected corresponds to the solution with the highest magnetic field strength. The resulting optimization variables for this solution are:

$$R_{\text{in}} = [0.289,\ 0.271,\ 0.277]\, m$$

$$z_{\text{pos}} = [-0.045,\ 0.019,\ 0.079]\, m$$

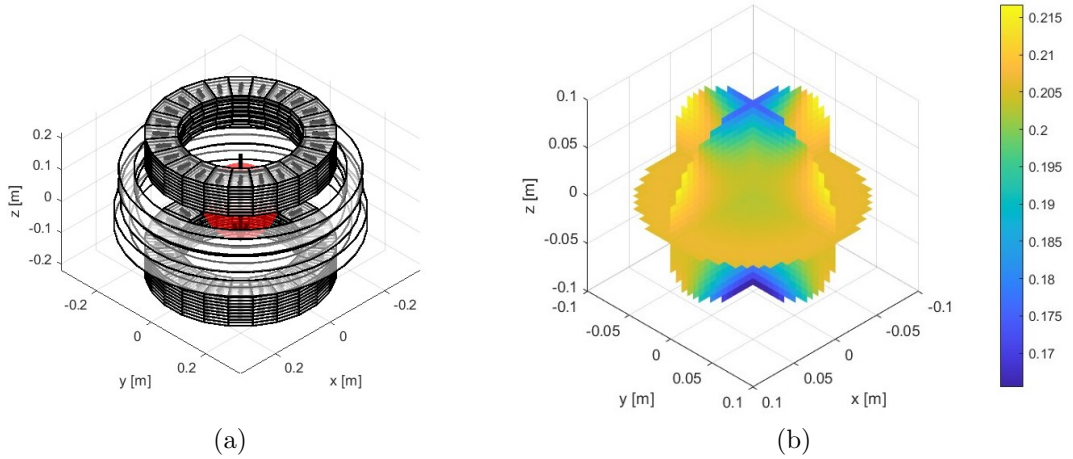

(a) (b)

Figure 4.9: Step 2 - 3D view of the optimized magnetic structure (a) and PmLab B_z field distribution in the FOV (b)

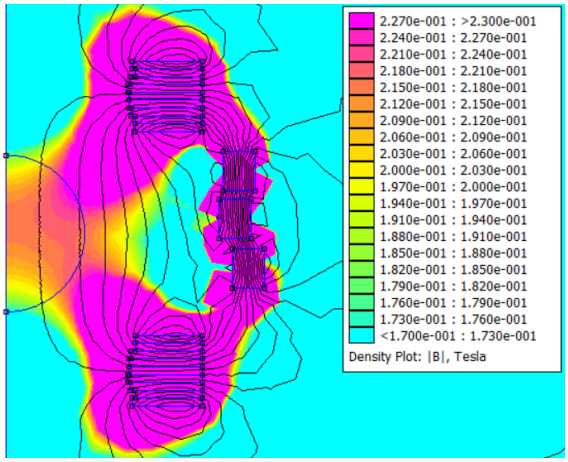

Figure 4.10: Step 2 - B_z field distribution (FEMM)

From Figures 4.9 and 4.10, an increase in field strength is evident, especially in the central FOV areas, near the $z = 0$ plane. However, this field increase does not follow the same course over the entire FOV, as there are areas that are much less affected by the influence of the magnets, such as the two spherical caps at the extremes along the z-axis. Therefore, even with additional material, it is difficult to significantly influence these peripheral areas.
An additional observation, the slight asymmetry of the magnetic field with respect to the xy-plane is due to the non-symmetrical arrangement of the outer cylinders.

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.2163	0.2203	+1.8%
$B_{z,\min}$ [T]	0.1659	0.1672	+0.8%
$B_{z,\text{mean}}$ [T]	0.2029	0.2101	+3.5%
Homogeneity [ppm]	233061	282941	+21%
Points in FOV	3544	4736	—

Table 4.5: Step 2 - Comparison between PmLab and FEMM results

Table 4.5 shows the results with the two calculation methods.
It is interesting to compare the results obtained before and after inserting the outer rings, as presented in Table 4.6.

PmLab results			
	Before	**After**	**Diff**
$B_{z,\max}$ [T]	0.1713	0.2163	+26%
$B_{z,\min}$ [T]	0.1443	0.1690	+17%
$B_{z,\text{mean}}$ [T]	0.1621	0.2029	+25%

Table 4.6: Comparison between PmLab results, before and after the external rings insertion

This non-uniform improvement indicates that the external cylinders do not affect all regions of the FOV uniformly. In particular, the minimum field values are typically located near the top and bottom edges of the FOV, along the z-axis. These are precisely the areas where the contribution of the external magnets is less effective.
This is the main cause of the substantial increase in inhomogeneities, and it is important to note that this has little to do with the optimization process employed. It is an inherent feature of this configuration that we will attempt to mitigate in other ways, as we will see in the next section.

4.5 Step 3 - Field-attenuating Rings

In this phase, the focus shifts entirely on improving the homogeneity of the magnetic field within the FOV, intentionally excluding field intensity from the set of optimization objectives.
To achieve this, the introduction of a new set of magnetic elements is proposed: two pairs of annular permanent magnets featuring radial magnetization in the opposite direction to that of the main magnetic layers.
Theses "attenuating rings" are placed on both sides of the central xy-plane, in correspondence with the regions where the magnetic field reaches its maximum (see Figures 4.9 and 4.10). Their function is not to contribute to the total field, but to help balancing it. The opposite radial magnetization creates localized magnetic vortices that, if placed correctly, introduce negative magnetic flux in areas where the main field it's too strong.
This technique does not represent the most efficient strategy for improving field uniformity, as it relies on reducing the field where it is too strong, rather than boosting it where it is too weak. As a result, it leads to a decrease in the overall field strength. Still, it remains the only feasible option in this context, as the insertion of magnetic material into the cylindrical bore is not possible, which limits the capability to influence the field at the caps of the FOV, where the field drops are

significant.

The population size was increased to 500 individuals to account for the higher dimensionality of the problem, which now involves 16 optimization variables, including: the inner radius, axial position, radial thickness, and axial thickness of each of the four attenuating rings. For manufacturability, the axial thickness was constrained to discrete values of either 0.01 m or 0.02 m.
Appropriate constraints have been applied to all design parameters to ensure that the resulting geometry is physically feasible, making sure to avoid the intersection between with any of the pre-existing magnetic structures or between themselves and preventing them to approach too closely the FOV, essential to preserve the physical space required for patient positioning.
The results of the single-objective optimization process are shown in Table 5.14 and, together with this, their representation in the *PmLab* environment (Figure 4.11a).

Ring #	$R_{in}\,[m]$	$t_r\,[m]$	$z_{pos}[m]$	$t_z[m]$
1	0.159	0.085	-0.084	0.02
2	0.153	0.046	-0.035	0.01
3	0.194	0.021	0.032	0.02
4	0.154	0.071	0.011	0.02

Table 4.7: Optimized parameters for the attenuating rings

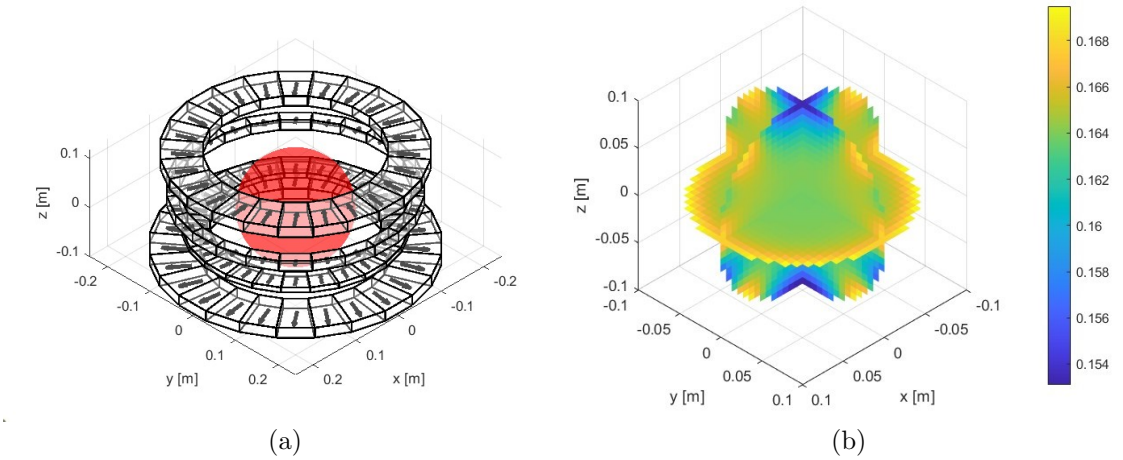

(a) (b)

Figure 4.11: Step 3 - 3D view of the optimized magnetic structure (a) and PmLab B_z field distribution in the FOV (b)[9]

[9]Magnetic layers and outer rings are excluded in the the left figure for representational clarity.

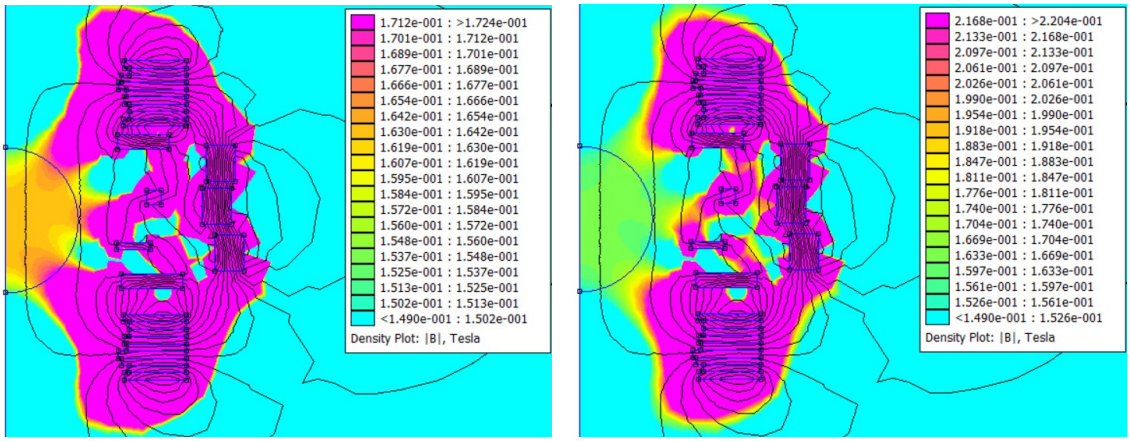

(a) (b)

Figure 4.12: Step 3 - B_z field distribution (FEMM) with two different color scales

Figure 4.12 shows the distribution of the B_z component of the magnetic field obtained from FEMM simulations after the inclusion of the attenuating rings. The figure 4.12a presents the field using a color scale automatically adjusted to the new field range, while 4.12b adopts the same maximum value used previously in Figure 4.10. This comparison allows for a clearer visualization of the mitigation effect introduced by the new ring structures; it is clear, indeed, how the maximum field intensity has been significantly reduced compared to the earlier configuration. At the same time, the uniformity of the color distribution in indicates a significantly improved level of homogeneity, as also confirmed by the numerical results reported in Table 4.8.

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\text{max}}$ [T]	0.1695	0.1723	+1.65%
$B_{z,\text{min}}$ [T]	0.1531	0.1491	-2.22%
$B_{z,\text{mean}}$ [T]	0.1638	0.1639	=%
Homogeneity [ppm]	86360	137933	+59
Points in FOV	3544	4736	—

Table 4.8: Step 3 - Comparison between PmLab and FEMM results

It is interesting to observe that the attenuating effect of the rings is quite pronounced, strong enough to bring the field magnitude back to levels comparable to those observed in Step 2 (before the insertion of the external boosting rings). This effect is especially evident when analyzing the average and the maximum value of B_z, which show a reduction of approximately 20% compared to the configuration in Step 3. On the other hand, the minimum field value decreases only slightly, by about 8%, suggesting that the attenuation effect primarily targets the regions of

highest magnetic flux density, which was the desired effect.

Although the qualitative outcome of this optimization step follows the intended design strategy, the quantitative results remain far from acceptable homogeneity levels. The main limiting factor continues to be the presence of inhomogeneities concentrated near the spherical caps of the FOV, located along the z-axis at the top and bottom of the imaging volume.
These areas represent a magnetic bottleneck, as they are inherently difficult to influence using the current magnet architecture, which lacks of components positioned near the central axis of the main magnet layers.
To overcome this limitation, a structural adaptation of the current configuration is necessary. Rather than a radical redesign, the next step involves a targeted refinement of the existing layout, carefully adjusted to better address the magnetic requirements of the system.

5 Redesign and Optimization

Starting again from the structure presented in Configuration 5 (Sec.3.3.2) and, following the modifications introduced in Chapter 4, the reference design is further revised by increasing the axial distance between the two magnetic assemblies.
In this updated configuration, the spacing between the inner rings of the two magnet layers is extended from $20\,cm$ to $32\,cm$ (while in Chapter 4 the extension was only $5\,cm$). As a result, the distance from the center of the FOV to the outer edge of each magnetic assembly increases from $19\,cm$ to $25\,cm$, as shown in Figure 5.1.

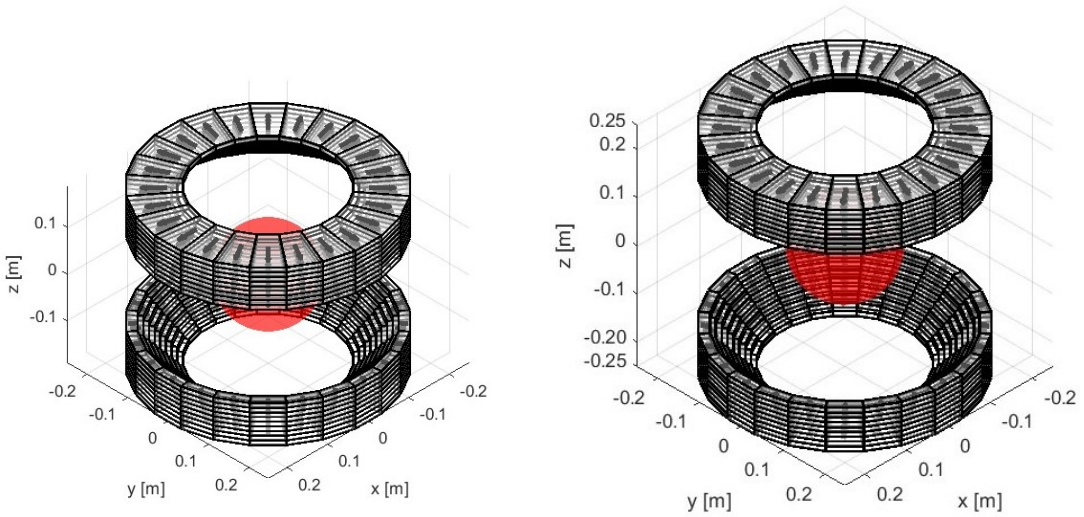

(a) (b)

Figure 5.1: Reference structure, Configuration 5 (a) and new extended setup (b)

Before proceeding with improvements to the structure and further analysis of the magnetic field, it is necessary to verify that this modification does not compromise the anatomical compatibility of the system. In particular, it must be ensured that the center of the FOV remains properly aligned with the region of interest, in this case the human brain, across the full range of possible users.
To this end, the analysis refers to typical anatomical proportions observed in adult individuals. The vertical distance between the center of the brain, typically lo-

cated near the midpoint of the cranial volume, and the shoulder plane generally falls within the range of 25 to 30 cm, depending on individual height, neck length, and posture. In certain cases, particularly in individuals of smaller stature or with shorter necks, this distance may be slightly less than 25 cm. Nevertheless, given that the imaging system is designed with a relatively large field of view, this does not represent a constraint on its usability.
Hence, by setting the distance between the FOV center and the outer edge of each magnetic assembly to 25 cm, the system remains fully compatible with the anatomical variability found in the general population. As such, the structural modification can be considered both safe and universally applicable, without loss of generality with respect to the original configuration.

The present section now provides both a quantitative and qualitative analysis of the magnetic field distribution resulting from the updated configuration, focusing, in particular, on how the increased spacing between the magnetic assemblies affects the uniformity, intensity and spatial characteristics of the field.
As shown in Figure 5.2, the increased distance between the two magnetic layers has significantly altered the magnetic field distribution within the FOV. In contrast to previous configurations, the new setup exhibits a reversed behavior. Here, the maximum field intensity is concentrated in the spherical caps, while the minimum field occurs near xy-plane.

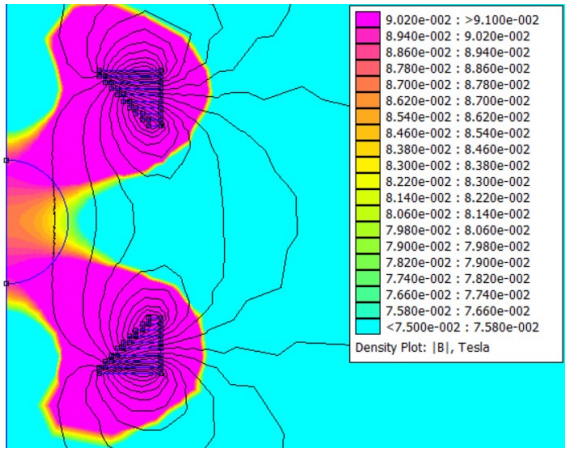

Figure 5.2: Magnetic field distribution in the new setup (FEMM)

However, the increased spacing does not come without trade-offs. As evidenced by the data in Table 5.3, the new configuration exhibits a general decrease in magnetic field strength throughout the FOV. This trend becomes even more apparent when looking at Table 5.4, which summarizes magnetic field values across the three

configurations examined.

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.0912	0.0932	+2.1%
$B_{z,\min}$ [T]	0.0783	0.0771	-1.5%
$B_{z,\text{mean}}$ [T]	0.0856	0.0881	+2.9%
Homogeneity [ppm]	150700	182945	+21%
Points in FOV	12712	4736	—

Table 5.1: Comparison between PmLab and FEMM results in the new setup

Parameter	**Ref. setup**	**1st Mod.**	**2nd** Mod.
$B_{z,\max}$ [T]	0.1264	0.1044	0.0892
$B_{z,\min}$ [T]	0.0909	0.0948	0.0756
$B_{z,\text{mean}}$ [T]	0.1191	0.1005	0.0828
Homogeneity [ppm]	298467	96189	163301

Table 5.2: Comparison of magnetic field PmLab results among the three setups

Despite the fact that the magnetic performance of this new configuration appears to be limited in terms of both field strength and homogeneity, this should not be regarded as a drawback. In fact, these results are not the primary focus at this stage. The real benefit of the current setup lies in a crucial improvement not seen in the previous versions: a substantial increase in magnetic field intensity within the two end-cap regions along the z-axis, which had been a major weak point of the system.
The fact that the field now shows a drop in the central xy-plane is not a major concern. Indeed, it is considerably easier to apply magnetic corrections in this portion of the FOV, as the insertion of additional material is only feasible in the space between the two magnetic assemblies, outside the central bore in which the FOV is contained (obviously without interfering with it).

For this reason, the analysis will proceed based on this newly established configuration, following the same methodology adopted in previous sections, firstly trying to maximize the strength of the field and, secondly, refining its uniformity through specific adjustments.

5.1 Step 1 - Inner Radii Optimization

After setting up the new configuration, the present step consists in optimizing the inner radii of the individual magnet layers. The optimization procedure, including parameters, variables, and constraints, follows the same methodology described in the previous chapter, in Section 4.3.
However, in this case, the analysis will be extended to include multiple solutions, trying to build a more comprehensive picture of how the magnetic performance evolves as additional design elements are introduced.
Rather than focusing on a single optimized configuration, the analysis explores a sequence of incremental refinements, each building upon the previous one, thus outlining an evolutionary path of the system.
In this context, the optimization process assumes a tree-like structure, where each branch represents a possible direction of development. By exploring these various paths, it can be better understood which changes actually improve the field, how those improvements build on each other and which design strategies offer the best compromise between field strength and homogeneity.
It is important to underline that this analysis is not intended to be exhaustive. Conducting a full optimization for each possible outcome of a previous iteration would result in computational demands that are far from the practical feasibility. Instead, the chosen strategy focuses on highlighting the most meaningful directions the design can take, enabling targeted improvements while maintaining a reasonable computational cost.

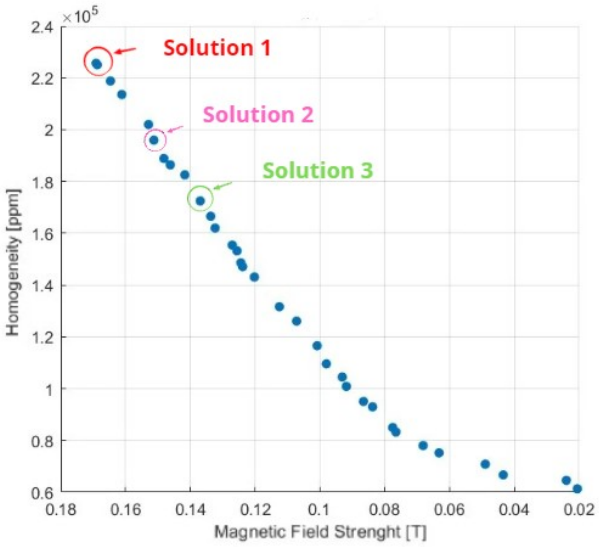

Figure 5.3: Pareto Front - Inner radii optimization

From the set of non-dominated solutions, shown in Fig.5.3, three initial configurations have been selected (highlighted in red, orange, and green in the figure) as starting points for further evolution and refinement. To this end, a lower limit of approximately $0.14\,T$ has been set on the average magnetic field strength to avoid excessively weak configurations that would be less suitable for imaging applications.The three selected configurations from the Pareto front are now presented in detail.

Solution 1

$$R_{\text{in}} = [0.15,\ 0.153,\ 0.15,\ 0.15,\ 0.15,\ 0.172,\ 0.15,\ 0.15,\ 0.15]\,m$$

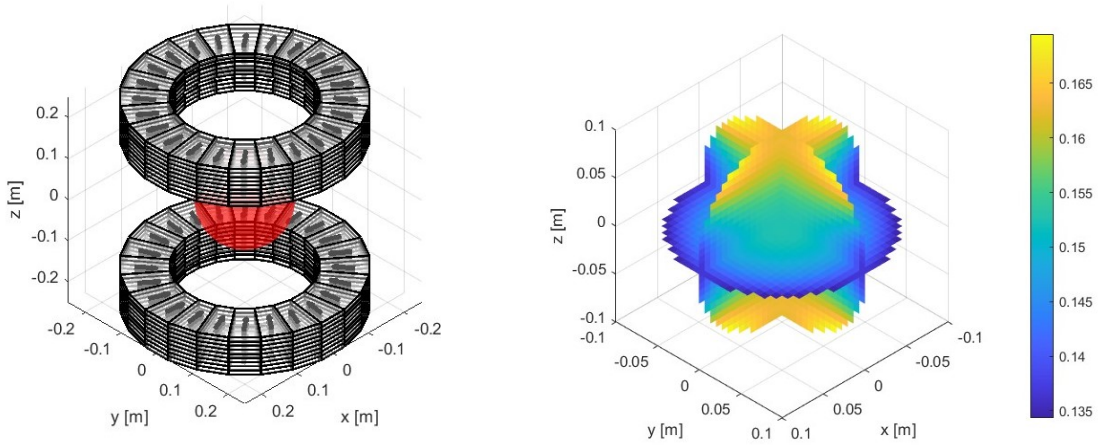

(a) (b)

Figure 5.4: Step 1 - Solution 1, structure (a) and PmLab B_z field distribution (b)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\text{max}}$ [T]	0.1695	0.1720	+1.4%
$B_{z,\text{min}}$ [T]	0.1343	0.1365	+1.6%
$B_{z,\text{mean}}$ [T]	0.1528	0.1577	+3.2%
Homogeneity [ppm]	230206	224749	-2.3%
Points in FOV	3544	4736	—

Table 5.3: Solution 1 - Comparison between PmLab and FEMM results

Solution 2

$$R_{\text{in}} = [0.162,\ 0.165,\ 0.15,\ 0.166,\ 0.156,\ 0.177,\ 0.155,\ 0.158,\ 0.163]\,m$$

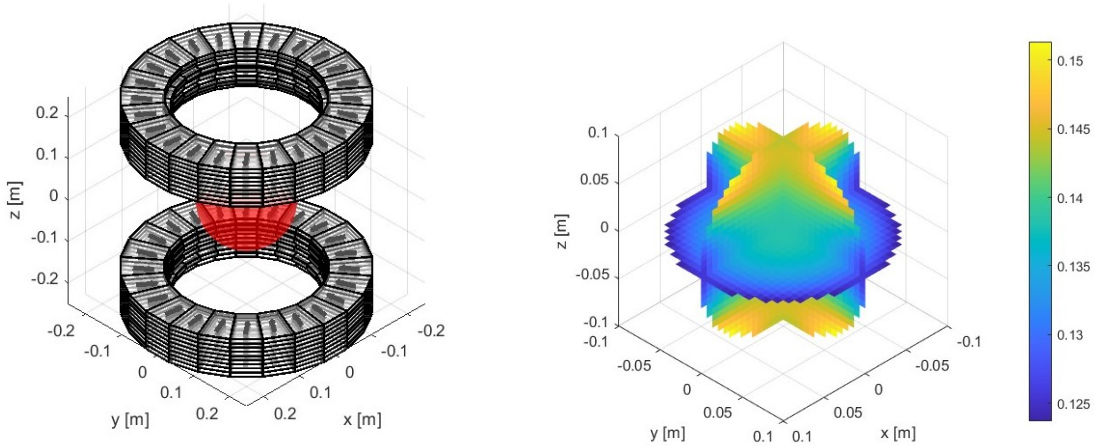

(a) (b)

Figure 5.5: Step 1 - Solution 2, structure (a) and PmLab B_z field distribution (b)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.1513	0.1541	+1.8%
$B_{z,\min}$ [T]	0.1237	0.1256	+1.5%
$B_{z,\text{mean}}$ [T]	0.1384	0.1423	+2.2%
Homogeneity [ppm]	199761	199056	-0.4%
Points in FOV	3544	4736	—

Table 5.4: Solution 2 - Comparison between PmLab and FEMM results

Solution 3

$$R_{\text{in}} = [0.156,\ 0.159,\ 0.187,\ 0.159,\ 0.167,\ 0.174,\ 0.171,\ 0.169,\ 0.167]\, m$$

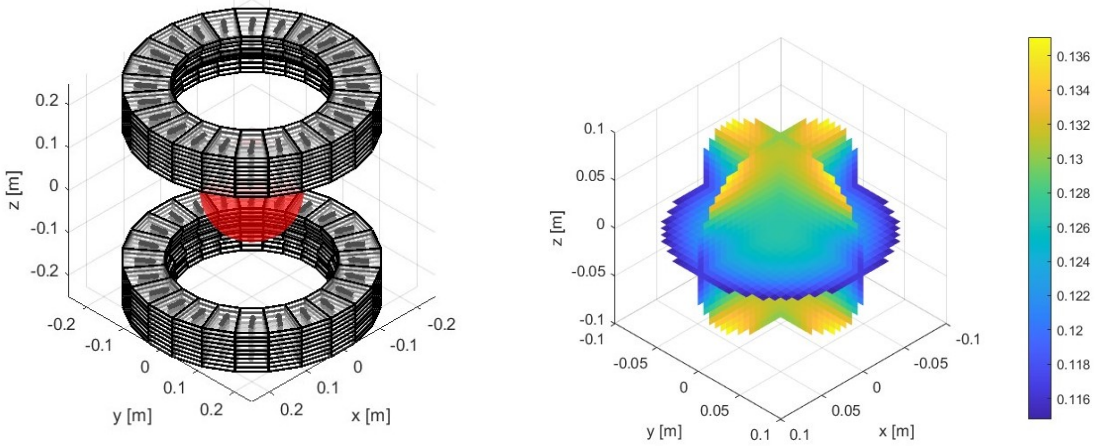

(a) (b)

Figure 5.6: Step 1 - Solution 3, structure (a) and PmLab B_z field distribution (b)

Metric	PmLab	FEMM	Difference
$B_{z,\max}$ [T]	0.1371	0.1386	+1.1%
$B_{z,\min}$ [T]	0.1148	0.1169	+1.8%
$B_{z,\text{mean}}$ [T]	0.1268	0.1296	+2.2%
Homogeneity [ppm]	175983	167443	-4.4%
Points in FOV	3544	4736	—

Table 5.5: Solution 3 - Comparison between PmLab and FEMM results

Given the increased axial distance between the two magnetic assemblies and the resulting change in the magnetic field distribution compared to the original setup, it is not surprising that the optimized solutions converge toward a relatively simple configuration in which all inner radii approach their maximum permitted value of $0.15\,m$.
While in the original configuration setting all $R_{in} = 0.15\,m$ would have increased the peak magnetic field at the expense of significantly reduced homogeneity (due to the limited influence on the end-caps), in this new setup the tendency towards this "trivial" solution does not lead to a dramatic loss in homogeneity, making it a reasonable and effective option.
Structurally, the three selected configurations are very similar to each other, with the main difference being the progressive enlargement of the magnet layers' diameter. This geometric variation naturally leads to a reduction in the magnetic field, due to the lower amount of magnetic material present in the system.
FEMM-based representations of the magnetic field have been omitted in this case, as their distribution patterns closely resemble those already shown in Figure 5.2, and do not provide additional insight beyond the data already discussed.

5.2 Step 2 - External Ring Magnets

As already explored in Section 4.4, this next step focuses on the strengthening of the magnetomotive force of the system by adding external magnet rings. A new optimization phase is set up to identify the best configuration of these rings, with some adjustments made to account for the geometric changes introduced in the current setup.
Specifically, four external rings are now considered instead of three, in order to make full use of the increased axial spacing between the two main magnet assemblies. Indeed, an excessive spacing between the rings could cause a dispersion of the magnetic field lines, reducing their effectiveness.

Unlike the approach adopted in the previous chapter, a symmetry constraint is introduced here, as there is no physical or functional reason to assume that the external rings should be asymmetrically placed.
The optimization process is carried out independently on each of the three previously explored configurations. The results of this optimization step are presented below.

5.2.1 Solution 1

Figure 5.7 shows the Pareto front obtained from the optimization process starting from Solution 1. From this set of non-dominated solutions, two configurations with opposite characteristics have been selected for further analysis.

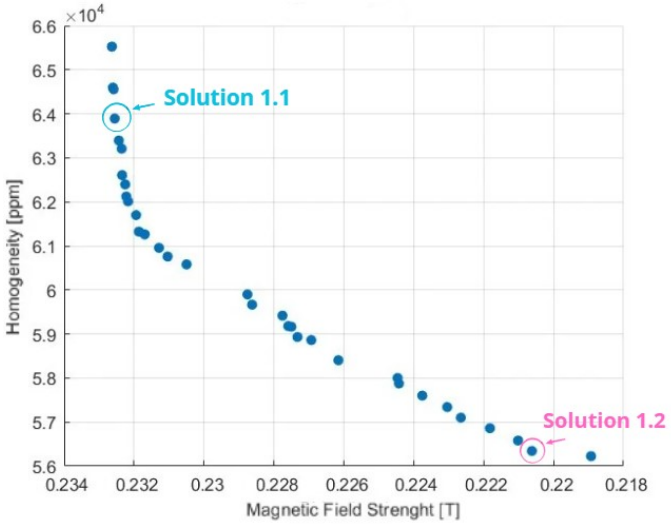

Figure 5.7: Pareto Front - Optimization results based on Solution 1

Solution 1.1

$$R_{\text{in}} = [0.254,\ 0.265,\ 0.265,\ 0.264]\, m$$

$$z_{\text{pos}} = [-0.096,\ -0.033,\ 0.033,\ 0.096]\, m$$

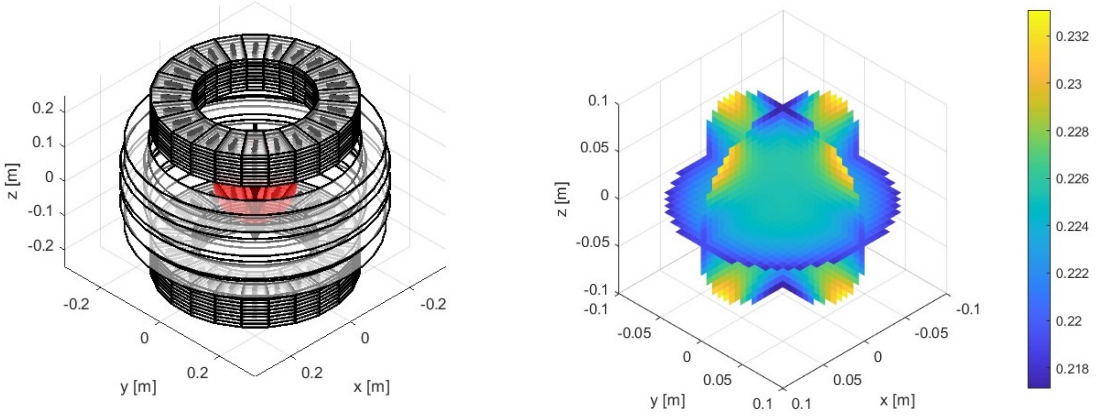

(a) (b)

Figure 5.8: Step 2 - Solution 1.1, structure (a) and PmLab B_z field distribution (b)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.2327	0.2359	+1.3%
$B_{z,\min}$ [T]	0.2182	0.2152	-1.4%
$B_{z,\text{mean}}$ [T]	0.2254	0.2272	+0.7%
Homogeneity [ppm]	64141	91252	+42%
Points in FOV	3544	4736	—

Table 5.6: Solution 1.1 - Comparison between PmLab and FEMM results

Solution 1.2

$$R_{\text{in}} = [0.339,\ 0.265,\ 0.265,\ 0.339]\, m$$

$$z_{\text{pos}} = [-0.096,\ -0.033,\ 0.033,\ 0.096]\, m$$

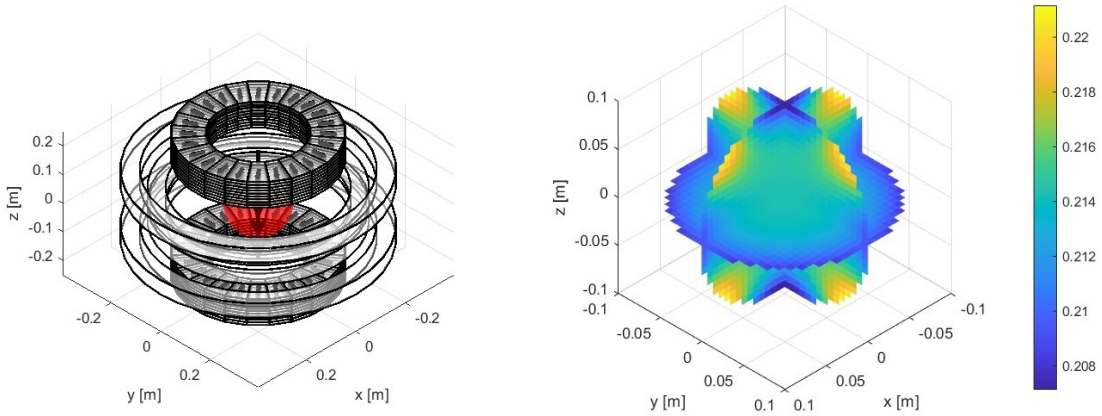

(a) (b)

Figure 5.9: Step 2 - Solution 1.2, structure (a) and PmLab B_z field distribution (b)

Metric	PmLab	FEMM	Difference
$B_{z,\max}$ [T]	0.2207	0.2238	+1.4%
$B_{z,\min}$ [T]	0.2086	0.2054	-1.6%
$B_{z,\text{mean}}$ [T]	0.2144	0.2163	+0.8%
Homogeneity [ppm]	56367	85103	+50%
Points in FOV	3544	4736	—

Table 5.7: Solution 1.2 - Comparison between PmLab and FEMM results

5.2.2 Solution 2

The Pareto front shown in Figure 5.13 corresponds to the optimization process initiated from Solution 2. As in the previous case, a wide range of trade-off solutions was generated by varying the configuration of the external magnetic rings.
Two representative points were selected for further analysis.

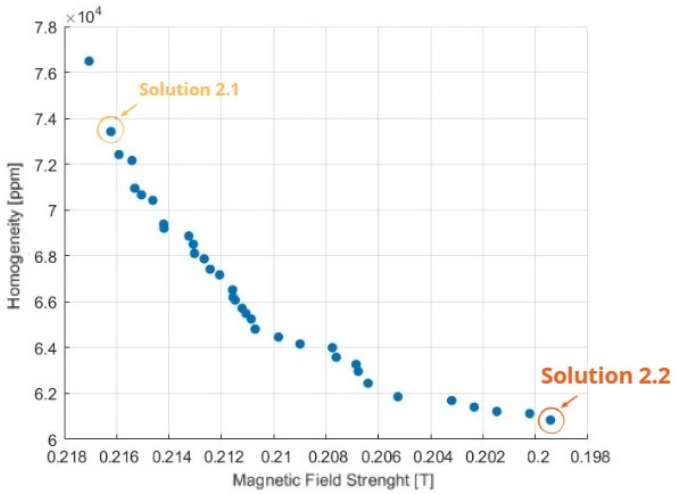

Figure 5.10: Pareto Front - Optimization results based on Solution 2

Solution 2.1

$$R_{\text{in}} = [0.253,\ 0.257,\ 0.257,\ 0.253]\, m$$

$$z_{\text{pos}} = [-0.105,\ -0.033,\ 0.033,\ 0.105]\, m$$

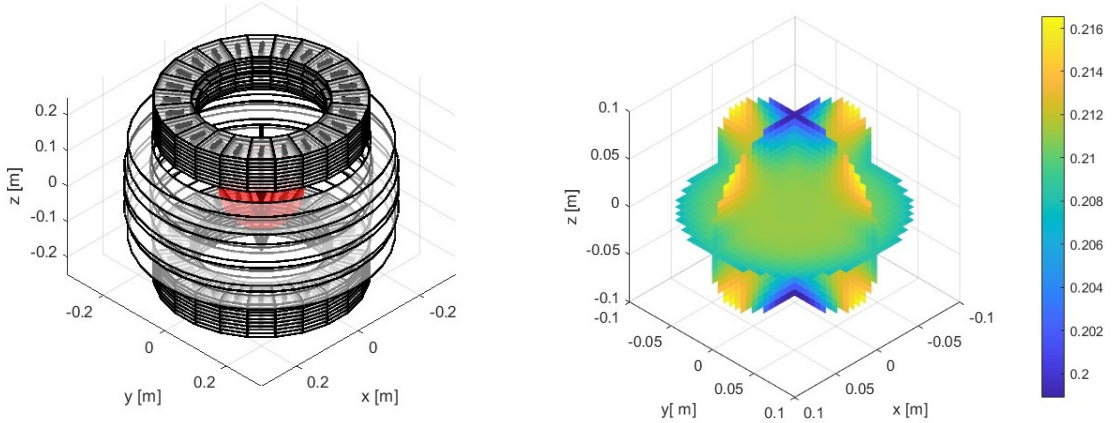

(a) (b)

Figure 5.11: Step 2 - Solution 2.1, structure (a) and PmLab B_z field distribution (b)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.2162	0.2192	+1.3%
$B_{z,\min}$ [T]	0.2007	0.1963	-2.2%
$B_{z,\text{mean}}$ [T]	0.2111	0.2120	+0.4%
Homogeneity [ppm]	73316	107833	+45%
Points in FOV	3544	4736	—

Table 5.8: Solution 2.1 - Comparison between PmLab and FEMM results

Solution 2.2

$$R_{\text{in}} = [0.332,\ 0.283,\ 0.283,\ 0.332]\, m$$

$$z_{\text{pos}} = [-0.114,\ -0.038,\ 0.038,\ 0.114]\, m$$

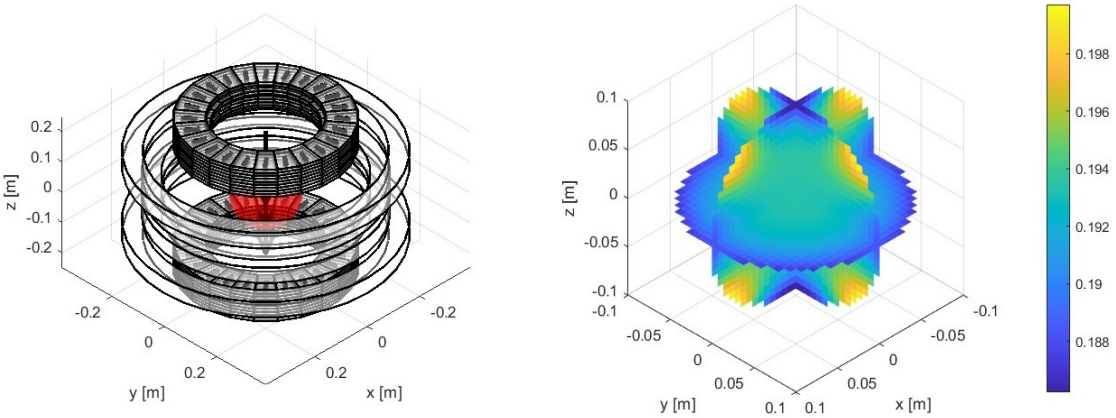

(a) (b)

Figure 5.12: Step 2 - Solution 2.2, structure (a) and PmLab B_z field distribution (b)

Metric	PmLab	FEMM	Difference
$B_{z,\max}$ [T]	0.1994	0.2020	+1.3%
$B_{z,\min}$ [T]	0.1876	0.1845	-1.7%
$B_{z,\text{mean}}$ [T]	0.1934	0.1945	+0.5%
Homogeneity [ppm]	60971	89974	+47%
Points in FOV	3544	4736	—

Table 5.9: Solution 2.2 - Comparison between PmLab and FEMM results

5.2.3 Solution 3

The third Pareto front originates from the optimization process applied to Solution 3, which was initially characterized by lower magnetic field values. Despite this, the optimization still yields a meaningful set of configurations with improved performance.

Among the results, two contrasting solutions have been selected: Solution 3.1, favoring stronger field generation, and Solution 3.2, prioritizing field uniformity.

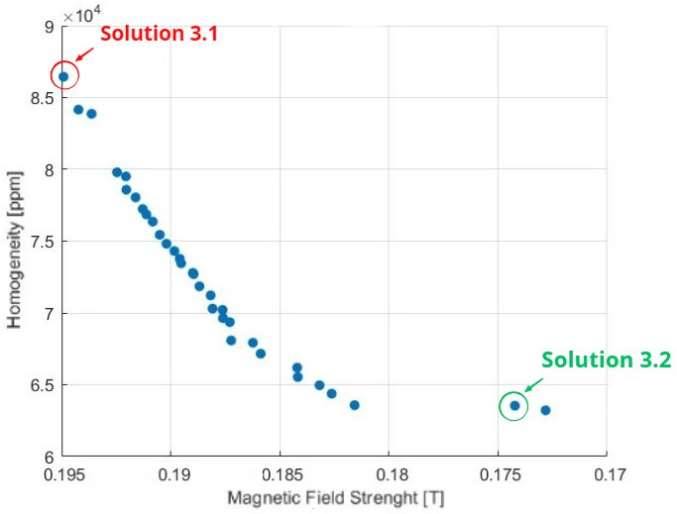

Figure 5.13: Pareto Front - Optimization results based on Solution 3

Solution 3.1

$$R_{\text{in}} = [0.298,\ 0.254,\ 0.254,\ 0.298]\, m$$

$$z_{\text{pos}} = [-0.104,\ -0.037,\ 0.037,\ 0.104]\, m$$

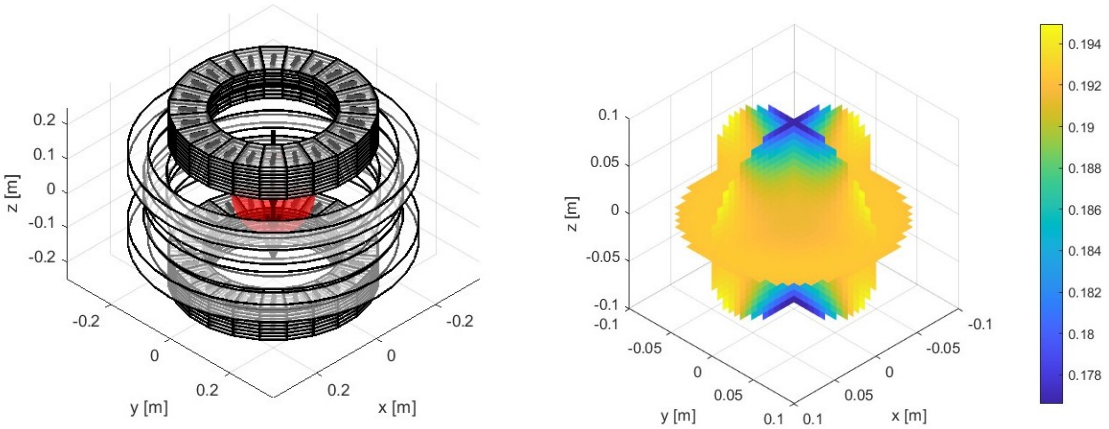

(a) (b)

Figure 5.14: Step 2 - Solution 3.1, structure (a) and PmLab B_z field distribution (b)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.1949	0.1968	+0.9%
$B_{z,\min}$ [T]	0.1784	0.1756	-1.6%
$B_{z,\text{mean}}$ [T]	0.1914	0.1909	-0.3%
Homogeneity [ppm]	86544	109321	+26%
Points in FOV	3544	4736	—

Table 5.10: Solution 3.1 - Comparison between PmLab and FEMM results

Solution 3.2

$$R_{\text{in}} = [0.316,\ 0.281,\ 0.281,\ 0.316]\, m$$

$$z_{\text{pos}} = [-0.105,\ -0.031,\ 0.031,\ 0.105]\, m$$

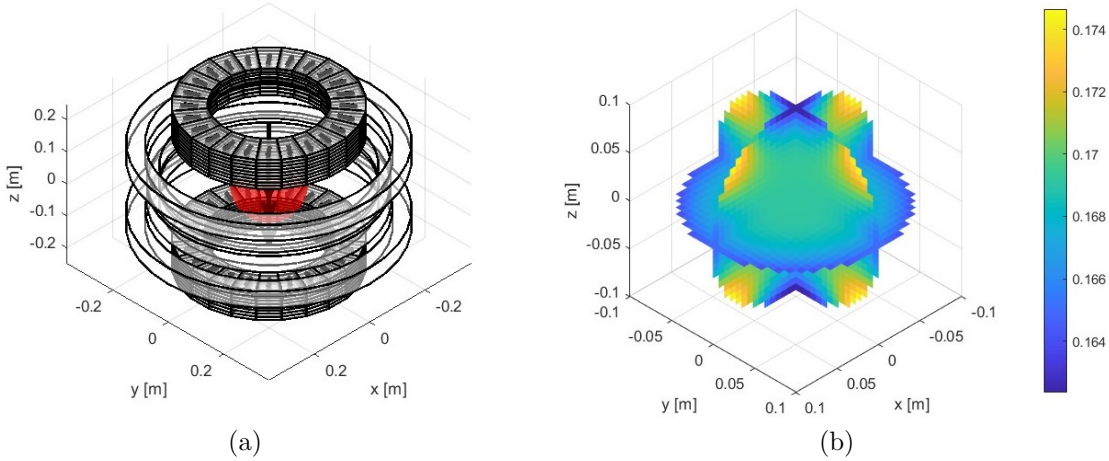

(a) (b)

Figure 5.15: Step 2 - Solution 3.2, structure (a) and PmLab B_z field distribution (b)

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.1744	0.1761	+0.9%
$B_{z,\min}$ [T]	0.1636	0.1589	-2.7%
$B_{z,\text{mean}}$ [T]	0.1692	0.1698	+0.3%
Homogeneity [ppm]	64047	101153	+57%
Points in FOV	3544	4736	—

Table 5.11: Solution 3.2 - Comparison between PmLab and FEMM results

The results reveal a clear and somewhat surprising trend: adding external rings consistently improves field homogeneity in each case examined. By reinforcing the magnetic field in regions that have become weakened, particularly near the central xy-plane, the rings effectively compensate for the low-field zones introduced by the increased spacing between layers. This effect becomes particularly evident and meaningful in the case of Solution 1, as shown in Figure 5.16.

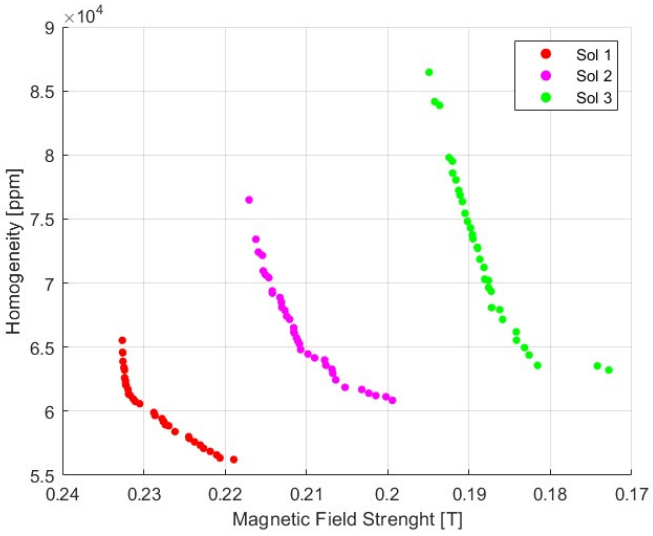

Figure 5.16: Comparison of Pareto fronts obtained from the three initial configurations.

In Figure 5.17, the same Pareto fronts previously discussed are shown, now with their starting configurations, added for comparison (marked as square points). This visualization clearly highlights the drastic reduction in magnetic field inhomogeneity achieved through the addition of external rings.

The improvement is particularly evident when looking at the results derived from Solution 1 (in red). In this case, the initial configuration presented a homogeneity of around $2.3 \times 10^5\ ppm$ while the optimized solutions reach average values around $6 \times 10^4\ ppm$, corresponding to a reduction of roughly 74%.

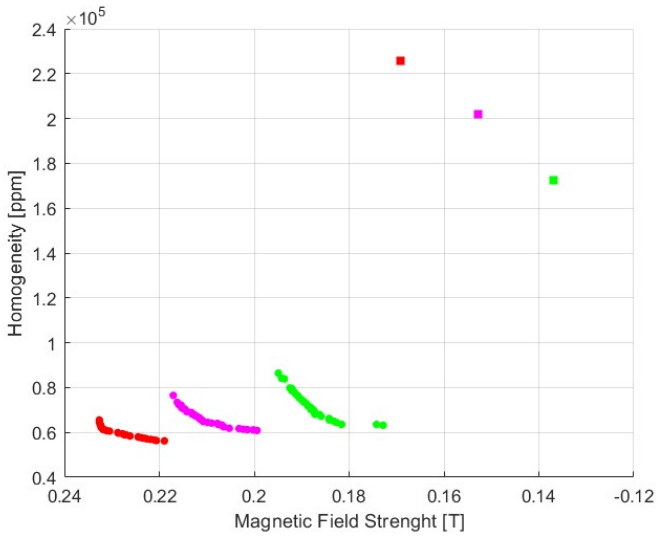

Figure 5.17: Pareto fronts obtained from the three initial configurations and their starting points .

It is exactly because of their strong positive response to the external magnetic rings that Solution 1.1 has been chosen for the next optimization step. As in the previous chapter, this new phase specifically targets the improvement of field homogeneity.

5.3 Step 3a - Field-attenuating Rings

The analysis proceeds with a single-objective optimization, similar to that carried out in Section 4.5, aimed at minimizing the magnetic field magnitude variation within the FOV. This is achieved through the insertion of two pairs of magnetic rings, magnetized inward and outward, but oriented in the opposite direction to the main magnet assemblies.
Differently from what was done in Chapter 4, in this case two parallel optimization processes are conducted. One explores asymmetric pairs of rings, where each ring can vary freely in shape and position, while the other focuses on symmetric configurations, assuming that the rings are mirrored across the central xy-plane.

The rationale behind performing both a symmetric and an asymmetric optimization is illustrated in Figure 5.18.

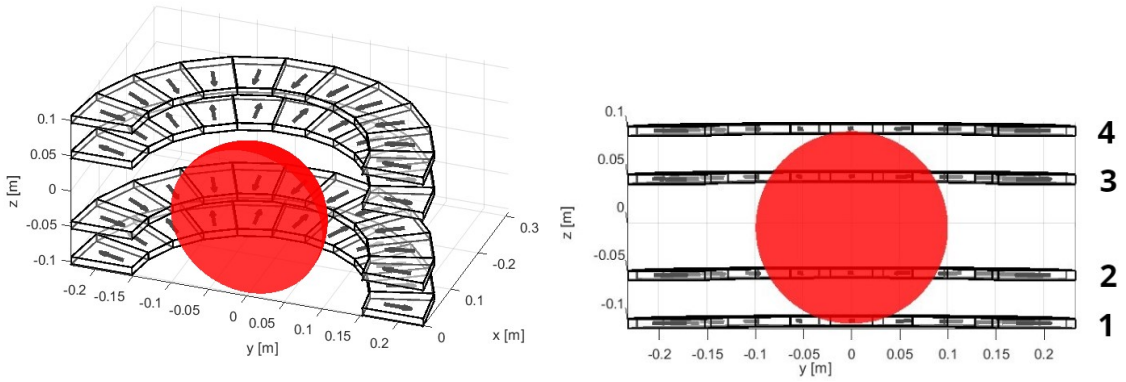

Figure 5.18: Cross-sectional view of the magnet structure with the four auxiliary rings

In this representation, it is assumed that the magnetic field is oriented along the z-axis, from bottom to top. To achieve a field-attenuating effect, rings 2 and 4 are magnetized inward, while rings 1 and 3 are magnetized outward (in the direction opposite to the main magnet assemblies).
It can be observed that each pair of rings sharing the same magnetization direction, rings 1 and 3, and rings 2 and 4, acts on different regions of the FOV. For instance, ring 3 influences a field portion located closer to the center, which has a wider cross-sectional area than the portion influenced by ring 1, which lies near the edge of the FOV. As a consequence, even though the rings are paired by magnetization direction, they operate under different geometric and field conditions.
This is the primary reason why an asymmetric analysis may be more appropriate in this context, allowing each ring to adapt more precisely to the specific region of the FOV it influences. However, given that the main structure is, up to this point, perfectly symmetric with respect to the central xy-plane, a symmetric ring configuration may be preferred for practical reasons. Therefore, both approaches will be carried out and compared, in order to assess their relative effectiveness in improving field uniformity.

The optimization process follows the same setup as described in Section 4.5, maintaining the same number of population elements, number of generations, optimization variables, and parameter limits. The results of both the symmetric and asymmetric configurations are presented below.

5.3.1 Asymmetric Rings Configurations

Ring #	R_{in} $[m]$	t_r $[m]$	$z_{pos}[m]$	$t_z[m]$
1	0.163	0.038	-0.126	0.02
2	0.166	0.017	-0.047	0.01
3	0.153	0.041	0.070	0.01
4	0.162	0.077	0.0124	0.02

Table 5.12: Solution 1.1 - Optimized parameters for the attenuating rings

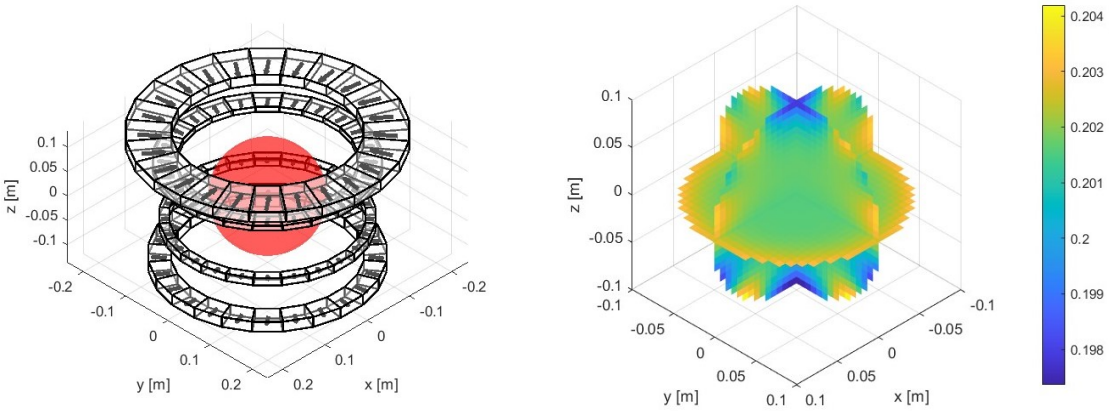

(a) (b)

Figure 5.19: Step 3 - Solution 1.1, structure (a) and PmLab B_z field distribution (b)

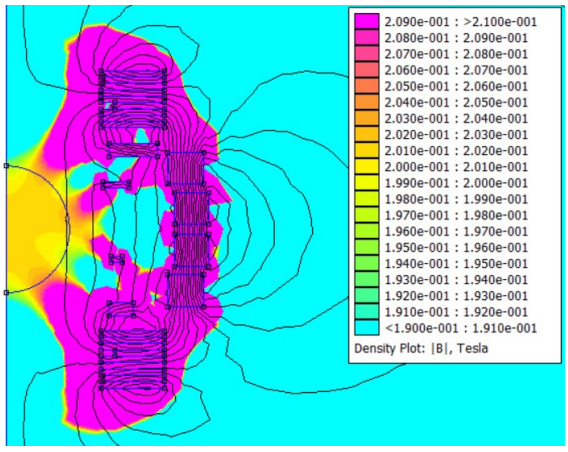

Figure 5.20: Step 3 - Solution 1.1, FEMM B_z distibution

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.2042	0.2047	+0.2%
$B_{z,\min}$ [T]	0.1974	0.1956	-0.9%
$B_{z,\text{mean}}$ [T]	0.2016	0.2011	-0.3%
Homogeneity [ppm]	33880	45625	+34%
Points in FOV	3544	4736	—

Table 5.13: Solution 1.1 - Comparison between PmLab and FEMM results

5.3.2 Symmetric Rings Configurations

Ring #	R_{in} $[m]$	t_r $[m]$	$z_{pos}[m]$	$t_z[m]$
1	0.151	0.089	-0.139	0.02
2	0.159	0.028	-0.069	0.01
3	0.159	0.028	0.069	0.01
4	0.151	0.089	0.139	0.02

Table 5.14: Solution 1.1 - Optimized parameters for the attenuating rings

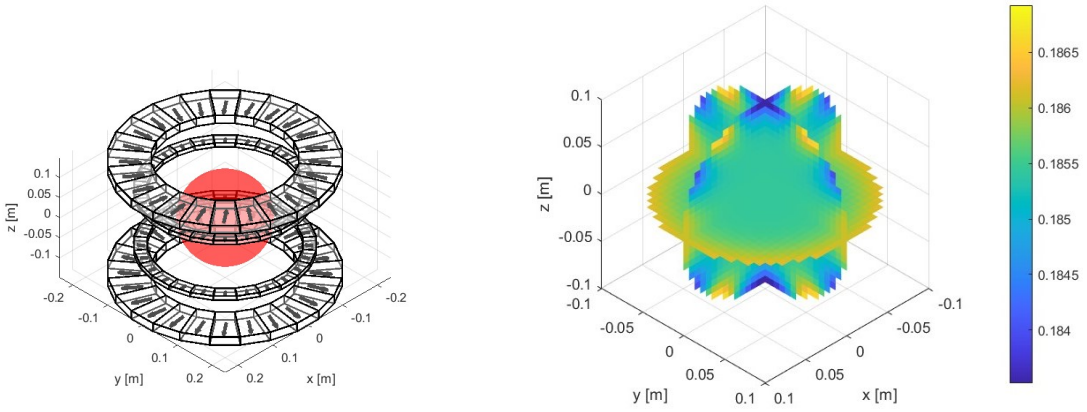

(a) (b)

Figure 5.21: Step 3 - Solution 1.1, structure (a) and PmLab B_z field distribution (b)

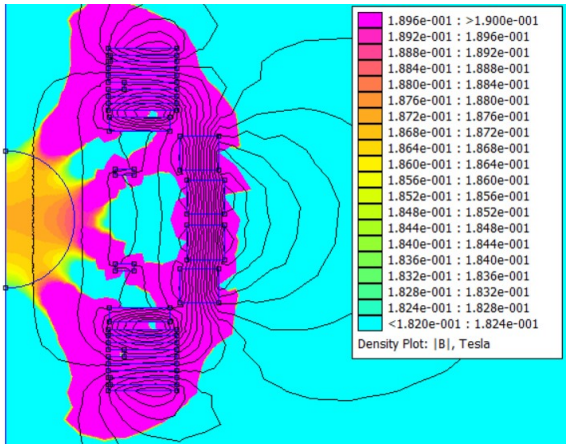

Figure 5.22: Step 3 - Solution 1.1, FEMM B_z distibution

Metric	**PmLab**	**FEMM**	**Difference**
$B_{z,\max}$ [T]	0.1865	0.1873	+0.4%
$B_{z,\min}$ [T]	0.1839	0.1828	-0.6%
$B_{z,\text{mean}}$ [T]	0.1854	0.1869	-0.8%
Homogeneity [ppm]	14352	24077	+67%
Points in FOV	3544	4736	—

Table 5.15: Solution 1.1 - Comparison between PmLab and FEMM results

As shown in Tables 5.13 and 5.15, the introduction of field-attenuating rings, whether symmetric or asymmetric, leads to significantly different outcomes, highlighting once again the intrinsic trade-off that governs magnetic field optimization. The implementation of asymmetric rings results in a reduction of field inhomogeneity within the FOV by approximately 50%, a substantial improvement in uniformity. However, this enhancement comes at the cost of a $\sim$ 11% decrease in the average magnetic field intensity.

In contrast, the use of symmetric rings yields a more pronounced improvement in homogeneity, with inhomogeneity levels reduced by up to 75%. Unfortunately, this comes with a more significant penalty in terms of magnetic field strength, resulting in an average field reduction of around 18%.

These results clearly illustrate the delicate balance between homogeneity and field intensity, emphasizing that optimizing one often requires a sacrifice in the other. The choice between symmetric and asymmetric ring configurations should therefore be guided by the specific application requirements and the acceptable tolerance in field strength degradation.

5.4 Step 3*b* - Ferromagnetic Profile Design

An alternative strategy is now explored, starting from the optimized Solution 1.1 obtained in Section 5.2. While the objective remains the same, improving field homogeneity within the FOV, this new approach replaces the external rings with a ferromagnetic structure designed to smooth out field variations through passive shaping.
The concept is to introduce a ferromagnetic component placed between the two magnet assemblies, designed to enclose the FOV. The internal boundary of this structure is defined by a continuous function $x = f(y)$, which shapes the profile in such a way as to guide and redistribute the magnetic flux.
The structures have been designed following these key steps:

- The function domain $y \in [0, 0.15]$ is discretized into multiple points where the profile function is evaluated using a polynomial expression. The resulting curve is then mirrored with respect to the central xy-plane to obtain a fully symmetric profile;
- A constraint is applied to limit the maximum penetration of the material, in order to avoid interference with the FOV and preserve the usable imaging volume;
- A further constraint is applied to the external boundary of the structure: the profile described by the function is not allowed to fall below a minimum radial coordinate, set to $x = 0.5\,cm$. If the computed profile violates this constraint at any point (i.e., if $x(y) < 0.5\,cm$), the function is clipped and forced to coincide with this lower bound. This condition is introduced to preserve the structural continuity and the mechanical integrity of the ferromagnetic component throughout its entire shape.

This type of analysis can only be carried out using the FEMM environment, as the current version of PmLab does not support the inclusion of ferromagnetic materials. For this reason, all simulations involving the shaped ferromagnetic structure will be performed in FEMM.
The starting point for this study consists in setting a fixed initial value of magnetic relative permeability equal to $\mu_r = 200$. This value is chosen as a realistic and practical estimate, consistent with the magnetic properties of commonly available ferromagnetic materials such as certain steels and soft iron alloys. Using this fixed property, different profile functions will be tested in order to evaluate their potential

effectiveness.
Due to the significantly higher computational cost associated with FEMM simulations, the number of generations in the optimization process has been reduced to 20. To compensate, the population size has been increased to 200, allowing for a broader exploration of the solution space. This strategy prioritizes diversity in the early stages of the search. The most promising solutions identified through this initial phase will then be further refined using a dedicated post-processing methodology.

Case 1a – Polynomial Function

In this first scenario, the internal boundary of the ferromagnetic structure is defined by a fourth-degree polynomial function of the form:

$$x(y) = a_4 y^4 + a_3 y^3 + a_2 y^2 + a_1 y + a_0$$

where the optimization variables are the five polynomial coefficients:

$$\text{coeff} = [a_4\,, a_3,\ a_2,\ a_1,\ a_0]$$

After completing the optimization cycle, the obtained results are:

$$\text{coeff} = [416.05\,, -97.419\,, -4.977\,, 0.017\,, 0.008]$$

These coefficients define the internal shape of the ferromagnetic structure, as shown in Figure 5.23. The computational results obtained using the FEMM environment for this optimized ferromagnetic structure are reported in Table 5.16, where these values are directly compared with those corresponding to the reference configuration presented in Section 5.3.2, which features symmetric magnetic rings to shape the field.

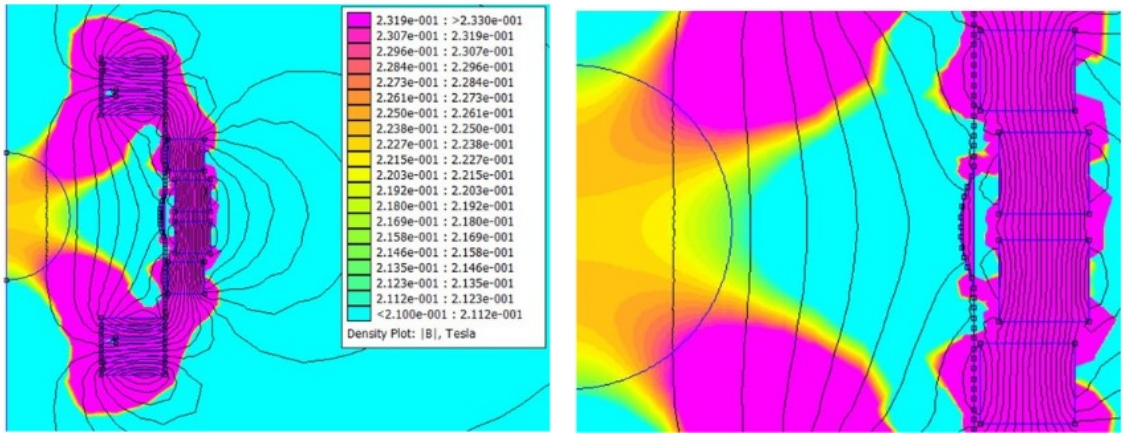

(a) (b)

Figure 5.23: Case 1 - Optimized ferromagnetic profile: full structure(a), zoomed view(b)

Parameter	**Value**	**Variation**
$B_{z,\max}$ [T]	0.2326 T	+24.7%
$B_{z,\min}$ [T]	0.2130 T	+15.8%
$B_{z,\text{mean}}$[T]	0.2236 T	+20.6%
Homogeneity [ppm]	87,635	+153%

Table 5.16: Case 1 - Magnetic field values and relative variation with Sec.5.3.2

As clear from both the magnetic field map and the numerical data, the interposed ferromagnetic structure placed between the two magnet layers proves to be largely ineffective in mitigating field inhomogeneities. No significant improvement has been made compared to the configuration with symmetric containment rings. A slightly improved performance can be noticed with respect to the solution employing only the external rings (see Section 5.2.1, Table 5.7). It can be attributed to the tendency of the ferromagnetic material to attract some magnetic field lines toward the central region. However, this influence is minimal and does not result in a meaningful enhancement of field homogeneity.

Case 2a - Polynomial Function for Both Profiles

In this configuration, both the internal and external boundaries of the ferromagnetic structure are described by two independent fourth-degree polynomial functions:

$$x_{\text{int}}(y) = a_4y^4 + a_3y^3 + a_2y^2 + a_1y + a_0$$
$$x_{\text{ext}}(y) = b_4y^4 + b_3y^3 + b_2y^2 + b_1y + b_0$$

The optimization variables in this case include the full set of polynomial coefficients for both profiles:

$$\text{coeff}_{\text{int}} = [a_4,\, a_3,\, a_2,\, a_1,\, a_0]$$

$$\text{coeff}_{\text{ext}} = [b_4,\, b_3,\, b_2,\, b_1,\, b_0]$$

Once the optimization is completed, the identified solutions are:

$$\text{coeff}_{\text{int}} = [115.16\,, -54.4\,, -6.266\,, -0.714\,, 0.034]$$

$$\text{coeff}_{\text{ext}} = [-493\,, 22.49\,, -7.499\,, -0.656\,, 0.023]$$

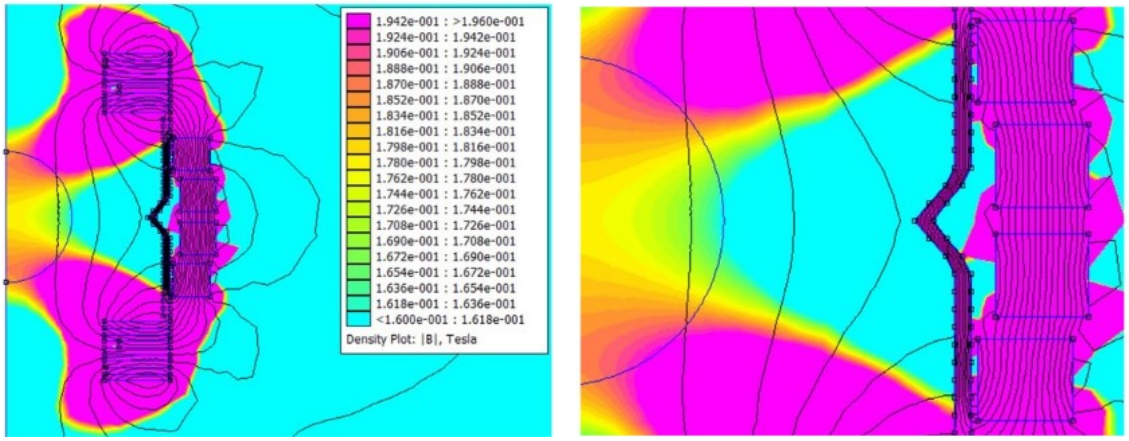

(a) (b)

Figure 5.24: Case 2 - Optimized ferromagnetic profile: full structure(a), zoomed view(b)

Parameter	**Value**	**Variation**
$B_{z,\text{max}}$ [T]	0.1950 T	+3.2%
$B_{z,\text{min}}$ [T]	0.1606 T	-12%
$B_{z,\text{mean}}$ [T]	0.1813 T	-3%
Homogeneity [ppm]	190070	+450%

Table 5.17: Case 2 - Magnetic field values and relative variation with Sec.5.3.2

In this case, as indicated in Table 5.17, the overall performance is significantly worse than in the previous configuration. This is primarily due to the high magnetic permeability of the ferromagnetic structure, which attracts and diverts the magnetic field lines produced by the external boosting rings. Instead of reinforcing the field in the central region, the flux is redirected toward the boundaries. As a result, the boosting rings are no longer able to effectively enhance and homogenize

the magnetic field within the FOV. Instead, their beneficial influence is suppressed, leading to both a reduction in field intensity and a substantial increase in inhomogeneity.

This observation suggests that a critical limitation of these solutions lies in the assumption of a fixed relative magnetic permeability μ_r for the ferromagnetic material. By prescribing $\mu_r = cost$, the optimization process is inherently constrained and driven toward configurations that may be structurally valid but ultimately ineffective in achieving the desired magnetic field. To overcome this issue, from now on the relative permeability μ_r will be treated as an additional optimization variable.

Case 1b - Polynomial Function

The function that describes the internal profile is the same of the previous case but, as anticipated, the set of optimization variables includes a new value:

$$\text{coeff} = [a_4\,, a_3,\ a_2,\ a_1,\ a_0,\ \mu_r]$$

The results of the optimization carried are:

$$\text{coeff} = [450.68,\ -88.19,\ 9.50,\ -1.21,\ 0.209,\ 11]$$

As can be observed, the relative magnetic permeability obtained from the optimization process is significantly lower than the fixed value used in the previous cases. This reduction of μ_r has a crucial impact: rather than strongly altering or concentrating the magnetic field, the ferromagnetic structure instead acts as a flux guide, as illustrated in Figure 5.25. Although the resulting value of $\mu_r = 11$ may appear low, it still corresponds to commercially available materials. For instance, NiZn ferrites typically have μ_r values between 10 and 15, and low-carbon manganese steels or non-oriented silicon steels conditions can also exhibit relative permeabilities in the range of 10–50.

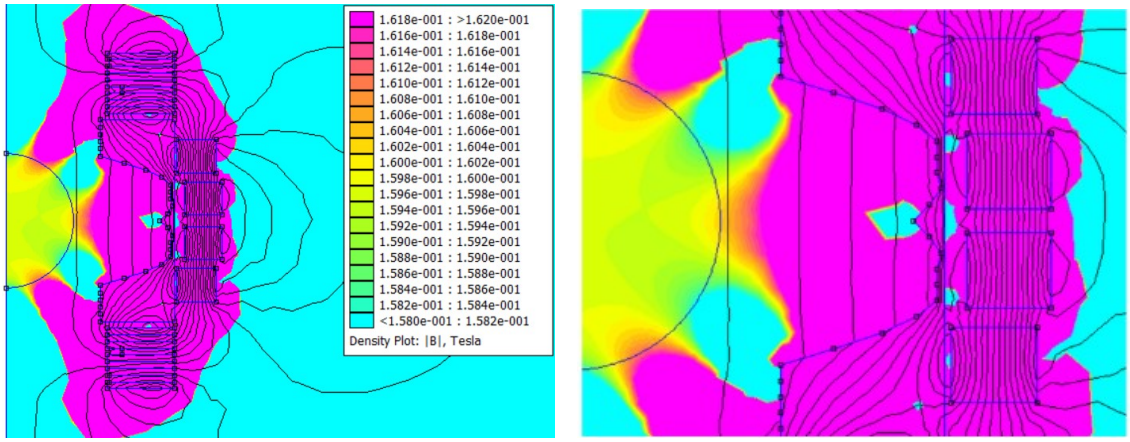

(a) (b)

Figure 5.25: Case 1b - Optimized ferromagnetic profile: full structure(a), zoomed view(b)

As one might expect, the optimized profile is shaped to redirect magnetic field lines toward the central region in such a way that the outer portions of the structure are thicker and penetrate deeper into the magnet assembly, in order to intercept and guide the magnetic flux. Toward the center, instead, the profile progressively narrows, leaving space for the field to redistribute more uniformly.
This trend toward improved field uniformity is even more evident when examining the numerical results reported in the Table 5.18.

Parameter	**Value**	**Variation**
$B_{z,\text{max}}$ [T]	0.1614 T	-14.5%
$B_{z,\text{min}}$ [T]	0.1581 T	-13.4%
$B_{z,\text{mean}}$ [T]	0.1597 T	-14.5%
Homogeneity [ppm]	20750	-15%

Table 5.18: Case 2 - Magnetic field values and relative variation with Sec.5.3.2

However, this improvement comes at a cost: a noticeable reduction in magnetic field strength, with the mean B_z value decreasing by approximately 14.5%. This loss in performance is a direct consequence of the magnetic flux behavior within the structure. A considerable portion of the field lines, attracted by the ferromagnetic material, is diverted away from the original central path.
Ultimately, this highlights the intrinsic nature of the problem as a trade-off: there is no universally optimal solution, but rather a range of configurations that may or may not align with the initial design priorities.

Case 2a - Polynomial Function for Both Profiles

The previous studied case, including both profiles, is here repeated, including μ_r an additional variable:

$$\text{coeff} = [a_4,\ a_3,\ a_2,\ a_1,\ a_0\,, b_4,\ b_3,\ b_2,\ b_1,\ b_0\,, \mu_r]$$

The process exhibits the following results:

$$\text{coeff} = [221.02,\ 68.79,\ 1.796,\ -1.211,\ 0.101,\ -326.43,\ 58.6,\ -1.344,\ -1.021\,, 0.089\,, 10.46]$$

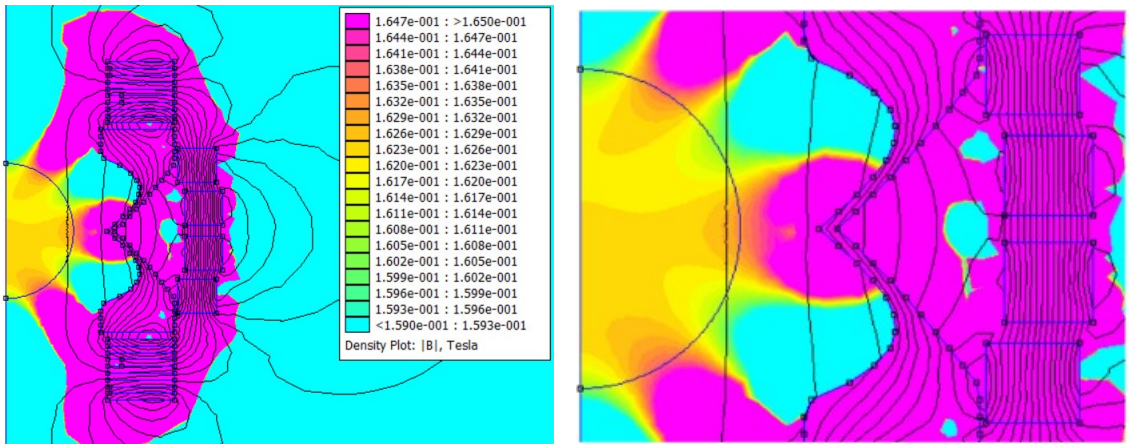

(a) (b)

Figure 5.26: Case 2b - Optimized ferromagnetic profile: full structure(a), zoomed view(b)

Parameter	**Value**	**Variation**
$B_{z,\text{max}}$ [T]	0.1652 T	-12.6%
$B_{z,\text{min}}$[T]	0.1608 T	-11.8%
$B_{z,\text{mean}}$[T]	0.1636 T	-12.4%
Homogeneity [ppm]	26356	+9%

Table 5.19: Case 2b - Magnetic field values and relative variation with Sec.5.3.2

Similar results are obtained when both the internal and external ferromagnetic profiles are included as variables in the optimization process. The optimized shape shares notable similarities with the profile obtained in *Case 1b*, especially in the outer regions, where the geometry is designed to channel magnetic flux toward the central area of the FOV.
However, the internal region of the profile differs slightly, as it extends more deeply

toward the xy-plane, forming a sharper inward indentation. While this geometry makes theoretical sense from a field-guiding perspective, it is clearly more challenging to implement from a practical manufacturing standpoint. In this case, the outcome does not meet the expected goals: modifying both profiles simultaneously leads not only to a general reduction in magnetic field strength, but also to a slight increase in field inhomogeneity. As a result, this configuration is considered suboptimal and is therefore discarded from further consideration.

Case 3 - Sigmoidal Functions

In addition to polynomial representations, other classes of functions were also explored to define the ferromagnetic profiles. Among these, the sinusoidal function was tested for its natural smoothness and symmetry. However, in practical terms, its effectiveness strongly depends on the application of upper and lower geometric bounds, which tend to distort the pure sinusoidal shape, reducing its overall efficiency and making the final design less intuitive.
A more promising alternative was found in the use of sigmoidal functions, which proved to be well suited for this kind of magnetic design. A typical sigmoidal function used in this context takes the form:

$$x(y) = A \cdot \left(\frac{1}{1 + e^{-B(y-C)}} \right)$$

where:

- A defines the amplitude (maximum extension);
- B controls the slope or steepness of the transition;
- C sets the inflection point (center of the transition zone).

This type of function naturally produces a smooth, progressive transition from one region to another with an S shaped curve first starting with a flat profile, then rising gradually, and eventually flattening again as it approaches its maximum value A.

In this case the optimization parameters and results are the following:

$$\text{coeff} = [A,\ B,\ C,\ \mu_r]$$

$$\text{coeff} = [0.1142,\ 174.29,\ 0.0834,\ 13.31]$$

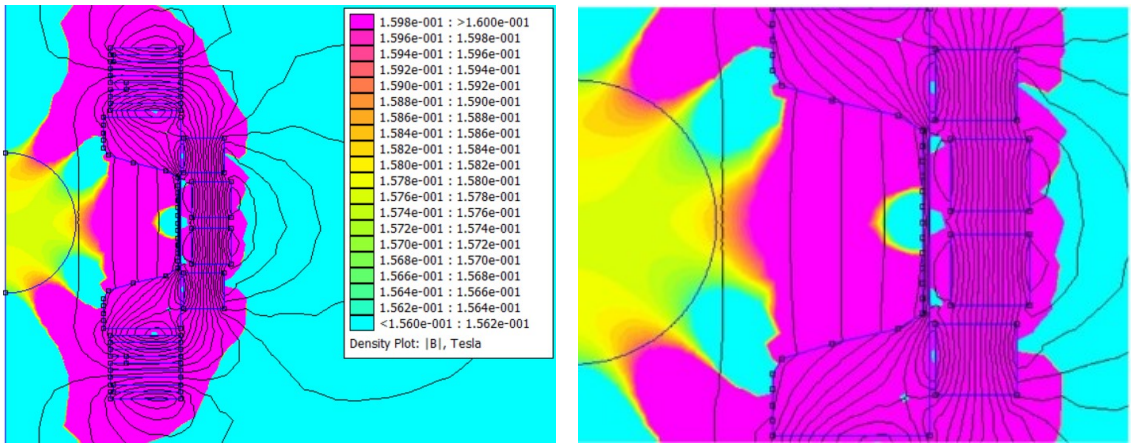

(a) (b)

Figure 5.27: Case 3 - Optimized ferromagnetic profile: full structure(a), zoomed view(b)

Parameter	**Value**	**Variation**
$B_{z,\max}$ [T]	0.1597 T	-15.5%
$B_{z,\min}$ [T]	0.1564 T	-14.3%
$B_{z,\text{mean}}$ [T]	0.1578 T	-15.5%
Homogeneity [ppm]	20812	-14%

Table 5.20: Case 3 - Magnetic field values and relative variation with Sec.5.3.2

As shown in the images, the optimized structure exhibits the same main features observed in the previous examples, particularly in Case 1b. However the profiles obtained through the sigmoidal function are more concave and gradual, facilitating the manufacturability of the piece.
In the end the resulting effects on the magnetic field remain very similar to those already obtained, both in terms of field homogeneity and field strength. This consistency strongly supports the convergence of the optimization process towards a class of effective and robust solutions, regardless of the specific functional form used to define the profile.

5.5 From ga to fmincon

As anticipated in Section 5.4, the process used here is structured as a two-step strategy combining global and local search techniques.
The first stage employs the genetic algorithm (*ga*), a stochastic global optimization method that does not require derivative information. This algorithm is particularly effective for exploring complex, high-dimensional spaces, where traditional gradient-

based methods might fail due to the presence of multiple local minima. As a result, the genetic algorithm is an effective tool for identifying promising regions within the solution space.
However, due to its stochastic nature, the genetic algorithm may produce solutions that are globally promising but lack of precision at the local level. To solve this limitation, a second step is added using MATLAB's `fmincon` function. This optimizer, deterministic and based on gradients (meaning it works best when the objective function is smooth and differentiable), it's especially effective at improving solutions within already promising areas found during the global search.
In this framework, the solutions produced by the genetic algorithm are passed directly to fmincon, which uses them as initial baseline. Thanks to this setup, fmincon can begin its local search in a region that is already promising, which helps speed up convergence and improves the final result. The result is a hybrid optimization strategy that combines the exploratory capabilities of the genetic algorithm with the precision of fmincon, bringing the solution as close as possible to the global minimum.
The table below reports some of the most remarkable the outcomes of the simulations carried out using the fmincon function. The results are sorted in descending order of field homogeneity, allowing for a clear comparison of the different configurations.

	$a4$	$a3$	$a2$	$a1$	$a0$	μ_r	**Hom [ppm]**
Run 1	971.7	11.14	11.48	-0.5227	0.0176	10.1	**25604**
Run 2	-35.39	-0.3784	95.68	-8.282	0.0152	16.4	**20687**
Run 3	-0.8802	358.152	61.09	-7.891	0.0959	12.6	**20610**
Run 4	-49.67	-42.88	42.18	-2.434	0.0133	10.8	**19849**
Run 5	802.3	341.2	27.28	-5.662	0.0313	17.7	**19804**
Run 6	622.2	28.06	92.148	-8.301	0.0962	13.2	**19773**
Run 7	-574.3	59.52	474.2	-42.66	0.0412	17.9	**19499**
Run 8	35.42	288.8	0.7790	-1.409	0.0230	11.0	**19286**
Run 9	526.1	314.6	44.14	-6.656	0.0103	18.4	**19276**
Run 10	1382	426.6	130.8	-15.79	0.0089	18.4	**19258**
Run 11	198.66	-969.3	473.7	-33.92	0.0097	18.3	**19218**
Run 12	515.7	405.4	52.01	-7.820	0.0059	18.5	**19162**
Run 13	1202	-940.9	392.7	-27.58	0.0052	18.6	**19105**
Run 14	75.30	300.4	46.03	-5.995	0.0708	13.1	**18855**
Run 15	797.7	724.0	456.6	50.28	0.0322	19.3	**17514**
Run 16	-1272	-689.3	424.3	-31.47	0.0061	19.9	**17287**
Run 17	751.9	380.2	300.4	-23.05	0.0241	15.6	**15804**
Run 18	1059	153.4	356.5	-32.48	0.00698	15.3	**15708**

Table 5.21: fmincon results

An interesting observation emerging from the table is the clustering of multiple runs (notably Run 6 through Run 14) around a local minimum in homogeneity, with values consistently falling in the range of 19100–19400 ppm. This phenomenon reflects the nature of the fmincon algorithm, which performs a local gradient-based search starting from the initial point provided by the genetic algorithm.

Given the limited exploration radius of fmincon, it tends to converge toward the nearest local minimum, especially when multiple initial conditions lie within the same basin of attraction. This clustering suggests that a shallow minimum exists around 19200 ppm, attracting several optimization paths. However, Run 15 through Run 18 demonstrate a clear descent toward significantly lower homogeneity values, reaching as low as 15700 ppm. This implies that those particular initial conditions were closer to a deeper local minimum, which fmincon was able to reach and exploit effectively.

To better understand the behavior of the optimized solutions and to visually highlight the differences among them, a selection of representative configurations is presented.

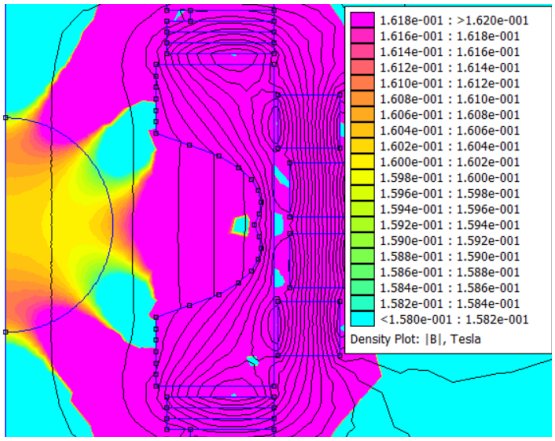

Figure 5.28: Run 1, final result

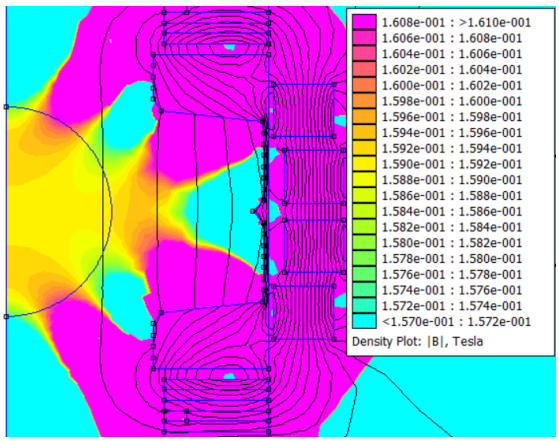

Figure 5.29: Run 2, final result

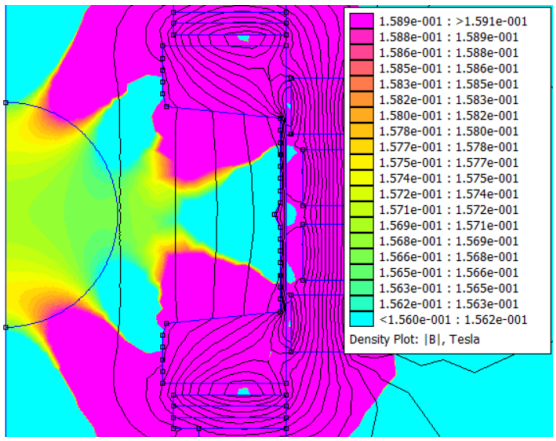

Figure 5.30: Run 9, final result

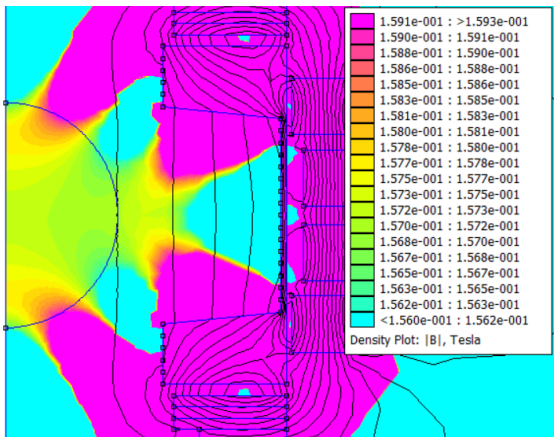

Figure 5.31: Run 13, final result

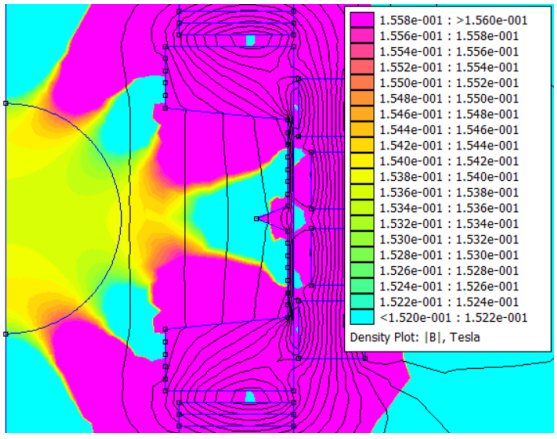

Figure 5.32: Run 16, final result

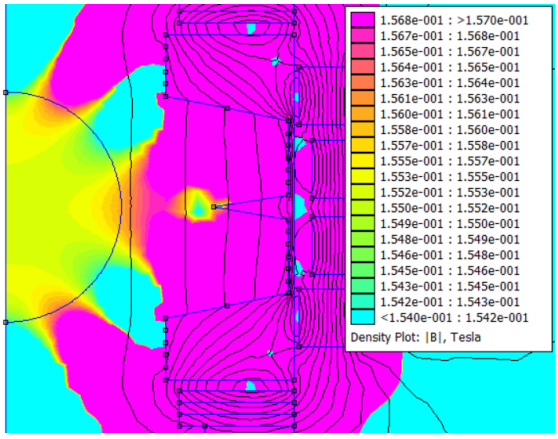

Figure 5.33: Run 19, final result

From a geometrical point of view, the progression across the selected configurations reveals a clear transformation of the ferromagnetic profile as the optimization

goes on.

In the initial solution (Run 1), the structure appears compact with a relatively smooth and rounded contour. As we move toward Run 9 and Run 13, the structure gradually becomes more concave, particularly along the central region, where the profile curves inward more noticeably. At the same time, the overall contour becomes less smooth and more angular, a direct consequence of the discretization adopted for the optimization process. In the final configurations (Run 16 and Run 19), this trend is further accentuated. The profile exhibits sharper transitions and increased irregularity, and most notably, a central protrusion begins to emerge. This feature becomes progressively more pronounced, and it serves a critical functional role: to increase the magnetic field intensity in the central area of the FOV, which is typically the most difficult to saturate magnetically.

It should be noted that all profiles exhibit a certain degree of angularity, due to the fact that they were constructed by interpolating only 15 discrete points along the profile curve. This simplification was adopted to reduce computational burden, given the high cost of repeated simulations in FEMM. As a result, the profiles show polygonal edges and sharp transitions, which could be smoothed in future iterations by refining the point density or applying spline-based reconstruction.

6 Comparative Analysis

One of the final steps before concluding this work is to conduct a quantitative comparison between the results obtained in the previous chapter and those presented in existing literature.
Throughout the thesis, no strict numerical targets were imposed for either magnetic field strength or homogeneity. Rather, the approach has been exploratory, the guiding principle has simply been: the stronger the field and the lower the inhomogeneities, the better the outcome.
However, evaluating the effectiveness of the proposed designs requires a point of reference. The most meaningful way to assess the significance of the results is by comparing them with those obtained in previous research aimed at the same goal. This comparison not only allows for a clearer understanding of the performance level achieved, but also serves to highlight the strengths and potential limitations of the methods adopted in this thesis.
It is worth noting that the specific topic addressed in this thesis is relatively underexplored in the current literature and, within this niche, further distinctions arise based on the type of magnetic technology employed: some studies focus on resistive coil systems, others on Halbach arrays. Moreover, many studies target very small FOVs, often only a few centimeters wide, which makes direct comparison difficult due to the inherently different design constraints and performance criteria of larger-scale configurations like those considered in this thesis.

Among the available works, one reference study has been identified that aligns closely with the methodology and objectives pursued here. The configuration selected as a benchmark for comparison is based on the study by Ren, Mu, and Huang [5], which had already been used as a reference model during the development of Configuration 5. Its compatibility in terms of geometry, field characteristics, and design philosophy makes it a natural and meaningful benchmark for evaluating the performance of the proposed solutions in this thesis. The initial structure used in the [5] is identical to that described in Section 3.3.2 and consists of a ring-pair

permanent magnet array composed of multiple Aubert ring pairs aligned along the z-axis. Each layer includes an external ring with a radius of 250 mm, and the system is arranged in 10 stacked layers, each with a thickness of 10 mm. The two magnetic assemblies are separated by a 200 mm air gap, and all magnets are uniformly magnetized with a remanence of 1.4 T (see Figure 3.18).
The main geometric difference between this work and the one presented in the thesis lies in the shape and volume of the FOV. While the reference work considers a cylindrical FOV with a diameter of 200 mm and a height of 50 mm located at the center of the magnetic bore, the thesis adopts a different FOV shape, which consequently influences the optimization strategy and performance.
As in the present work, the study by Ren et al. employs a genetic optimization strategy to identify the optimal magnet configuration while for the evaluation of the magnetic field within the FOV, a current model is adopted. The algorithm begins by generating an initial population of N candidate solutions, each defined by a set of 11 optimization variables: the internal radii of the 10 inner rings and the inter-layer distance (d) between the two magnet assemblies. These variables are then progressively evolved through the genetic algorithm in a multi-objective framework, with the goal of identifying configurations that simultaneously optimize magnetic field strength and field homogeneity within the defined FOV.

The outcome of the optimization process conducted in the article is presented in Figure 6.1, which also illustrates the resulting internal ferromagnetic profile derived from the final configuration.

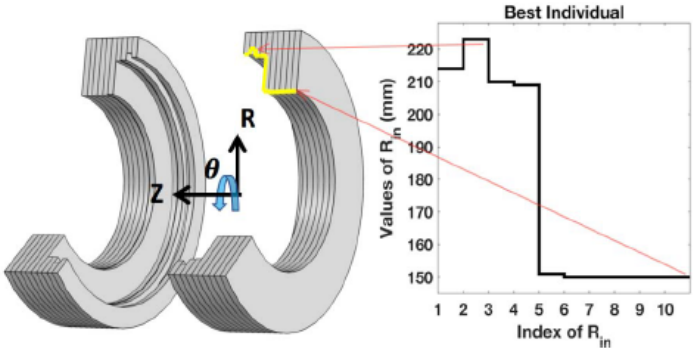

Figure 6.1: Optimized configuration and internal ferromagnetic profile as presented in [5]

$$R_{\text{in}} = [214,\ 223,\ 210,\ 209,\ 151,\ 150,\ 150,\ 150,\ 150,\ 150]\, mm$$

$$d = 200\, mm$$

According to the results reported in the study, the optimized structure achieves a field homogeneity of 24,786 ppm and an average magnetic field strength of 169.7 mT in the cylindrical FOV.
However, in order make the design physically realizable, the structure has to be discretized into individual magnet circular segments. Specifically, the configuration shown in Figure 6.1 is segmented into 12 fan-shaped magnets, each uniformly magnetized. This fragmentation process, necessary for manufacturing feasibility, introduces discontinuities that degrade performance: as a result, the field homogeneity worsens to 32,511 ppm, while the average magnetic field decreases slightly to 167.6 mT.

When comparing the results obtained in this thesis with those reported by Ren et al., it becomes evident that, in terms of magnetic field strength, the performances are broadly comparable.
In fact, if we consider the results presented in Section 5.3.2, relatively to the configuration with symmetric compensating rings (see Table 5.15), the proposed design in this thesis demonstrates that it is possible to achieve even higher average magnetic field values, reaching approximately 185 mT, while maintaining lower field inhomogeneities (24007 ppm as computed via FEMM and 14352 ppm in PmLab simulations).
Similarly, in Section 5.4, using an optimized ferromagnetic profile slightly lower average field values are observed (around 158 mT), with inhomogeneities reduced to approximately 15700 ppm in the most refined version.

At first glance, the results obtained in this thesis may appear to outperform those presented in the reference study, both in terms of field strength and homogeneity. However, a strictly quantitative comparison is neither feasible nor entirely appropriate, due to two fundamental considerations. The first is the size and geometry of the FOV, which significantly influence the optimization strategy and ultimately affect the resulting magnetic performance. In this thesis, the FOV is defined as a sphere with a diameter of 200 mm, corresponding to a volume of $V = \frac{4}{3}\pi r^3 = 4186\,mm^3$, whereas the reference work considers a cylindrical FOV with the same diameter but a height of only 50 mm, resulting in a smaller volume of $V = \pi r^2 h = 1570\,mm^3$. This means that the cylindrical FOV has slightly more than one third the volume of the spherical FOV used in the current study. As a result, the optimization strategy adopted in the reference work is fundamentally different from the one required here. In fact, ensuring sufficient field coverage over a smaller cylindrical volume does not

necessarily require any structural modifications to the base magnet configuration.

In contrast, the larger and axially extended FOV considered in the developed configurations demanded two major structural adjustments. First, the two assemblies of magnet layers had to be moved farther apart to achieve adequate coverage throughout the extended volume. This increased separation inevitably led to a significant drop in magnetic field intensity, which was subsequently compensated by the introduction of external "boosting" rings, specifically designed to reinforce the field. In the end, this arrangement led to an unacceptable level of inhomogeneity within the FOV. To address this, an additional set of components was introduced just to enhance field uniformity, without affecting the overall intensity.
Therefore, while the achieved magnetic performance can be considered promising, particularly in light of the fact that field homogeneity was improved despite the enlargement of the FOV, this result came at a cost. Specifically, the successive introduction of additional magnetic and ferromagnetic components led to a progressive increase in structural complexity, material usage, and overall cost. Most importantly, it also resulted in a significant increase in weight, which inevitably compromises one of the original goals of the design: the portability of the device.

As has become clear throughout the course of this thesis, the design of permanent magnet systems is fundamentally a matter of trade-offs. There is no single "correct" configuration; rather, different designs may prove more or less suitable depending on the specific application requirements. Every modification, whether aimed at improving field homogeneity, increasing intensity, or preserving portability, comes with its own set of advantages and limitations. Ultimately, the effectiveness of a given configuration must be evaluated in light of the intended use case, balancing competing objectives such as performance, simplicity, cost, and practicality.

7 Further Developments

The optimization of the static magnetic field to achieve the highest possible strength and homogeneity represents only the initial stage in the design process of an MRI system. While a strong and uniform static field is essential to ensure a good signal-to-noise ratio, it is not sufficient, on its own, to produce spatially resolved images. Indeed even though a well-designed static field is important for getting a strong signal and good image contrast, it does not provide any information about the spatial origin of the signal.
Spatial encoding in MRI is based on the principle that the resonance frequency of nuclear spins is determined by the local magnetic field. By introducing controlled spatial variations in the magnetic field across the sample, it becomes possible to establish a direct correlation between the characteristics of the observed signal and the spatial position of the spins. This spatial modulation of the field can be achieved in two fundamentally different ways.
In conventional high-field MRI systems, spatial encoding is achieved through dedicated gradient coils, which are activated to superimpose linear magnetic field variations along the x, y, and z axes.
Alternatively, particularly in low-field MRI systems based on permanent magnets, spatial encoding can arise intrinsically from the geometry of the magnet structure itself. These systems utilize what are known as Self-Encoding Magnetic Fields (SEM), static yet spatially varying magnetic field distributions that inherently encode spatial information into the signal, eliminating the need for time, varying gradient fields.

Depending on how many spatial dimensions are encoded, MRI acquisition strategies are generally categorized as 1D, 2D, or 3D imaging.
In 1D imaging, the system captures spatial information along just one axis, so the result is a signal profile instead of a full image. This technique is mainly used to check how uniform the magnetic field is or to test how the system performs along a certain direction, rather than for producing diagnostic images.

2D imaging involves spatial encoding along two orthogonal directions, typically within a selected slice. This technique is the most commonly employed method in both clinical and low-field MRI systems. It provides high spatial resolution in a plane, while maintaining reasonable acquisition times and moderate hardware demands.
3D imaging, on the other hand, encodes spatial information across all three dimensions simultaneously. Although this makes it possible to achieve uniform resolution in every dimension and to reconstruct full 3D volumes, it generally requires longer acquisition times and more complex reconstruction algorithms, making it less practical for fast or portable systems.
In this study a 2D imaging strategy turns out to be the most suitable option as it is compatible with both gradient-based encoding and SEM approaches. Starting from that a 3D representation can be achieved by acquiring multiple adjacent slices and combining them, without increasing hardware complexity. This approach enables volumetric imaging while avoiding the technical challenges associated with full 3D spatial encoding.
A series of possible solutions for implementing spatial encoding in the proposed magnetic configuration are now presented. Some of these approaches are well-established and technically more accessible, while others are less conventional but may offer interesting perspectives and opportunities for future research and development.

Gradient Coils

Gradient coils are specially designed conductor windings used to generate controlled, spatially varying magnetic fields that are superimposed on the main static field. Typically, gradient coils are activated sequentially during the imaging sequence, allowing spatial information to be encoded independently along the x, y, and z axes. In 2D imaging, a gradient is first applied along the z-axis during the RF excitation to select a specific slice within the volume (slice selection). Then, a phase encoding gradient is applied along the y-axis, introducing position-dependent phase shifts among the spins. Finally, during signal acquisition, a frequency encoding gradient is applied along the x-axis, so that spins located at different x-positions precess at different frequencies. This combination enables the reconstruction of a two-dimensional image of the selected slice.
An interesting example of gradient coil redesign specifically tailored for low-field MRI systems was presented in [13], where a custom solution adapted to a Halbach array configuration is developed. This work highlights the need to revisit traditional

gradient design principles when dealing with compact, non-conventional magnet geometries.

In the specific magnetic configuration developed here, the static magnetic field generated by the permanent magnets presents symmetry with respect to the xy-plane. As a result, spatial encoding along the z-axis becomes ambiguous, since spin systems located at symmetric positions above and below the central plane experience identical field strengths and thus resonate at the same frequency. Gradient coils along the x and y axes, instead, do not suffer from this limitation and can be designed following conventional procedures typically adopted in low-field MRI systems.
Two potential strategies, which could serve as starting points for future investigations, may be considered to overcome this issue. The first consists of dividing the field of view into two hemispheres, acquiring and reconstructing each half separately. While this approach could potentially mitigate the ambiguity introduced by field symmetry, it also presents limitations in terms of imaging flexibility and may lead to uncertainty around the symmetry plane.
Another interesting solution could be the implementation of two independent z-gradient coils, each operating exclusively on one half of the field of view. This configuration would allow selective control over the upper and lower hemispheres, effectively breaking the field symmetry and enabling unambiguous spatial encoding. By activating one coil at a time, it becomes possible to directly associate the acquired signal with a specific region of the semi-FOV, simplifying the reconstruction process and reducing encoding ambiguity.

Irregular Shaped I/O ring pair

An alternative solution to the use of conventional gradient coils along the x and y axes is proposed in the study by Ren et al. in [6], which presents a further development of the magnetic structure originally designed in [5].
The study presents the design and development of an irregular-shaped inward–outward ring-pair magnet array, capable of generating a 1D monotonic field pattern for 2D head imaging in low-field portable systems. In this configuration, the magnetic layers are discretized into circular segments with variable outer radii, which are determined through an optimization process. As shown in Figure 7.1, the inner radii of the tiles are fixed, while the outer radii are allowed to vary within each ring, but constrained to be equal across different rings. The position of the FOV is also reconsidered and modified, being shifted a few millimeters in the $-x$ direc-

tion. The optimization process is guided by three main objectives: maximizing the field strength, achieving controlled field inhomogeneity, and producing a monotonic, near-linear field pattern to enable effective spatial encoding.

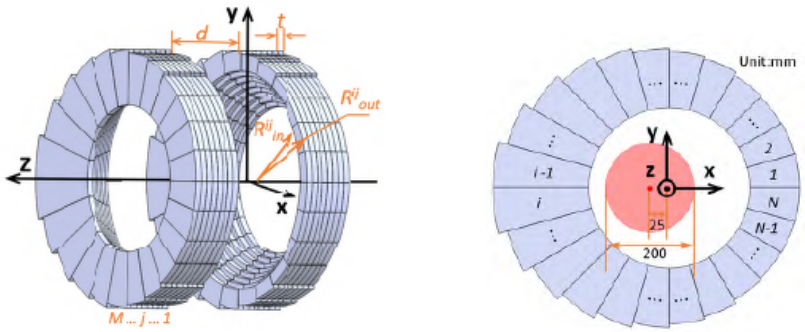

Figure 7.1: Irregular-shaped IO ring pair, 3D and front view. Adapted from [6]

The result of the design process is illustrated in Figure 7.2, where the monotonicity of the magnetic field along the x-direction is clearly visible. In order to enable 2D imaging, the system would need to be mechanically rotated, acquiring one projection at each angular position. While this configuration does not address the issue of z-axis symmetry, which would still be present, it offers the advantage of eliminating the need for active gradient coils along the x and y directions, as these components are inherently encoded in the magnetic structure itself.

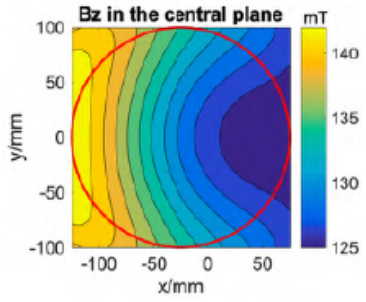

Figure 7.2: Simulated B_0 field in the central xy-plane. Adapted from [6]

Disturbing Magnets

A further strategy to introduce spatial encoding in the presence of field symmetry involves the use of disturbing magnets, which are auxiliary magnetic blocks strategically placed to locally perturb the field distribution. In the configuration developed in this thesis, the main magnetic field is axially symmetric, and the geometry of the system is more naturally described in polar coordinates (r, θ, z) rather than in Cartesian space: the magnetic field distribution depends indeed primarily on the

radial distance r and the axial position z, while remaining invariant with respect to the angular coordinate θ. This symmetry causes ambiguity in spatial encoding, as multiple points in the field of view experience the same magnetic field strength and therefore produce identical resonance frequencies, making them impossible to distinguish in the reconstructed image.
By introducing small permanent magnets the resulting superposition of the primary and disturbing fields leads to a modified field profile in which the axial symmetry is intentionally disrupted. To fully exploit this approach for angular encoding, the additional magnet has to be rotated around its central axis, allowing the disturbing field to progressively "illuminate" different angular sections of the FOV. By doing so, multiple projections can be acquired at different orientations, enabling the reconstruction of spatially resolved information across the circular cross-section of the FOV.

As a conceptual extension of the single disturbing magnet approach, it can be considered the idea of multiple disturbing magnets positioned at staggered heights along the z-axis, each specifically designed to locally perturb the magnetic field within a distinct axial slice, as shown in Figure 7.3.

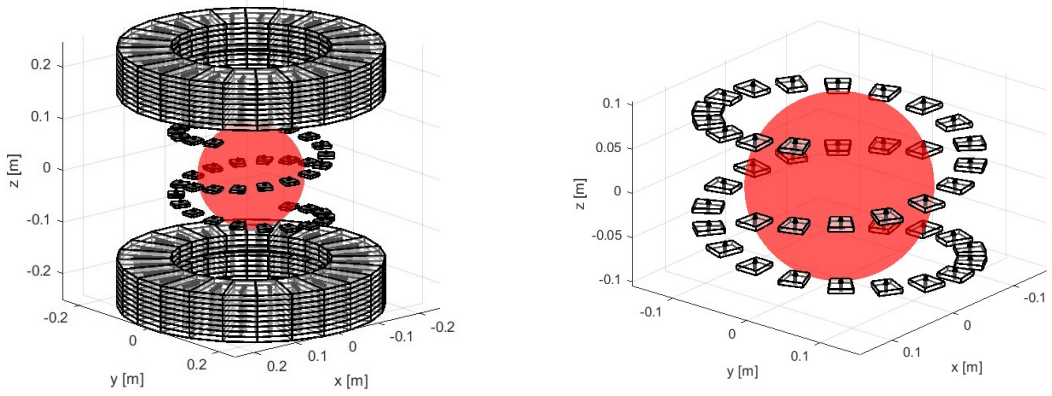

Figure 7.3: Possible configuration with staggered disturbing magnets

When combined with rotational motion it enables the acquisition of a complete 2D image for each axial layer. By repeating this process across multiple layers, each perturbed by a disturbing magnet placed at a different z-position, a full 3D image can be reconstructed through the stacking of individually encoded 2D slices.
This technique could serve as a passive alternative to z-gradient coils, avoiding the need for active components. In this setup, each magnet is associated with a specific axial slice: when a signal is detected under its influence, it can be directly linked to the corresponding z-position. In other words, the spatial localization along z is

not achieved through a linear gradient, but through the known and localized effect of each disturbing magnet, which acts as a spatial marker.
It is important to underline that this approach currently remains a conceptual proposal, and several technical challenges must be overcome before it can be seen as a viable solution. For instance:

- Field overlap and interference: disturbing magnets placed too close along the z-axis may generate overlapping perturbations, making it difficult to isolate individual slices clearly;
- Axial resolution limitations: the spacing between the disturbing magnets directly defines the achievable resolution along the z-axis. For example, if the magnets are placed in planes separated by $5\,mm$, the maximum theoretical resolution along z would also be $5\,mm$. While this level of detail may be acceptable for certain applications, it could be insufficient in scenarios requiring finer anatomical resolution or more precise spatial discrimination;
- Impact on overall field homogeneity: introducing localized distortions may degrade the field uniformity in adjacent regions, potentially affecting image quality.

These limitations highlight the need for a much more in-depth investigation to assess the actual effectiveness and practical utility of this approach. Nevertheless, it represents an intriguing passive alternative for enabling axial and rotational encoding in low-field, coil-free MRI architectures.

Acknowledgments

A conclusione di questa tesi un primo ringraziamento non può che andare al Prof.Freschi. La sua attenzione e la sua pazienza hanno reso possibile tutto ciò, dandomi l'opportunità di vivere un'esperienza in Austrlia che mi ha profondamente arricchito sia dal punto di vista personale che professionale.
Grazie a tutti i miei amici, tanti, tanto che racchiudervi in questa semplice definizione sarebbe addirittura riduttivo.
Quelli "vecchi", quelli di casa, ormai colonne portanti della mia vita che, sono sicuro, nè la distanza nè il tempo riusciranno ad erodere.
Quelli "nuovi", con cui, tra gioie e dolori, ho condiviso questo percorso che si chiude oggi e che, posso dirlo con certezza, hanno reso questo periodo il più bello della mia vita.
Il ringraziamento più importante, tuttavia, va senz'altro alla mia famiglia che mi ha accompagnato, con la più assoluta fiducia, lungo qualunque strada abbia deciso di intraprendere.
Ai miei nonni, che senza mai interferire mi hanno sempre fatto sentire il proprio sostegno. La notizia di ogni esame passato veniva da loro accolto come una laurea. Ora posso dire loro che questa effettivamente lo è.
A Patrizia, Antonino e al piccolo Nunzio, sperando per di poter rappresentare per quest'ultimo, oggi come nel futuro, una fonte di ispirazione.
Infine a mio padre, mia madre, mia sorella, a loro va il ringraziamento più importante. Diversi sono gli ostacoli che ci sono stati posti dinnanzi e diversi ancora ne avremo da saltare. Sono sicuro che riusciremo a farlo come fino ad ora abbiamo fatto: insieme. A voi dedico questa tesi, a voi dedico questa laurea.
Io ci ho messo solo qualche ora in aula studio o qualche pagina di appunti scritti malamente, tutto resto ce lo avete messo voi.

Bibliography

[1] Jianming Jin. Electromagnetic Analysis and Design. Wiley-IEEE Press, 1st edition, 2008. ISBN 978-0-470-13755-5.

[2] Christophe Rubeck, Jean-Paul Yonnet, Hicham Allag, Benoît Delinchant, and Olivier Chadebec. Analytical calculation of magnet systems: Magnetic field created by charged triangles and polyhedra. IEEE Transactions on Magnetics, 57(3):1–8, 2021. doi: 10.1109/TMAG.2020.3041504.

[3] Peter Blümler and Helmut Soltner. Practical concepts for design, construction and application of halbach magnets in magnetic resonance. Applied Magnetic Resonance, 2023. doi: 10.1007/s00723-023-01602-2. URL `https://doi.org/10.1007/s00723-023-01602-2`.

[4] Clarissa Z Cooley, Melissa W Haskell, Stephen F Cauley, Charlotte Sappo, Cristen D Lapierre, Christopher G Ha, Jason P Stockmann, and Lawrence L Wald. Design of sparse halbach magnet arrays for portable mri using a genetic algorithm. IEEE Transactions on Magnetics, 54(1):1–12, 2018. doi: 10.1109/TMAG.2017.2751001.

[5] Zhi Hua Ren, Wen Chuan Mu, and Shao Ying Huang. Design and optimization of a ring-pair permanent magnet array for head imaging in a low-field portable mri system. IEEE Transactions on Magnetics, 55(1):1–8, 2019. doi: 10.1109/TMAG.2018.2876679.

[6] Tianqi Ren, Zixuan Gong, and Yansheng Huang. Compact low-field mri system with optimized permanent magnet design using a self-shielding halbach array. IEEE Transactions on Instrumentation and Measurement, 71:1–10, 2022. doi: 10.1109/TIM.2022.3156733.

[7] Davide Vai. Progettazione di array halbach a magneti permanenti per la risonanza magnetica a basso campo. Master's thesis, Politecnico di Torino, 2023. Supervisor: Prof. Fabio Freschi.

[8] Clarissa Zimmerman Cooley, Jason P. Stockmann, Brandon D. Armstrong, Mathieu Sarracanie, Michael H. Lev, Matthew S. Rosen, and Lawrence L. Wald. Two-dimensional imaging in a lightweight portable mri scanner without gradient coils. Magnetic Resonance in Medicine, 73(2):872–883, 2015. doi: 10.1002/mrm.25147.

[9] Michael Poole. Improved Equipment and Techniques for Dynamic Shimming in High Field MRI. Ph.d. thesis, University of Nottingham, 2007. Accessed on request.

[10] Clarissa Z. Cooley, Patrick C. McDaniel, Jason P. Stockmann, Sai Abitha Srinivas, Stephen F. Cauley, Monika Śliwiak, Charlotte R. Sappo, Christopher F. Vaughn, Bastien Guerin, Matthew S. Rosen, Michael H. Lev, and Lawrence L. Wald. A portable scanner for magnetic resonance imaging of the brain. Nature Biomedical Engineering, 5:229–239, 2021. doi: 10.1038/s41551-020-00641-5. URL https://www.nature.com/articles/s41551-020-00641-5.

[11] Fernando Galve, Eduardo Pallás, Teresa Guallart-Naval, Pablo García-Cristóbal, Pablo Martínez, José M. Algarín, Jose Borreguero, Rubén Bosch, Francisco Juan-Lloris, José M. Benlloch, and Joseba Alonso. Elliptical halbach magnet and gradient modules for low-field portable magnetic resonance imaging. NMR in Biomedicine, 37:e5258, 2024. doi: 10.1002/nbm.5258. URL https://doi.org/10.1002/nbm.5258.

[12] Ting-Ou Liang, Yan Hao Koh, Tie Qiu, Erping Li, Wenwei Yu, and Shao Ying Huang. High-performance permanent magnet array design by a fast genetic algorithm (ga)-based optimization for low-field portable mri. Journal of Magnetic Resonance, 345:107309, 2022. doi: 10.1016/j.jmr.2022.107309. URL https://doi.org/10.1016/j.jmr.2022.107309.

[13] Charles Z. Cooley, Patrick C. McDaniel, Jason P. Stockmann, Thomas Witzel, and Lawrence L. Wald. Gradient coil design for low-field mri using halbach permanent magnet arrays. Magnetic Resonance in Medicine, 85(1):646–661, 2021.