



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Permanent-Magnet Optimization for Stellarators as Sparse Regression

Alan A. Kaptanoglu, Tony Qian, Florian Wechsung, and Matt Landreman

Phys. Rev. Applied **18**, 044006 — Published 4 October 2022

DOI: [10.1103/PhysRevApplied.18.044006](https://doi.org/10.1103/PhysRevApplied.18.044006)

Permanent magnet optimization for stellarators as sparse regression

Alan A. Kaptanoglu*

*Institute for Research in Electronics and Applied Physics,
University of Maryland, College Park, MD, 20742, USA*

Tony Qian

*Princeton Plasma Physics Laboratory,
100 Stellarator Road, Princeton, NJ, 08540, USA*

Florian Wechsung

*Courant Institute of Mathematical Sciences,
New York University, New York, NY, 10012, USA*

Matt Landreman

*Institute for Research in Electronics and Applied Physics,
University of Maryland, College Park, MD, 20742, USA*

A common scientific inverse problem is the placement of magnets that produce a desired magnetic field inside a prescribed volume. This is a key component of stellarator design, and recently permanent magnets have been proposed as a potentially useful tool for magnetic field shaping. Here, we take a closer look at possible objective functions for permanent magnet optimization, reformulate the problem as sparse regression, and propose an algorithm that can efficiently solve many convex and nonconvex variants. The algorithm generates sparse solutions that are independent of the initial guess, explicitly enforces maximum strengths for the permanent magnets, and accurately produces the desired magnetic field. The algorithm is flexible, and our implementation is open-source and computationally fast. We conclude with two new permanent magnet configurations for the NCSX and MUSE stellarators. Our methodology can be additionally applied for effectively solving permanent magnet optimizations in other scientific fields, as well as for solving quite general high-dimensional, constrained, sparse regression problems, even if a binary solution is required.

Keywords: permanent magnets, stellarators, sparse regression, optimization

I. INTRODUCTION

Magnet design is required in a large number of scientific domains, but it is a fundamentally ill-posed problem because many different magnet designs can produce an identical target magnetic field via the Biot-Savart law. An extreme case for magnetic design comes from plasma experiments that are investigating nuclear fusion; these experiments often require very strong and complex magnetic coils. One class of plasma experiments, stellarators, particularly relies on sophisticated coil design algorithms in order to produce ideal magnetic fields for confining plasma [1, 2]. These three-dimensional magnetic fields must be carefully shaped in order to provide high-quality confinement of charged particle trajectories and many other physics objectives. Stellarator optimization is usually divided into two stages. The first is a configuration optimization using fixed-boundary magnetohydrodynamic (MHD) equilibrium codes to obtain MHD equilibria with desirable physics properties [3–5].

After obtaining the optimal magnetic field in this first stage, coils must be designed to produce these fields, subject to a number of engineering constraints such as a minimum coil-to-coil distance, maximum forces on the coils [6], maximum curvature on the coils, and many other requirements [7]. The result is that stellarator coils are often very complex 3D shapes,

raising the cost and difficulty of manufacturing. The primary cost of the W-7X stellarator program was the manufacture of these complex coils with very tight engineering tolerances [8]. The NCSX stellarator was never finished in large part because of similar fabrication and assembly obstacles [9].

A. A role for permanent magnets

Until recently, the highly-shaped and precise nature of stellarator equilibria implied the necessity of these very complex coils. One way to circumvent this requirement, first proposed in Helander et al. [10], is to simplify stellarator coil designs by surrounding a stellarator with a manifold of permanent magnets that can provide significant portions of the magnetic field. These permanent magnets cannot be used to generate a net toroidal flux, so traditional magnetic coils are still required. Instead, the permanent magnets allow for significant reductions in the coil complexity and cost. Moreover, permanent magnets operate without power supplies, require minimal cooling, ameliorate magnetic ripple due to discrete coils, and facilitate improved diagnostic access to the plasma chamber. However, there are also some potential disadvantages, including the inability to turn off the field, the possibility of demagnetization, and an upper limit on the achievable field strength; material science advances [11] for permanent magnets could significantly address the latter two disadvantages. Despite these potential setbacks, the low cost and simple man-

* Corresponding author (akaptano@umd.edu).

ufacture of permanent magnet stellarators are tantalizing, especially for university-lab-scale experiments. There is already a promising experimental effort to produce a very cheap but practical permanent magnet stellarator [12]. In the present work, we will show that permanent magnet stellarators have an additional advantage. Unlike filamentary coil optimization, the permanent magnet optimization problem can be expressed as sparse regression. Thanks to the prolific scientific interest in variations of sparse regression, permanent magnet optimization can subsequently be relatively well understood and rapidly solved. Moreover, traditional stellarator coil optimization using a winding surface [13], even with complicated requirements on coil-coil forces [6], can also be formulated as sparse regression.

B. Motivation

Simplifying coils by utilizing permanent magnets comes with its own challenges. Sophisticated algorithms are still required to find high-quality configurations of permanent magnets. Optimization for determining the optimal placement of permanent magnets (subject to minimizing the cost of the magnets and various engineering constraints) is in a somewhat early stage. There are currently several different formulations and associated algorithms for addressing the permanent magnet optimization problem [14–19], but the relationships between them are often unclear. Some of the optimization problems are multi-stage or use discrete optimization, and the “best” set of loss terms is an open question. These nonconvex problems seem to exhibit many local minima and high sensitivity to the initial conditions. Often additional post-processing optimization steps are taken to further improve the initial optimization solutions. Despite these methods relying on some heuristics and post-processing, they often produce very high-quality solutions. Similarly, genetic algorithms have been used to generate permanent magnet configurations in the magnetic resonance imaging (MRI) community [20, 21]. Here too, there have been few conclusions about optimality, and the algorithms are not always extensible when future work necessitates higher-dimensional problems or additional constraints.

If the permanent magnet optimization could be better understood and more efficiently solved in the stellarator community, “off-the-shelf” permanent magnet stellarators could be rapidly designed for cheap university-lab-scale experiments with very simple toroidal field (TF) coils. If such stellarators can be constructed and widely distributed, stellarator expertise could be rapidly cultivated and the parameter space of quasi-symmetric and quasi-isodynamic stellarators could be rapidly explored. This is especially important, as multi-stage stellarator optimization has recently generated highly quasi-symmetric configurations [22, 23] and near-axis expansions [24–28] have facilitated efficient numerical explorations and new discoveries within a very large parameter space of quasi-symmetric stellarator configurations. Outside the stellarator optimization community, improvements in understanding and algorithms for permanent magnet optimization prob-

lems would have significant repercussions for better permanent magnet configurations in MRI, automobile design, and other industrial uses.

C. Contributions of the present work

Towards rapid design and distribution of cheap permanent magnet stellarators, this work provides a foundation for understanding the permanent magnet optimization problem by showing it can be reformulated as sparse regression. Then, we provide a computationally efficient, easy-to-use, generally applicable, and open-source code for computing permanent magnet configurations for stellarators. The algorithm produces sparse solutions that are independent of the initial guess, explicitly enforces hard constraints on the dipole moment magnitudes, and illustrates that continuous optimization is sufficient for generating high-performance permanent magnet stellarators. As far as we are aware, it joins a very small and recent list of sparse regression algorithms that can effectively solve problems with $\mathcal{O}(10^6)$ optimization variables, and may be the first algorithm capable of handling a similar number of convex constraints and producing binary solutions. We can regularly solve permanent magnet optimizations that depend on dense matrices of billions of elements.

In particular, we provide a new relax-and-split method for general permanent magnet optimization in the open-source SIMSOPT code [5]. Unlike previous work, the entirety of the permanent magnet pipeline, i.e. the geometry, optimization tools, post-processing, etc., is contained in this single open-source tool. The effectiveness of this new algorithm is demonstrated by finding and illustrating two high-performance permanent magnet stellarator configurations. All of the present work’s methodology and results can be found in the SIMSOPT code.

D. Permanent magnet optimization as sparse regression

The sparsity-promoting optimization problems that appear in this work are the foundation of the field of sparse regression, which encompasses sparse system identification [29–31], compressed sensing [32], and many other tasks in signal and image processing [33]. This fundamental relationship with optimization problems occurring in sparse regression means that efficient and well-understood algorithms can be immediately brought to bear on the permanent magnet problem.

There are two primary differences between the optimization problem that we solve in this paper and the many corresponding sparse regression problems in other scientific fields. First, the latter typically has $\mathcal{O}(10^2)$ or fewer unknowns, while realistic permanent magnet coil optimizations can easily have $\mathcal{O}(10^5)$ unknowns (with a corresponding $\mathcal{O}(10^5)$ constraints). It follows that, unlike most sparse regression applications, permanent magnet optimization requires algorithms that scale well with the number of unknowns and the number of constraints. A notable exception is that sparse regression for

image and signal processing is often very high-dimensional, and in this case approaches to this problem typically either (1) convexify the problem, e.g. with the l_1 norm, or (2) use a “greedy” algorithm that iteratively and rapidly solves the problem, albeit often with weak guarantees on optimality [33]. Both of these approaches can struggle to solve binary problems, e.g. the desired solution for permanent magnet optimization is not just a sparse set of magnets but one in which the magnets are all either *maximum-strength* or *exactly zero*.

Second, even the *convex* formulation of permanent magnet optimization, which omits any engineering requirements on the permanent magnets, often gives rise to different initial conditions converging to different final solutions. This pseudo-paradox is simply a result of the strong ill-posedness of the optimization problem; one should imagine the optimization landscape as a convex space but with a very large, flat valley containing many quasi-minima. It is well known that this issue is addressed with regularization. Indeed, without sufficient regularization, each new initial condition will result in a quasi-minima solution that is correct to numerical precision. This problem is so ill-posed that the quasi-minima often look like extremely different configurations of permanent magnets. Fig. 1 shows that modest Tikhonov regularization fully circumnavigates this qualitative behavior, although the regularization may make it more difficult to access the desired parameter space of sparse, binary solutions.

This type of ill-posedness also occurs in varieties of sparse regression. A notable example occurs in sparse system identification, where the goal is to search through very large libraries of candidate functions. In that case, the goal is to discover *the* underlying equations of a dynamical system; there is a *right* answer. In the context of permanent magnet optimization, this strong ill-posedness can be seen as a strength rather than a weakness since (modulo engineering constraints) there is no *right* answer; *any* permanent magnet configuration that produces negligible magnetic field errors on the plasma surface is a suitable solution. In other words, there are significant degrees of freedom available for incorporating additional requirements regarding the configuration of magnets.

However, even in the permanent magnet problem there are some potential downsides of ill-posedness. Coupled with true local minima in the nonconvex setting, ill-posedness can have significant consequences regarding the number, manufacture, and cost of the permanent magnets. It also impedes conclusions about the possible parameter space of configurations; it is possible that higher-quality or easier-to-engineer configurations are available but have not been found. These problems can be compounded by a poor choice of loss terms in the objective function or because such configurations cannot be easily found with the algorithms used to minimize the objective.

A relax-and-split algorithm is ideal for high-dimensional, constrained sparse regression. In a relax-and-split approach, an iterative solve is set up between a convex subproblem over any set of convex constraints, and an “easy” nonconvex subproblem that renders the overall approach very computationally efficient. The relax-and-split iterations continue until the convex and nonconvex subproblems both converge. The relax-and-split formulation exhibits excellent and well-

understood local convergence properties [34], circumnavigates the initial guess sensitivity of previous work (justified later in Sec. II E), and can be easily extended in various ways from the baseline implementation in the open-source PySINDy code [35, 36] that is used for sparse system identification. Variants of relax-and-split were used in Champion et al. [30] to solve system identification problems with affine constraints in the optimization variables, extended for additional constraints and nonconvex loss terms in Kaptanoglu et al. [37, 38], recently used as a denoising strategy in Hokanson et al. [39], and increasingly contribute to optimization problems found outside the field of system identification [40–42]. Our new implementation for permanent magnets can also be used to rapidly solve complex sparse regression problems with $\mathcal{O}(10^4) - \mathcal{O}(10^6)$ unknowns and a similar number of constraints. To our knowledge, solving such large, constrained optimization problems with the l_0 norm has been largely infeasible with current algorithm implementations in the sparse regression field until very recent work [43–45]. By framing permanent magnet optimization as sparse regression, we build a bridge to these powerful tools.

II. METHODOLOGY

Permanent magnets can be approximated as magnetic dipoles if the distance between the field evaluation point and the center of the permanent magnet is substantially larger than the size of the permanent magnet. This is a substantial advantage for optimization because the magnetic field generated by a magnetic dipole exhibits a simple analytic expression that is independent of the permanent magnet’s geometry and linear in the dipole moment \mathbf{m} . Green’s functions can also be used for fully calculating the near-fields of permanent magnets with simple shapes like spheres and cubes, while retaining the linearity in \mathbf{m} . Moreover, the dipole moment of the permanent magnet is approximately independent of the external field if the permanent magnet permeability is low.

Together, these approximations facilitate a very simple form for the total magnetic field. It follows that a straightforward way to imagine placing permanent magnets is to approximate the magnets as magnetic dipoles, discretize a large computational domain into elements of a mesh, and represent each element as a single dipole at the center (again assuming that the element size is substantially smaller than the distance between the element and the plasma). If the total number of discrete dipoles \mathbf{m}_i is D , then the total magnetic field \mathbf{B}_M is simply a sum over all the dipoles,

$$\mathbf{B}_M = \frac{\mu_0}{4\pi} \sum_{i=1}^D \left(\frac{3\mathbf{m}_i \cdot \mathbf{r}_i}{\|\mathbf{r}_i\|_2^5} \mathbf{r}_i - \frac{\mathbf{m}_i}{\|\mathbf{r}_i\|_2^3} \right). \quad (1)$$

Here μ_0 is the vacuum permeability, \mathbf{r}_i is the vector between the evaluation point and the i^{th} dipole, and SI units will be used throughout this work. Critically, Eq. (1) is linear in the dipole moments \mathbf{m}_i . The primary target for stage-2 coil optimization is to, together with the traditional magnetic coils, use the permanent magnets to fully match the desired mag-

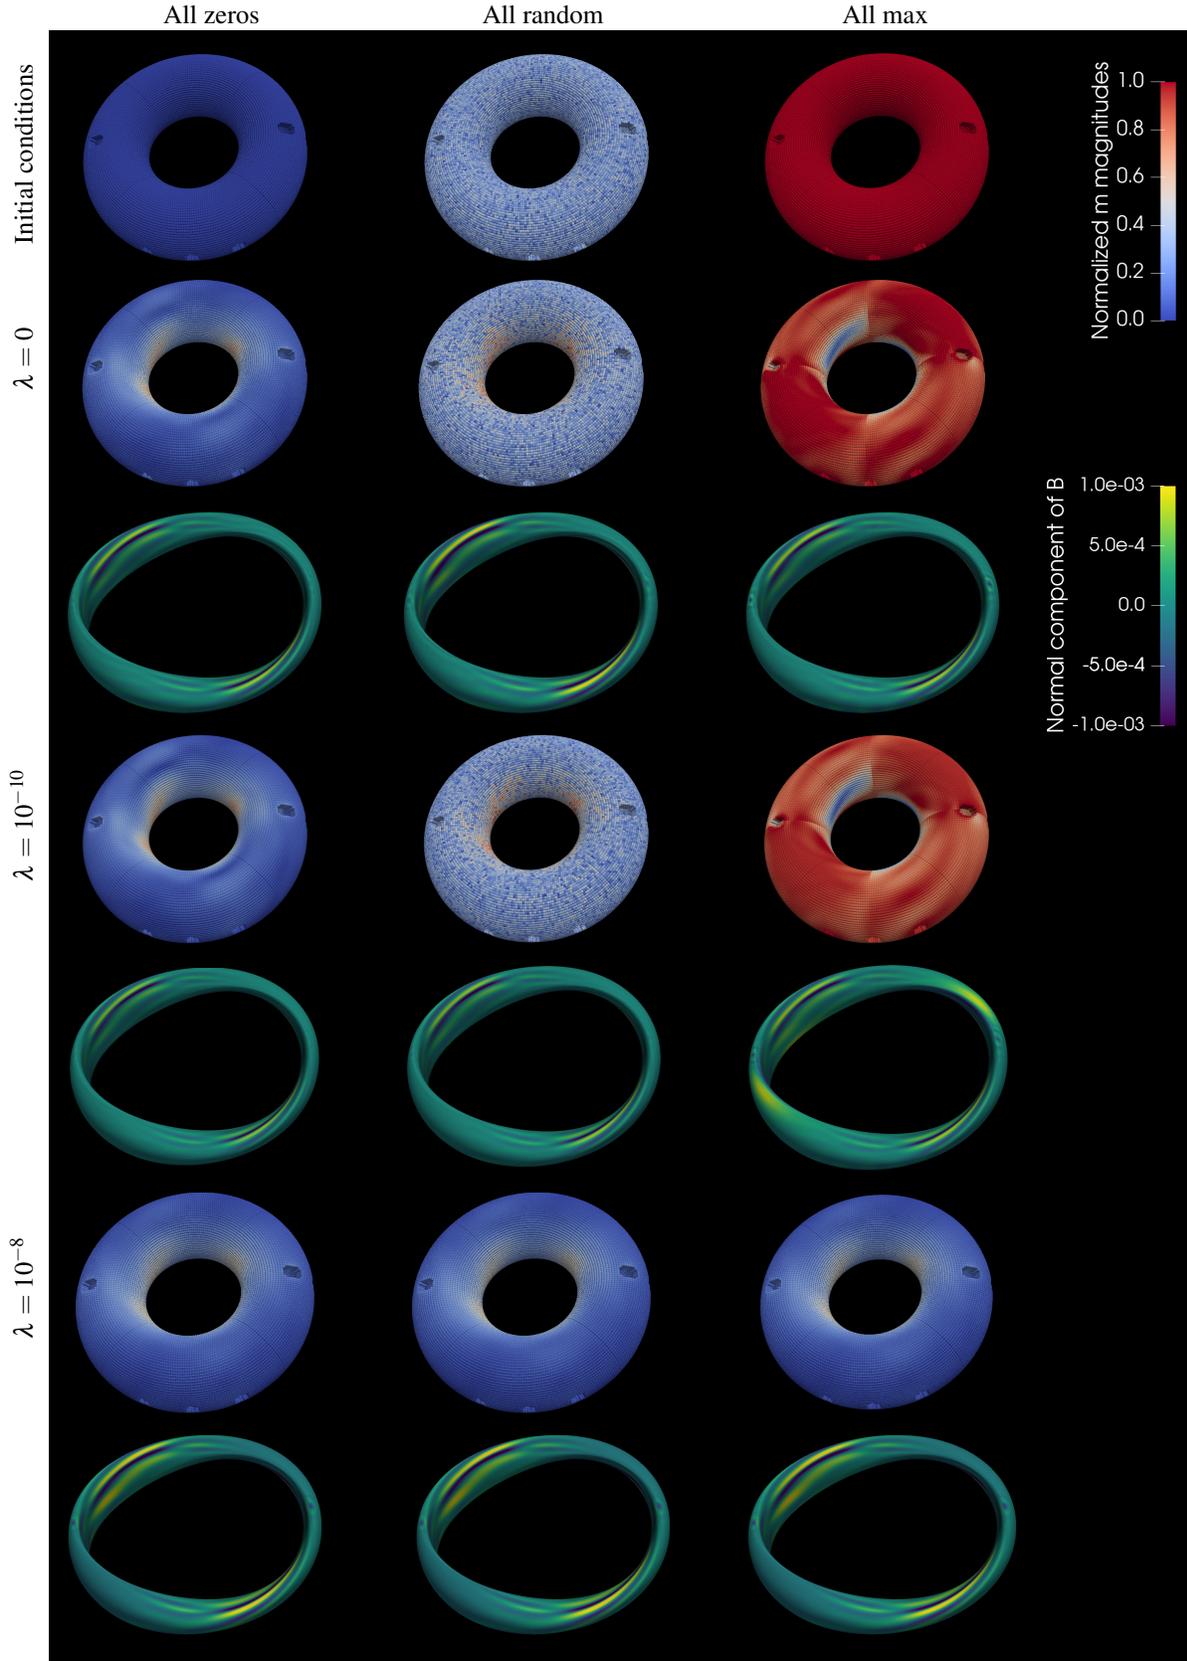


FIG. 1: Permanent magnet manifolds (rows 1, 2, 4, 6) and $\mathbf{B} \cdot \hat{\mathbf{n}}$ on the plasma surface (rows 3, 5, 7), generated by solving (5) for the MUSE permanent magnet stellarator. The toroidal field coils are omitted from the illustration. Tikhonov regularization with $\lambda = \{0, 10^{-10}, 10^{-8}\}$ and three different initial conditions were used, until a high performance value $f_B \sim 10^{-8}$ was achieved. Despite the convexity, the ill-posedness of the optimization results in vastly different permanent magnet configurations (and some variation in $\mathbf{B} \cdot \hat{\mathbf{n}}$ on the plasma surface) depending on the initial condition. Only with sufficient regularization can we break the degeneracy from the ill-posedness.

netic fields at the plasma boundary surface,

$$f_B = \int ((\mathbf{B}_M + \mathbf{B}_P + \mathbf{B}_C) \cdot \hat{\mathbf{n}})^2 d^2x. \quad (2)$$

That is, we use the permanent magnets to minimize the value of f_B . Here $\hat{\mathbf{n}}$ is the normal vector to the plasma boundary, and \mathbf{B}_P and \mathbf{B}_C are the magnetic fields generated by plasma current and the traditional coils, respectively. For the purposes of this work, \mathbf{B}_P is the field from a stage-1 optimized plasma boundary and \mathbf{B}_C is the field generated from a minimal set of basic coils that produce the net toroidal flux. Eq. (2) is convex because it can be written as a linear least-squares term in the \mathbf{m}_i , as we show explicitly in Appendix A.

Next, we assume a straightforward grid of permanent magnet locations. Cartesian (x, y, z) , cylindrical (r, ϕ, z) , and simple toroidal $(r_{\text{minor}}, \phi, \theta)$ coordinate systems are implemented in the code. Rectangular cubes, curved square bricks, and simple toroidal shells are ideal shapes for permanent magnets because they facilitate cheap, mass manufacturing and straightforward assembly. Cylindrical coordinates are used in Sec. III A and simple toroidal coordinates are used in Sec. III B. Eq. (2) is independent of the coordinate-system choice; the coordinate system here primarily serves to specify which coordinate directions are considered grid-aligned, since this is an attractive engineering property to promote.

As in Hammond et al. [46], we define the volume as the space between two toroidal limiting surfaces. The inner surface that encloses the plasma is usually chosen to be the experimental vacuum vessel. If there is no designed vacuum vessel, a simple transformation can be used to generate an inner toroidal surface as follows. The plasma boundary normal vectors in cylindrical coordinates are projected onto the (r, z) plane of the corresponding quadrature point (so that the new surface locations are defined at identical poloidal cross-sections), and then multiplied by an overall offset value to generate a surface for the inner toroidal boundary of the permanent magnet configuration. The outer limiting surface is generated similarly, using the projected normal vectors on the inner toroidal surface. For moderately shaped equilibria, these simple transformations work well to generate the permanent magnet volume. Any curved square bricks that are not between the inner and outer surfaces are eliminated. Custom grids are also admissible, and future work could straightforwardly implement additional grid generation schemes.

As discussed in Sec. I D, the linear least-squares part of the permanent magnet problem is strongly ill-posed. Typically Tikhonov regularization with strength λ is added to the stage-2 permanent magnet optimization for regularization of f_B ,

$$f_m = \lambda \sum_{i=1}^D \|\mathbf{m}_i\|_2^2, \quad (3)$$

and this term is also convex in the \mathbf{m}_i . A downside of this term is that it tends to select for solutions with weak but nonzero

magnets. For more general regularization terms, we consider,

$$f_m = \lambda \sum_{i=1}^D R(\mathbf{m}_i), \quad (4)$$

where the sum over i could also be chosen to be inside the regularizer R . Here $R(\mathbf{m}_i)$ is smooth and convex, such as the commonly used l_2 norm; nonsmooth and nonconvex regularizers are discussed below but we denote these separately.

A. Optimization objectives

Before formulating the full optimization problem to solve, it is illustrative to be explicit about the goals of the optimization. Consider the following objectives that we would like to promote or impose in the optimization problem,

1. Fit the MHD equilibrium fields: minimize f_B .
2. Regularize f_B : minimize f_m .
3. Avoid magnetic dipole values above the maximum magnetization of the material: impose the hard constraints that $\|\mathbf{m}_i\|_2 \leq m_i^{\text{max}}$ where the maximum value for dipole moment i is defined through the remanence field B_{rem} and the cell volume V_i^{cell} , $m_i^{\text{max}} = \mu_0^{-1} B_{\text{rem}} V_i^{\text{cell}}$. For this work, we assume $B_{\text{rem}} = 1.465$ T to match the value used for FAMUS runs of the MUSE stellarator. MUSE uses commercial ‘‘N52’’ neodymium-iron-boron (Nd2Fe14B) magnets.
4. Produce a sparse, binary solution. Bias the dipoles towards either satisfying $\|\mathbf{m}_i\|_2 = 0$ or $\|\mathbf{m}_i\|_2 = m_i^{\text{max}}$ for all i , so that permanent magnets are either omitted or placed with approximately maximum strength. Moreover, minimize the cost of the permanent magnetic configuration by using as few magnets as possible.
5. Promote properties that reduce the engineering complexity of the magnet configuration, such as constraining each dipole’s direction to be ‘‘grid-aligned’’, i.e. perpendicular to the face of each grid element.

Below, we start with a convex formulation that is sufficient for objectives 1–3 in Sections II B–II C, and then add in non-convexity for promoting sparse configurations of permanent magnets that address objectives 4–5 in Sections II D–II E.

B. A first attempt at formulating the optimization

We now propose a first objective function that accomplishes the optimization objectives 1–3. First, define the vector of optimization variables, $\mathbf{m} = [m_1^x, m_1^y, m_1^z, m_2^x, \dots, m_D^z]$. Then the proposed objective function can be formulated,

$$\arg \min_{\mathbf{m}} \frac{1}{2} \|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2^2 + \lambda R(\mathbf{m}) \quad (5)$$

$$\|\mathbf{m}_i\|_2^2 \leq (m_i^{\text{max}})^2, \quad i = 1, \dots, D.$$

The quadratic constraints are convex since each \mathbf{m}_i vector must be contained in a l_2 ball in \mathbb{R}^3 . For the remainder of the work, we will denote as \mathbb{S}_m the hypersurface spanned by the intersection of the spherical constraints. Critically for efficiency, the constraints are separable. Unlike most stellarator optimization problems, we do not convert these constraints into a related loss term – we impose them as hard constraints. Interior point methods (IPMs) may also be used here, where the inequality constraints would be converted into barrier functions in the loss terms, but IPMs usually do not scale very well to high-dimensional settings. Fortunately, for convex problems with a large number of separable, convex constraints, there are other fast algorithms.

Before moving on to an algorithm for Eq. 5, we note here that the demagnetization effects from finite permanent magnet coercivity could also be modeled as convex constraints. These constraints could be critical for accurate modeling of stellarators with large magnetic field strengths. If the dipole at position \mathbf{x}_i demagnetizes if the external field strength is greater than B_{\max} , this can be posed as

$$\|\mathbf{B}_M(\mathbf{m}, \mathbf{x}_i) + \mathbf{B}_P(\mathbf{x}_i) + \mathbf{B}_C(\mathbf{x}_i)\|_2^2 \leq B_{\max}^2, \quad (6)$$

It should be clear that the field from the i th dipole is excluded from expression (6) when evaluated at \mathbf{x}_i , and that for permanent magnets that are very close to the i th permanent magnet, there are significant errors introduced by the dipole approximation (however, these errors can be avoided with spherical permanent magnets or accounted for if other simple shapes are used). There are constraints of the Eq. (6) type from all the dipole locations, adding another D constraints to the values in \mathbf{m} if each constraint is active. Each constraint is a convex but not separable constraint on the \mathbf{m}_i and could be implemented with barrier functions in future work.

C. An algorithm for simple permanent magnet optimization

The optimization problem in (5) is a convex problem with separable, convex, spherical constraints. Therefore, it can be solved with many different algorithms, and the global minimum can be found. The primary issues are the high-dimension of the problem and the large number of constraints. The problem in (5) forms the backbone of the present work, so we first provide an algorithm that can efficiently solve this problem despite the high dimension.

Many algorithms have been considered for quadratic programs with quadratic constraints [47]. It is only recently that fast solvers have been devised for exactly the case of high-dimensional convex optimization problems with a large number of spherical constraints [48]. In particular, an algorithm is proposed in Bouchala et al. [49] that can optimally solve these problems, and we adopt this algorithm for our purposes. An extended algorithm and other variations [50] are also available if additional equality constraints are present. The algorithm is essentially projected-gradient descent (the standard algorithm for high-dimensional, smooth, constrained optimization) with a conjugate-gradient acceleration, and as such, it relies on a

explicit form for the projection operator,

$$P_{\mathbb{S}_m}(\mathbf{m}) = \arg \min_{\mathbf{y} \in \mathbb{S}_m} \|\mathbf{m} - \mathbf{y}\|_2^2. \quad (7)$$

The projection can be decomposed as,

$$P_{\mathbb{S}_m}(\mathbf{m}) = [P(\mathbf{m}_1), \dots, P(\mathbf{m}_D)], \quad (8)$$

by the separability of the constraints, and each of the projections satisfies,

$$P(\mathbf{m}_i) = \arg \min_{\|\mathbf{y}\|_2^2 \leq (m_i^{\max})^2} \|\mathbf{m}_i - \mathbf{y}\|_2^2 = \frac{\mathbf{m}_i}{\max\left(1, \frac{\|\mathbf{m}_i\|_2}{m_i^{\max}}\right)}. \quad (9)$$

Lastly, we typically initialize the algorithm with an initial guess for the dipoles as all zeros. The natural guess is $\mathbf{m}^{(0)} = P_{\mathbb{S}_m}(\mathbf{A}^\dagger \mathbf{b})$, where \mathbf{A}^\dagger denotes the pseudo-inverse of \mathbf{A} , but we found that this often generates very poor initial guesses when the constraints are very active. Alternatively, when the problem is strongly ill-posed, it is often useful to start with the dipoles at maximum strength to preference the algorithm towards solutions with strong magnets. Regardless, the algorithm requires that $\mathbf{m}^{(0)} \in \mathbb{S}_m$. Many other initial conditions were tested to verify that, with sufficient regularization, the algorithm correctly converges to the minimum of this convex problem, as in Fig. 1.

D. Sparse optimization for permanent magnet design

Unfortunately, the formulation in (5) is not sufficient for a sparse, binary solution. These are attractive properties to promote, since it means only a sparse collection of maximum-strength permanent magnets can be used. We begin this section with a discussion of one of the foundations of sparsity-promoting optimization: the l_0 “norm”, which is simply an operator that counts up the number of nonzero terms in a vector. The l_0 norm has not been used before for permanent magnet optimization because it appears very challenging to address with traditional optimization. For instance, the l_0 norm is non-smooth, preventing the straightforward application of generic algorithms for nonconvex optimization based on gradient and Hessian information such as BFGS [51]. The high-dimension of the problem and large number of constraints further reduces the number of fast algorithms that are suitable for addressing objective functions with the l_0 norm. Nonetheless, we define a new objective function with the l_0 norm and borrow recent ideas from the field of sparse regression to effectively solve it,

$$\arg \min_{\mathbf{m} \in \mathbb{S}_m} \frac{1}{2} \|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{m}\|_2^2 + \alpha \|\mathbf{m}\|_0. \quad (10)$$

We will show below that the formulation in (10) now satisfies all of the requirements for the permanent magnet optimization problem and the partial convexity of the objective function can be utilized. Moreover, the objective function in (10) can be altered without significant changes in procedure; substitutions for the l_0 norm in (10) can be made for other regularizers,

depending on the optimization goals. Examples include the l_1 , group $l_{2,1}$, and group $l_{0,1}$ norms. For instance, the *non-overlapping* group $l_{0,2}$ or $l_{0,1}$ operators [52] could be used to enforce an l_0 norm-like structure on the magnitudes of each dipole, allowing for dipoles that do not align with any of the grid directions*. These operators can also be used in the algorithm proposed below but we focus on the l_0 norm in this work both for concreteness and because it satisfies all the requirements for generating sparse, high-quality configurations of permanent magnets. As we detail our algorithm, it should become apparent that the l_1 loss term, the typical choice for a relaxation of the l_0 loss term, is not sufficient for generating sparse and *binary* solutions.

Beyond the connection with sparse regression, it may not be clear yet why this formulation of the permanent magnet optimization problem could be preferable to formulations proposed in previous works. In Section II E we will show that we can use the special structure appearing in our formulation to design a fast, flexible, easy-to-understand solution to placing a sparse set of permanent magnets for stellarator field shaping.

E. Proposed relax-and-split algorithm

An algorithm is required for effectively solving the various proposed objective functions that can be used for permanent magnetic optimization. The backbone of all these objective functions is a linear least-squares problem with a convex regularizer, subject to a large number of separable, spherical constraints. Nonconvex optimization problems with convex constraints can be effectively solved with the relax-and-split formulation [34] of the optimization problem if the proximal operator,

$$\text{prox}_{\alpha N}(\mathbf{m}) \equiv \arg \min_{\mathbf{y}} \left[\frac{1}{2} \|\mathbf{m} - \mathbf{y}\|_2^2 + \alpha N(\mathbf{y}) \right], \quad (11)$$

of the nonconvex (and/or nonsmooth) part of the objective function, $N(\mathbf{m})$, is known or easily approximated. In the present work, it is assumed that an analytic expression is known for the proximal operator (as is the case for the l_0 , l_1 , non-overlapping group $l_{0,1}$, etc.), which significantly simplifies our task and additionally implies that, despite the notation, $N(\mathbf{m})$ is not an arbitrary function. There are also nonconvex loss terms without known analytic proximal operators that nonetheless can be rapidly numerically computed.

To see how the proximal operator enters into the optimization, we first introduce a proxy variable $\mathbf{w} \sim \mathbf{m}$. This proxy is kept reasonably close to the values in \mathbf{m} by the inclusion

* Consider a set of subsets $\{G_1, G_2, \dots\}$ such that the union of the subsets is the entire index set of the optimization variables, $G_1 \cup G_2 \cup \dots = \{1, \dots, D\}$. The group $l_{0,1}$ norm is defined as $\|\|\mathbf{m}\|_{G_1}^{G_1}, \|\mathbf{m}\|_{G_2}^{G_2}, \dots, \|\mathbf{m}\|_0\|_0$. For overlapping groups, this can be a very complicated operator. For the permanent magnet problem, the groups are the non-overlapping triplets consisting of the components of each dipole vector. In the non-overlapping case, these operators can be addressed like the l_0 is addressed in Appendix B.

of a new least-squares term $\mathbf{v}^{-1} \|\mathbf{m} - \mathbf{w}\|_2^2$, with the value of the hyperparameter \mathbf{v} determining how closely the two variables should match. The goal is now to solve the following optimization problem,

$$\arg \min_{\mathbf{w}, \mathbf{m} \in \mathbb{S}_m} \left[\frac{\|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2^2}{2} + \frac{\|\mathbf{m} - \mathbf{w}\|_2^2}{2\mathbf{v}} + \lambda R(\mathbf{m}) \right] + \alpha N(\mathbf{w}). \quad (12)$$

We can solve this problem effectively by using variable projection [53]. We solve the inner optimization problem over \mathbf{m} at fixed \mathbf{w} , then solve the outer optimization over \mathbf{w} at fixed \mathbf{m} , and iterate between these solves until convergence is found. The inner optimization problem is convex over the original dipole variables in \mathbf{m} . Fix the initial guess for $\mathbf{w}^{(0)}$ and denote the solution after the first convex iteration

$$\mathbf{m}^{(1)} \equiv \arg \min_{\mathbf{m} \in \mathbb{S}_m} \left[\frac{\|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2^2}{2} + \frac{\|\mathbf{m} - \mathbf{w}^{(0)}\|_2^2}{2\mathbf{v}} + \lambda R(\mathbf{m}) \right], \quad (13)$$

so that the remaining optimization problem is

$$\begin{aligned} \mathbf{w}^{(1)} &\equiv \arg \min_{\mathbf{w}} \left[\frac{1}{2\mathbf{v}} \|\mathbf{m}^{(1)} - \mathbf{w}\|_2^2 + \alpha N(\mathbf{w}) \right] \\ &= \text{prox}_{\mathbf{v}\alpha N}(\mathbf{m}^{(1)}). \end{aligned} \quad (14)$$

In other words, the remaining optimization problem is equivalent to computing the proximal operator. The proximal operator for the l_0 norm is *hard-thresholding*, detailed briefly in Appendix B. The hard thresholding naturally generates sparse vectors. Now we can iterate between the convex and nonconvex optimization subproblems so that the full algorithm can be summarized in Algorithm 1.

Algorithm 1 Relax-and-split with the l_0 loss term

Input: Magnet geometry encoded in \mathbf{A} , coil and plasma magnetic fields in \mathbf{b} , initial guess $\mathbf{w}^{(0)}$, error tolerance ϵ , initial hard threshold α_0 , and strength of any additional regularization via λ .

Output: Solutions \mathbf{m}^* and \mathbf{w}^* .

procedure RELAX-AND-SPLIT(\mathbf{A} , \mathbf{b} , λ , \mathbf{v} , α_0)

for $\alpha = \alpha_0, 2\alpha_0, \dots$

for $k = 0, \dots, k_{max}$

$$\mathbf{m}^{(k+1)} \equiv \arg \min_{\mathbf{m} \in \mathbb{S}_m} \left[\frac{\|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2^2}{2} + \frac{\|\mathbf{m} - \mathbf{w}^{(k)}\|_2^2}{2\mathbf{v}} + \lambda R(\mathbf{m}) \right]$$

$$\mathbf{w}^{(k+1)} = \text{prox}_{\mathbf{v}\alpha\|\cdot\|_0}(\mathbf{m}^{(k+1)})$$

$$\text{if } \|\mathbf{w}^{(k+1)} - \mathbf{w}^{(k)}\|_2 \leq \epsilon$$

break

$$\mathbf{w}^{(0)} = \mathbf{w}^{(k+1)}$$

return $\mathbf{m}^{(k+1)}, \mathbf{w}^{(k+1)}$

end procedure

In other words, for each value of the hard threshold $\propto \sqrt{\alpha_0}$, solve the convex part of the optimization, take the proximal operator of the resulting vector, and iterate until convergence is found. Typically $\mathbf{v} \sim \|\mathbf{A}\|_2^{-2}$ works well, so λ and α can be considered the primary hyperparameters. λ tunes the strength of the regularization and α tunes the strength of the nonconvexity.

We found empirically that it often works very well to start with weak values of the hard threshold used in the l_0 norm (the threshold is proportional to $\sqrt{\alpha}$), converge this problem, and then use the solution as an initial guess to a new problem with a larger value of α . This is repeated until α is large enough to threshold off essentially all of the magnets that are not maximum strength. Since the magnets are all essentially maximum strength by the final optimization loop, they are also all grid-aligned by virtue of the l_0 loss term. To see this, note that the algorithm thresholds off components of each \mathbf{m}_i that are below some minimum strength, subject to the constraints on the maximum dipole magnitudes. For instance, if components of \mathbf{m}_i with below 97.5% of the maximum strength are thresholded, it follows that this dipole must either be exactly zero or have a single nonzero vector component. As far as we are aware, this iterative approach has not been used before in the field of sparse regression.

There are some significant advantages of Algorithm 1. The convex and nonconvex parts of the optimization problem can both be efficiently solved, additional convex equality and inequality constraints can be added straightforwardly, and the requirements on \mathbf{m} are split between \mathbf{m} and the proxy variable \mathbf{w} , allowing the algorithm to “relax” into satisfying the constraints. The formulation as sparse regression allows for the rapid adoption of concepts, algorithms, and expertise from the prolific field of sparse regression.

Moreover, if the problem is reasonably well-posed (e.g. Tikhonov regularized), the initial permanent magnet configuration is unimportant – whatever it is, it is erased by the convex optimization in the first step, fully circumnavigating part of the initial guess issues present in the other nonconvex algorithms for permanent magnet optimization. It may appear at first that this insensitivity to the initial condition may cause the algorithm to only explore a reduced space of configurations. But we now also have additional hyperparameters, and by sweeping the value of these hyperparameters, we have a systematic way to explore a very wide space of permanent magnet configurations.

III. RESULTS

We now present optimized permanent magnet configurations for the NCSX and MUSE stellarators. Both examples are run with a set of high-resolution quadrature points on the unique part of the plasma boundary – 64 points in θ , and 64 points in ϕ ($\times 2N_{fp}$ for all N_{fp} field periods). Both of the following considered stellarators are stellarator-symmetric and field-period symmetric, so it is only required to design dipoles for $1/2N_{fp}$ of the toroidal angle extent, and then to repeat this configuration around the torus.

A. NCSX

NCSX was a planned quasi-axisymmetric stellarator that was partially built at the Princeton Plasma Physics Laboratory.

It was originally designed with 18 modular coils and 18 planar coils. The equilibrium of interest, C09R00, was scaled to have an on-axis magnetic field strength of 0.5 T, which is the maximum field produced by the existing planar coils. C09R00 also exhibits a three-fold field symmetry, major radius of 1.44 m, minor radius of 0.32 m, and volume-averaged plasma beta $\langle\beta\rangle = 4.09\%$. However, for direct comparison with the FAMUS solution in [16], the C09R00 shape is used but with no plasma current, i.e. $\langle\beta\rangle = 0$, and the toroidal field was taken to be perfectly toroidal with no ripple.

For a direct comparison with the FAMUS [14] method, Algorithm 1 is used in a cylindrical coordinate system with the l_0 norm and Tikhonov regularization is omitted. Tikhonov regularization was simply not required for a high-quality solution (and the relax-and-split term, $\propto \|\mathbf{m} - \mathbf{w}\|_2^2$, is an additional source of regularization).

The FAMUS grid of allowed dipole locations has a resolution of 14 points radially, so $\mathbf{A} \in \mathbb{R}^{4096 \times 172032}$. We use the same grid for our optimizations in SIMSOPT. FAMUS is used with a level of regularization for a plausible solution $f_B \approx 1.6 \times 10^{-6} \text{ T}^2\text{m}^2$. An effective volume of the permanent magnet region can be defined by

$$V_{\text{eff}} = \sum_{j=1}^D \frac{\|\mathbf{m}_j\|_2}{M_0}, \quad M_0 \equiv \frac{B_{\text{rem}}}{\mu_0}, \quad (15)$$

and the result is 2.29 m^3 out of a maximum available volume of 3.23 m^3 . To measure how binary a solution is, we also define the binary fraction

$$f_\delta = 1 - \frac{\#\{i|\delta \leq \|\mathbf{m}_i\|_2 \leq 1 - \delta\}}{D}. \quad (16)$$

There is no post-processing optimization done to the FAMUS solution, so there is significant “pile-up” near dipole magnitudes at 0 and 1 and $f_{0.01} = 0.57$.

A representative SIMSOPT \mathbf{m}^* solution achieves a similar $f_B \approx 1.6 \times 10^{-6} \text{ T}^2\text{m}^2$ with a very similar effective volume of 2.34 m^3 and a substantially improved $f_{0.01} = 0.84$. This improvement with SIMSOPT occurs despite the fact that the relax-and-split method also produces a cylindrically grid-aligned solution. This solution is compared with the FAMUS solution in Figures 2 and 3. The \mathbf{w}^* solution achieves a poor value of f_B , but nonetheless still successfully pulls \mathbf{m}^* towards a high-quality solution that is approximately sparse. The relax-and-split and FAMUS solutions have some qualitative similarities in the strength of the permanent magnets. It is noteworthy that in Fig. 3, both solutions appear to require a set of weak magnets for low-error configurations of this NCSX example. This lends some evidence to the following propositions: that the permanent magnet grid choice can be reasonably important for finding high-performance solutions, and that the low-errors in f_B are quite sensitive to small changes in the magnets. Overall, the major takeaway is that we have found a relax-and-split solution that provides similar f_B performance as the FAMUS solution, despite being significantly sparser and grid-aligned.

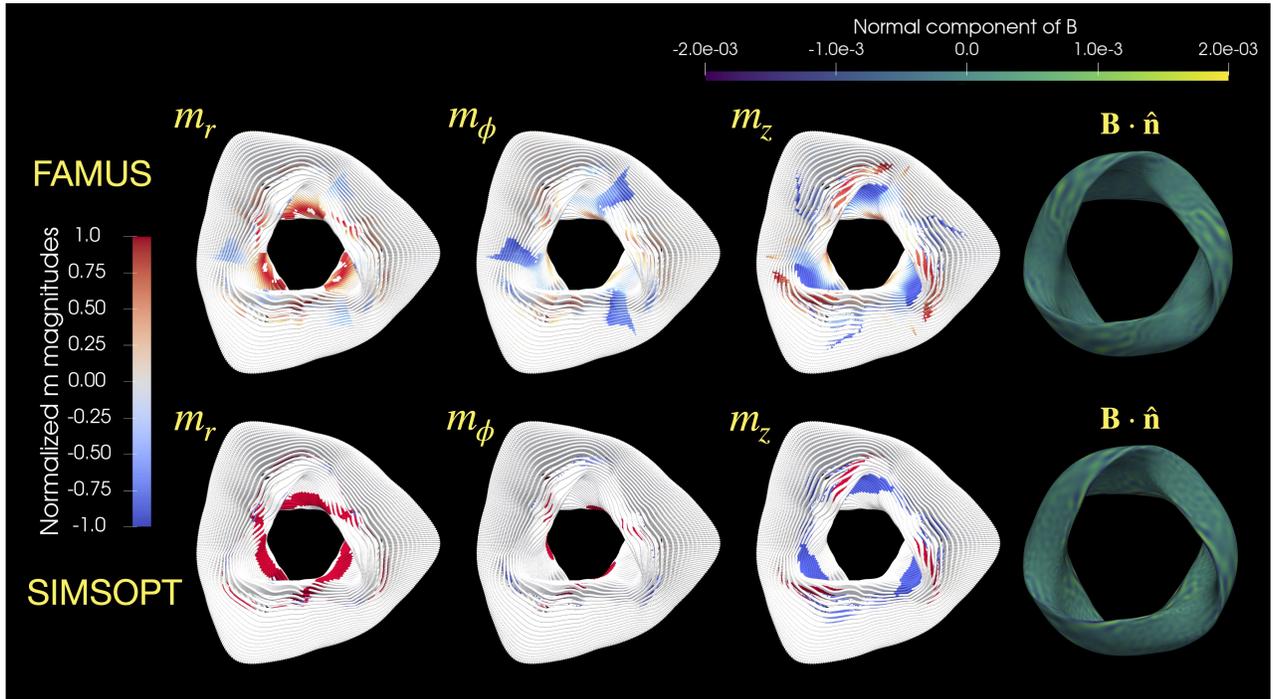


FIG. 2: Comparison between FAMUS and relax-and-split solutions on an example NCSX permanent magnet manifold. Note that the relax-and-split dipoles are aligned with the cylindrical grid.

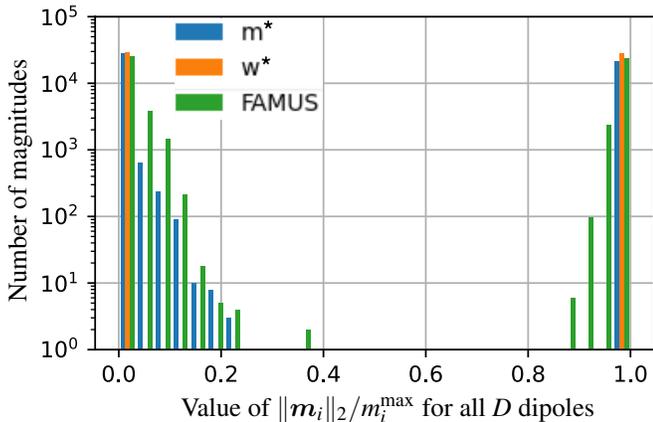


FIG. 3: Distributions of the dipole magnitudes for the relax-and-split solutions and FAMUS. Significant numbers of weak magnets seem to be required for low-error configurations with this NCSX permanent magnet manifold.

B. The MUSE stellarator

MUSE is a table-top stellarator experiment using permanent magnets that is currently under construction [19]. MUSE was optimized for a high degree of quasi-symmetry, and the experiment’s permanent magnet configuration was optimized to have good flux surfaces with the FAMUS code [12] to $f_B \approx 5.17 \times 10^{-8} \text{ T}^2 \text{ m}^2$. The dipoles were constrained to point only in the minor radial direction in simple toroidal coordinates. Diagnostic ports and other spatial restrictions are

used to represent the real permanent magnet configuration. Spatial restrictions only eliminate a portion of the grid from consideration, and importantly do not change the properties of the optimization problem. After the permanent magnets are optimized, the combined magnetic field from the magnets and coils are used to compute Poincaré plots in order to verify the quality of the flux surfaces. Substantial discrete optimization was used in post-processing steps to achieve engineering constraints while preserving the physics objectives, as described in Qian et al. [19].

We use the same stage-1 optimized plasma surface, the same 16 planar TF coils, and the same MUSE permanent magnet grid array of four toroidal quadrants [19] (i.e. the dipole locations are identical in the FAMUS and SIMSOPT optimizations) with a simple toroidal coordinate system. The grid provides for the inclusion of four vertical ports and twelve horizontal ports. We solve both the convex and nonconvex optimization problem from (5) and (10) and compare with the FAMUS solution. For these optimizations, the number of dipoles is 75,460 ($\times 4$ via symmetries), each with three vector components, so $\mathbf{A} \in \mathbb{R}^{4096 \times 226380}$. Tikhonov regularization is used with strength $\lambda = 10^{-8}$ for the convex problem; these results are identical to the last two rows of Fig. 1. Tikhonov regularization is omitted for the SIMSOPT solution with the l_0 term included, for the same reason as in Sec. III A.

The results comparing the convex optimization in (5), the FAMUS solution, and the relax-and-split solutions \mathbf{m}^* and \mathbf{w}^* are shown in Fig. 4. The toroidal field coils are shown in green, the normalized permanent magnet strength is indicated in the blue-red colorbar, and the normal component of

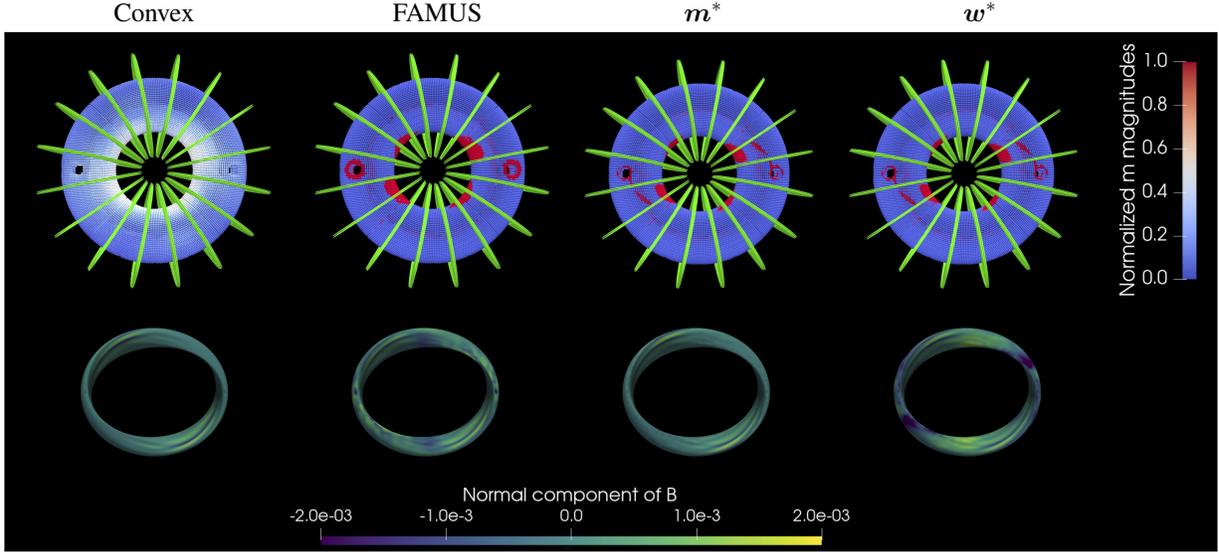


FIG. 4: Comparison between the convex algorithm, the FAMUS solution, and the relax-and-split solutions for MUSE. The sparse relax-and-split w^* produces slightly larger f_B error than FAMUS, but uses fewer magnets.

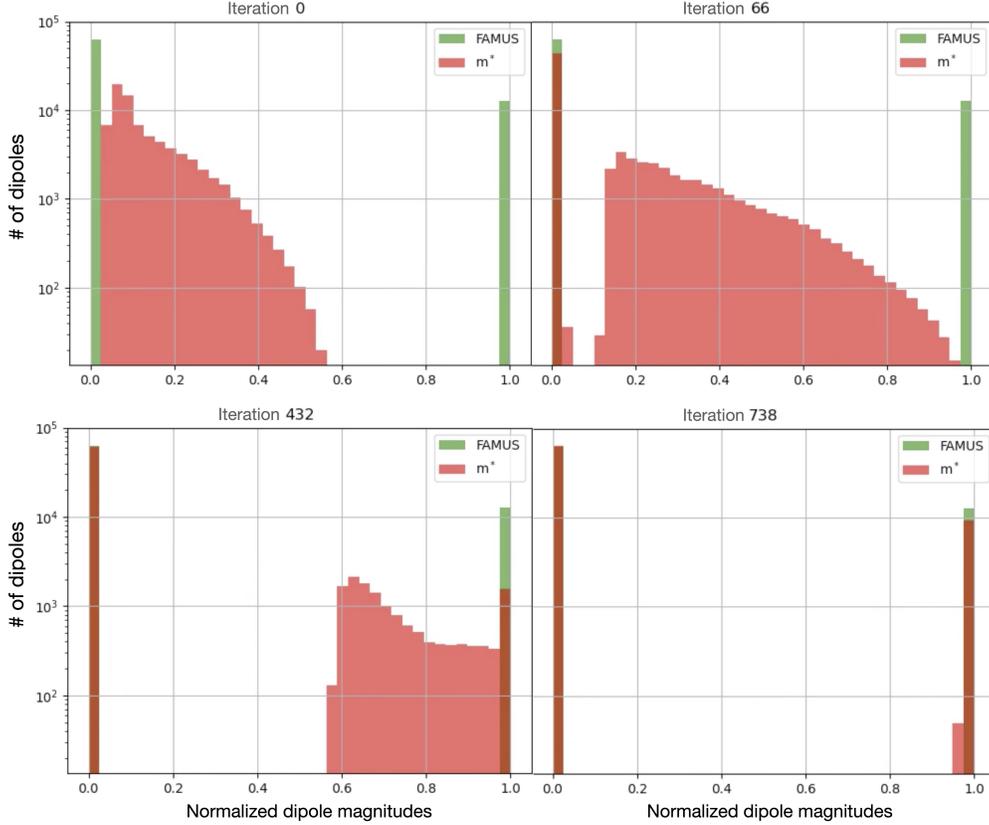


FIG. 5: Histograms of the dipole magnitudes as the algorithm progresses (red) compared with the optimized FAMUS solution for MUSE (green). Slowly increasing the thresholding has the effect of pushing all the magnets to magnitudes of zero or one.

the residual field errors is illustrated in the purple-yellow colorbar. The optimization began by truncating magnets with strengths below 5% and the threshold parameter (proportional to α_0) was increased until magnets below 97.5% strength were truncated. The algorithm progress is visualized in Figures 5

and 6, showing that the optimization successfully transforms a nonsparse solution of weak dipoles to a very sparse solution of maximum-strength dipoles.

The m^* solution achieves $f_B \approx 2.8 \times 10^{-8}$ and the w^* solution achieves $f_B \approx 1.5 \times 10^{-7}$, using nonzero dipole magni-

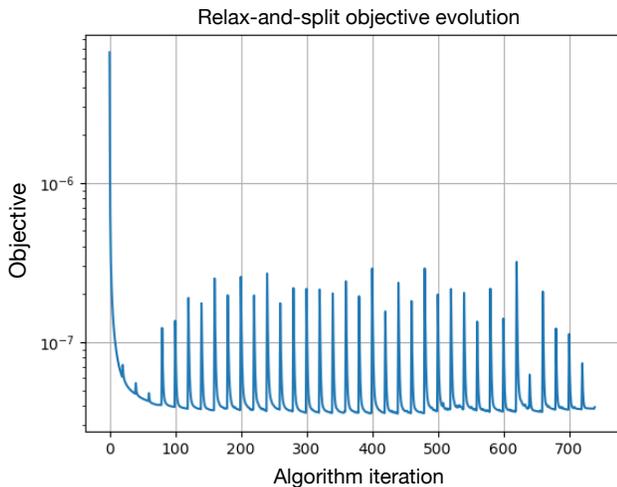


FIG. 6: Illustration of the optimization objective as the relax-and-split algorithm progresses. Increases in the hard threshold causes sudden spikes in the objective value, but the algorithm quickly recovers to pre-thresholding error levels. At the end of the algorithm, when only magnets of strength 97.5% remain in the MUSE permanent magnet manifold, the error is roughly equal to the error at the beginning, when dipole magnitudes could vary between 0 and 1.

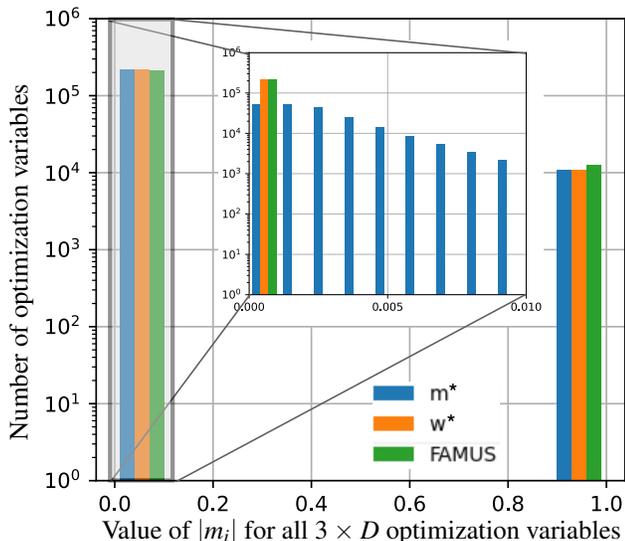


FIG. 7: Distributions of all the permanent magnet optimization variables for the relax-and-split and FAMUS solutions for MUSE. The vast majority of the dipoles are zeroed and relax-and-split uses fewer full-strength magnets. A zoomed view illustrates the magnet “stragglers” in m^* , while w^* fully zeroes out magnets below the hard threshold.

tudes for only 14.6% of the permanent magnets. It is interesting to note that w^* produces five times the f_B value compared to m^* . Yet the m^* and w^* distributions appear at first to be virtually identical in Fig. 7, and in 3D visualization as in the top row of Fig. 4. Zooming into the origin of Fig. 7, we can see that m^* actually breaks the grid-alignment of the dipoles in w^* , with a large number of very weak components. Summing

over the absolute values of all of these very weak but nonzero components actually amounts to the equivalent contribution of approximately 600 ($\times 4$ for symmetries) full-strength, grid-aligned permanent magnets (modulo the important directionality but this is a rough estimate anyways). This “straggler” issue is seen elsewhere in sparse regression, and in fact motivates the post-processing and post-optimization that is performed to improve FAMUS permanent magnet solutions. It is a substantial strength that w^* does not suffer from this issue, but it comes at the cost of increased f_B error. Moreover, this sensitivity to very weak, nonzero components in the otherwise grid-aligned dipoles motivates a stochastic optimization approach, outlined in Appendix C, in order to search for configurations that are robust to small changes in the solution.

FAMUS uses approximately 16.8% of the permanent magnets, or exactly 6608 additional magnets* with respect to the relax-and-split solution w^* . Additional hyperparameter tuning could be done to get a comparison with FAMUS using the exact same number of dipoles, but the relax-and-split solution is already sparse, high-quality, and therefore suitable enough for this example configuration. The extra FAMUS post-processing precludes a perfect comparison with the relax-and-split method anyways. Moreover, as is illustrated in Fig. 8, the FAMUS solution constrains the dipoles to point in the minor radial direction, while the relax-and-split solution produces dipoles that are aligned with one of the three simple toroidal directions. Despite these extra degrees of freedom, the solution qualitatively reproduces the dipole structures seen in the FAMUS solution. In contrast, the relax-and-split solution with Cartesian grid-aligned dipoles in Fig. 8 produces another excellent solution but looks qualitatively quite different. The relax-and-split and FAMUS solutions for both examples are summarized in Table I and Poincaré plots for each of the solutions are illustrated in Fig. 9. The surfaces generated with m^* look slightly improved over those of FAMUS but the w^* surfaces look slightly degraded.

Consider the remarkable fact that the FAMUS and relax-and-split toroidally-aligned solutions qualitatively match. The two algorithms have different degrees of freedom and solve very different nonconvex optimization problems, yet the solutions are qualitatively very similar. There are two effects that seem like plausible reasons for this convergence between algorithms. First, certain grid locations may be essential for properly minimizing f_B , e.g. locations near highly-shaped magnetic fields. Second, the requirement of maxi-

*The experimental MUSE device actually uses 9,736 magnet “towers”. The total magnet count is much lower than is indicated by the grid of 75, 460 \times 4 locations because the final permanent magnet configuration for MUSE is generated by concatenating together identical toroidal slices (each modeled as an ideal dipole) into towers. For instance, if there are 14 slices next to each other, each 1/16” thick, a single 7/8” thick magnet is used instead. Future work could investigate promoting solutions with large, continuous groups of magnets so that this concatenation can be done effectively. Here, the FAMUS and relax-and-split solutions are compared by quoting the difference in the total number of permanent magnets with reference to the original grid locations. The most intuitive metric might be the effective permanent magnet volume in Eq. (15). The relax-and-split solution uses 13% less effective permanent magnet volume than the FAMUS solution.

Stellarator	Solution	f_B	$V_{\text{eff}} \text{ (m}^3\text{)}$	$f_{0.01}$	Grid-alignment
NCSX	FAMUS	1.6×10^{-6}	2.29	0.57	None
NCSX	SIMSOPT m^*	1.6×10^{-6}	2.34	0.84	cylindrical
NCSX	SIMSOPT w^*	4.7×10^{-4}	2.30	1	cylindrical
MUSE	FAMUS	5.7×10^{-8}	3.245×10^{-3}	1	r_{minor}
MUSE	SIMSOPT m^*	2.8×10^{-8}	2.906×10^{-3}	0.92	toroidal
MUSE	SIMSOPT w^*	1.5×10^{-7}	2.822×10^{-3}	1	toroidal

TABLE I: Summary of the comparison between FAMUS and relax-and-split on the NCSX and MUSE examples.

imum strength, binary magnets effectively regularizes much of the permanent magnet optimization space. In other words, there may be a vast number of permanent magnet configurations that minimize f_B to high-performance levels, but far fewer such configurations have binary distributions.

IV. CONCLUSION

We have shown that permanent magnet optimization can be formulated as sparse regression. This scientific problem is common to many fields, including automobile manufacturing, MRI, ship de-gaussing, and many other applications. We have classified a number of loss terms that can be effectively used in permanent magnet optimization, and proposed a relax-and-split algorithm that takes advantage of the partial convexity, addresses an important class of nonsmooth and nonconvex optimization loss terms, and allows for additional convex equality and inequality constraints to be added to future objective functions. The algorithm is generally applicable, and our implementation is computationally efficient and open-source. It can be immediately used or extended for winding-surface stellarator coil optimization and high-dimensional sparse regression across many scientific domains. We concluded by discovering accurate and sparse solutions for two stellarators. Interestingly, the relax-and-split algorithm qualitatively matches the FAMUS solution for the MUSE stellarator example, despite the significant differences between the algorithms and their associated optimization problems. This may be an indication that the requirement of sparse, binary magnet solutions may regularize the optimization enough to leave only a few high-quality minima, despite the nonconvexity of the problem.

Future work should further explore the parameter space of possible permanent magnet configurations on a variety of stage-1 plasma equilibria. More sophisticated physical and engineering objectives can be built into the optimization by modeling the permanent magnet deviation from a dipole field, adding constraints for the demagnetization upper bound in Eq. (6), calculating the errors induced by the permeability of the permanent magnet material, and using stochastic optimization as outlined in Appendix C. Soft rather than hard constraints on the dipole magnitudes might allow for a significantly faster convex step in the iterative relax-and-split algorithm, although we found empirically that often high-quality configurations can be computed without requiring that full convergence is achieved during each convex solve.

Addressing the remaining optimization questions is an important task – beyond this barrier, permanent magnet stellarators would be ideal laboratory experiments that can be constructed easily and inexpensively, while providing vital small-scale stellarator insights for informing the design of full-scale stellarators that can have construction costs in the billions of dollars.

ACKNOWLEDGEMENTS

AAK would like to acknowledge important conversations with Christopher Hansen, Aleksandr Aravkin, and Steven Brunton. Thanks to Caixiang Zhu for help with FAMUS. This work was supported by the U.S. Department of Energy under award DEFG0293ER54197 and through a grant from the Simons Foundation (560651, ML). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

Appendix A: Least-squares form of f_B

Here we derive how the f_B term can be put into a least-squares form, following Zhu et al. [14] but with some minor changes. First, we define the geometric factor,

$$g_i(\phi_{i_\phi}, \theta_{i_\theta}) = \frac{\mu_0}{4\pi} \left(\frac{3\mathbf{r}_i \cdot \hat{\mathbf{n}}}{\|\mathbf{r}_i\|_2^5} \mathbf{r}_i - \frac{\hat{\mathbf{n}}}{\|\mathbf{r}_i\|_2^3} \right) \sqrt{\Delta\phi_{i_\phi} \Delta\theta_{i_\theta} \|\mathbf{n}\|_2}. \quad (\text{A1})$$

Here $(\phi_{i_\phi}, \theta_{i_\theta})$ define a set of quadrature points on the plasma surface in toroidal coordinates, (i_ϕ, i_θ) index these points, $(\Delta\phi_{i_\phi}, \Delta\theta_{i_\theta})$ denote the grid spacing, and \mathbf{n} is shorthand for the normal vector at these locations, $\mathbf{n}_{i_\phi, i_\theta}$. The number of points in each direction are N_ϕ and N_θ . Stack all of these factors together for an overall matrix $\mathbf{A} \in \mathbb{R}^{N_\phi N_\theta \times 3D}$,

$$\begin{aligned} \mathbf{A}(\phi_{i_\phi}, \theta_{i_\theta}) &= [\mathbf{g}_1 \ \dots \ \mathbf{g}_D] & (\text{A2}) \\ &= [\mathbf{g}_1^x \ \mathbf{g}_1^y \ \mathbf{g}_1^z \ \dots \ \mathbf{g}_D^x \ \mathbf{g}_D^y \ \mathbf{g}_D^z] \\ &= \begin{bmatrix} g_1^x(\phi_1, \theta_1) & g_1^y(\phi_1, \theta_1) & g_1^z(\phi_1, \theta_1) & \dots \\ \vdots & \vdots & \vdots & \\ g_{N_\theta}^x(\phi_{N_\theta}, \theta_{N_\theta}) & g_{N_\theta}^y(\phi_{N_\theta}, \theta_{N_\theta}) & g_{N_\theta}^z(\phi_{N_\theta}, \theta_{N_\theta}) & \dots \end{bmatrix}. \end{aligned}$$

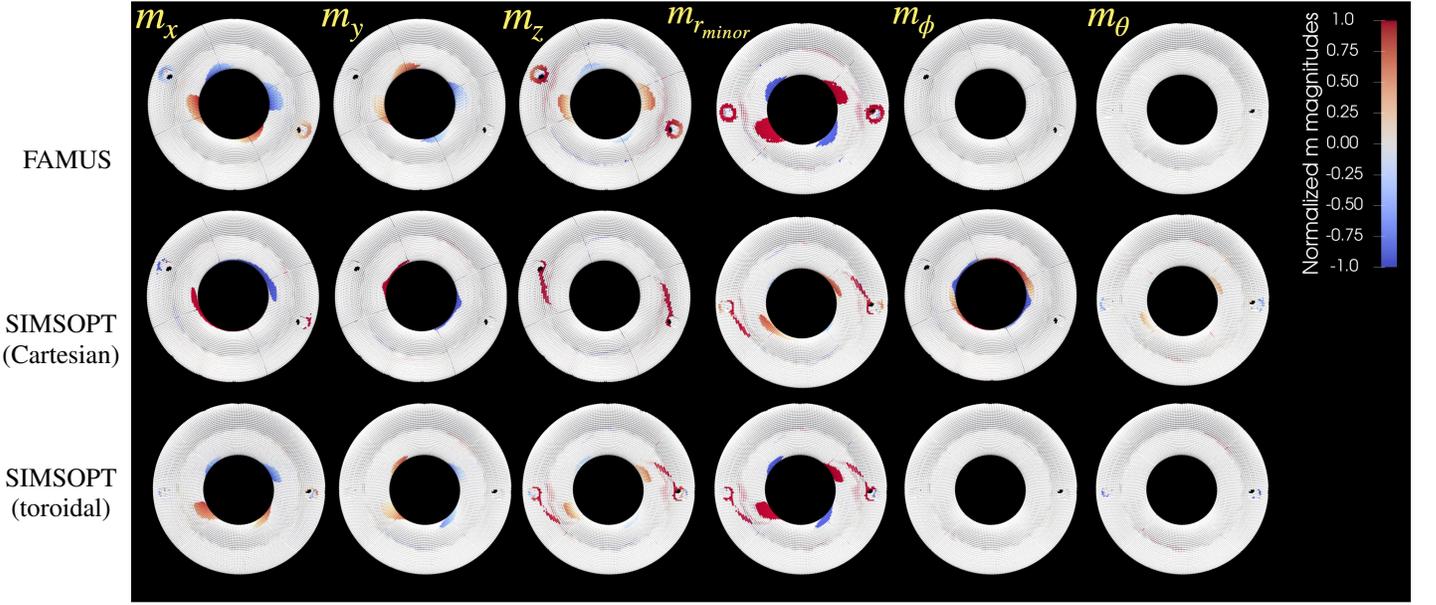


FIG. 8: The different components of the MUSE \mathbf{m} solutions in Cartesian and simple toroidal coordinates illustrate the different grid-alignments used in FAMUS (top row) and in SIMSOPT (middle and bottom rows are the relax-and-split dipole solutions grid-aligned with Cartesian coordinates and grid-aligned with simple toroidal coordinates, respectively). The FAMUS magnets are constrained to point inwards or outwards in the minor radial direction, and in the toroidal relax-and-split case, magnets point in one of the three simple toroidal directions yet qualitatively reproduce the magnet structures seen by FAMUS.

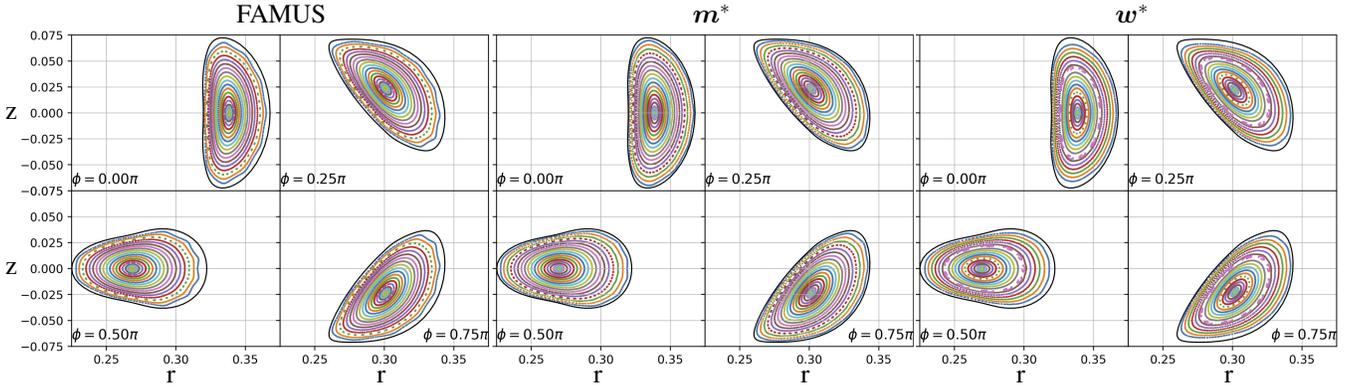


FIG. 9: Poincaré plots for MUSE solutions with the target plasma boundary in black.

This factor encodes all of the geometry, is independent of the dipoles \mathbf{m}_i , and can be computed once before optimization begins. To see why this is a useful definition, note that if $\mathbf{m} = [m_1^x, m_1^y, m_1^z, m_2^x, \dots, m_D^z]$, we can write,

$$\begin{aligned} f_B &= \int ((\mathbf{B}_M + \mathbf{B}_P + \mathbf{B}_C) \cdot \hat{\mathbf{n}})^2 d^2x, \\ &= \sum_{i_\phi} \sum_{i_\theta} \Delta\phi_{i_\phi} \Delta\theta_{i_\theta} (\mathbf{B}_M \cdot \hat{\mathbf{n}} - B_n)^2 \|\mathbf{n}\|_2, \end{aligned} \quad (\text{A3})$$

where the normal component of the coil plus plasma magnetic fields (on the plasma surface) is denoted $-B_n \in \mathbb{R}^{N_\phi N_\theta}$ here. Then convert the double sum to a single sum with the relabel-

$$\text{ing } i = i_\phi + N_\phi(i_\theta - 1),$$

$$\begin{aligned} f_B &= \sum_{i_\phi} \sum_{i_\theta} \left((\mathbf{B}_M \cdot \hat{\mathbf{n}} - B_n) \sqrt{\Delta\phi_{i_\phi} \Delta\theta_{i_\theta} \|\mathbf{n}\|_2} \right)^2 \\ &= \sum_i (\mathbf{A}_{ij} m_j - b_i)^2 = \|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2^2. \end{aligned} \quad (\text{A4})$$

Therefore, f_B can be written as a least-squares term in the \mathbf{m} vector representing the optimization variables. We can enforce field-period and stellarator symmetries into these equations while using only the original \mathbf{m} representing a half-period of the permanent magnet manifold. The coordinate system for the permanent magnet manifold is chosen to inherit the symmetries. For \mathbf{B}_M to exhibit field-period symmetry, \mathbf{m} is rotated so that the angle between \mathbf{m} and \mathbf{r}_i is preserved at

the new point.

If both coordinate systems are also stellarator symmetric, it follows that the new vector $\mathbf{r}'_i = [x_i, -y_i, -z_i]$ is the transformation of \mathbf{r}_i under the symmetry. If each \mathbf{m}_i is chosen to be stellarator symmetric, then $\mathbf{m}'_i = [-m_x, m_y, m_z]$ and this is sufficient for \mathbf{B}_M to be stellarator symmetric.

The transformation of \mathbf{A} is the inverse of the transformation of \mathbf{m} . For instance, if the new vector \mathbf{m}'_i is the result of flipping the m_r component from a reflection \mathbf{R}_S via stellarator symmetry, followed by a rotation \mathbf{R}_{fp} by the appropriate angle for field-period symmetry, then \mathbf{A} should be rotated by the negative angle, and then the resulting vector should be transformed by the stellarator symmetry. In total then, the least-squares term $\mathbf{A}\mathbf{m}$ that is actually used in the code is

$$\left[\mathbf{A}(\mathbf{r}_i) + \mathbf{A}(\mathbf{R}_S \mathbf{r}_i) \mathbf{R}_S^T + \mathbf{A}(\mathbf{R}_{fp} \mathbf{r}_i) \mathbf{R}_{fp}^T + \dots \right] \mathbf{m}, \quad (\text{A5})$$

which now is calculating the contributions of all of the magnets (around the full torus) to a half-period of the plasma surface, while only using the half-period variables \mathbf{m} .

Appendix B: The l_0 proximal operator

Proximal operators are useful tools for optimization over nonsmooth and/or nonconvex terms that otherwise have relatively simple structure. The proximal operator is defined as,

$$\text{prox}_{\alpha N}(\mathbf{m}) \equiv \arg \min_{\mathbf{y}} \left[\frac{1}{2} \|\mathbf{m} - \mathbf{y}\|_2^2 + \alpha N(\mathbf{y}) \right]. \quad (\text{B1})$$

The advantage of this formulation is that it can be analytically or rapidly numerically computed for a number of common nonconvex functions. The proximal operator of the l_0 “norm” is well-known as *hard-thresholding*,

$$\text{prox}_{\alpha \|\cdot\|_0}(\mathbf{m}) = \mathcal{H}_{\sqrt{2\alpha}}(m_1) \times \dots \times \mathcal{H}_{\sqrt{2\alpha}}(m_{3D}), \quad (\text{B2})$$

$$\mathcal{H}_{\sqrt{2\alpha}}(m_i) = \begin{cases} 0 & |m_i| < \sqrt{2\alpha} \\ m_i & |m_i| \geq \sqrt{2\alpha} \end{cases}. \quad (\text{B3})$$

In practice, we actually hard threshold the normalized \mathbf{m} (note that the l_0 norm is invariant to a global rescaling of all the \mathbf{m} components before optimization begins), since otherwise the hard thresholding will preferentially remove small magnets, even if those magnets are at full strength and very

important to the solution. The maximum strength of the permanent magnets is usually proportional in some way to the geometry (for instance, directly proportional to the cylindrical radius), and therefore an unnormalized hard-thresholding would tend to remove the geometrically-small magnets first.

Appendix C: Stochastic optimization with permanent magnets

It may not be immediately clear that the relax-and-split algorithm can be extended for more advanced optimization techniques such as stochastic optimization. Stochastic optimization is a useful technique for generating configurations that are robust to random and systematic errors in the coil shapes (traditional coil optimization) and the permanent magnet magnitudes and directions (permanent magnet optimization). In this section, we briefly illustrate that typical forms of stochastic optimization preserve all the loss term structure needed for relax-and-split. Consider a vector of independent, normally-distributed random variables $\boldsymbol{\xi}$ as in Wechsung et al. [54]. The goal is to minimize the expectation \mathbb{E} of the $\mathbf{B} \cdot \hat{\mathbf{n}}$ errors over the distribution of perturbations $\boldsymbol{\xi}$,

$$\min_{\mathbf{m}} \mathbb{E}(f_B(\mathbf{m} + \boldsymbol{\xi})) + \lambda R(\mathbf{m}) + \alpha N(\mathbf{m}). \quad (\text{C1})$$

Alternatively, one can optimize for the worst case scenario,

$$\min_{\mathbf{m}} \max_{\boldsymbol{\xi}} \mathbb{E}(f_B(\mathbf{m} + \boldsymbol{\xi})) + \lambda R(\mathbf{m}) + \alpha N(\mathbf{m}). \quad (\text{C2})$$

The point of the regularization terms is to simplify the manufacturing, so they are relevant *before* any perturbations are introduced. For instance, if a dipole has zero magnitude, a magnet will not be placed, and therefore there are no manufacturing errors to consider at this location. The expectation value can be approximated using the sample average approximation; drawing S independent realizations of $\boldsymbol{\xi}$ and using,

$$\mathbb{E}(f_B(\mathbf{m}, \boldsymbol{\xi})) \approx \frac{1}{S} \sum_{s=1}^S f_B(\mathbf{m} + \boldsymbol{\xi}_s), \quad (\text{C3})$$

to achieve $\mathcal{O}(S^{-\frac{1}{2}})$ approximation error. The $\boldsymbol{\xi}_s$ are fixed, so (C1) and (C2) are optimizations using deterministic sums of the original objectives. Sums of convex terms are also convex, and thus relax-and-split can be applied in this setting.

-
- [1] G. Grieger, W. Lotz, P. Merkel, J. Nührenberg, J. Sapper, E. Strumberger, H. Wobig, R. Burhenn, V. Erckmann, U. Gasparino, *et al.*, Physics optimization of stellarators, *Physics of Fluids B: Plasma Physics* **4**, 2081 (1992).
- [2] L.-M. Imbert-Gerard, E. J. Paul, and A. M. Wright, An introduction to stellarators: From magnetic fields to symmetries and optimization, arXiv preprint arXiv:1908.05360 (2019).
- [3] M. Drevlak, C. Beidler, J. Geiger, P. Helander, and Y. Turkin, Optimisation of stellarator equilibria with ROSE, *Nuclear Fu-*

sion **59**, 016010 (2018).

- [4] S. Lazerson, J. Schmitt, C. Zhu, J. Breslau, S. Developers, *et al.*, *Stellopt*, Tech. Rep. (Princeton Plasma Physics Lab.(PPPL), Princeton, NJ (United States), 2020).
- [5] M. Landreman, B. Medasani, F. Wechsung, A. Giuliani, R. Jorge, and C. Zhu, SIMSOPT: A flexible framework for stellarator optimization, *Journal of Open Source Software* **6**, 3525 (2021).

- [6] R. Robin and F. A. Volpe, Minimization of magnetic forces on stellarator coils, *Nuclear Fusion* (2022).
- [7] C. Zhu, S. R. Hudson, Y. Song, and Y. Wan, New method to design stellarator coils without the winding surface, *Nuclear Fusion* **58**, 016008 (2017).
- [8] V. Erckmann, H.-J. Hartfuss, M. Kick, H. Renner, J. Sapper, F. Schauer, E. Speth, F. Wesner, F. Wagner, M. Wanner, *et al.*, The W7-X project: Scientific basis and technical realization, in *17th IEEE/NPSS Symposium Fusion Engineering (Cat. No. 97CH36131)*, Vol. 1 (IEEE, 1997) pp. 40–48.
- [9] R. Strykowski, T. Brown, J. Chrzanowski, M. Cole, P. Heitzenroeder, G. H. Neilson, D. Rej, and M. Viol, Engineering cost & schedule lessons learned on NCSX, in *2009 23rd IEEE/NPSS Symposium on Fusion Engineering* (IEEE, 2009) pp. 1–4.
- [10] P. Helander, M. Drevlak, M. Zarnstorff, and S. Cowley, Stellarators with permanent magnets, *Physical review letters* **124**, 095001 (2020).
- [11] J.-P. Wang, Environment-friendly bulk Fe16N2 permanent magnet: Review and prospective, *Journal of Magnetism and Magnetic Materials* **497**, 165962 (2020).
- [12] T. Qian, D. Bishop, A. Chambliss, A. Dominguez, C. Pagano, D. Seidita, M. Zarnstorff, and C. Zhu, Stellarator fields without stellarator coils: MUSE a table top PM stellarator, *Bulletin of the American Physical Society* **66** (2021).
- [13] M. Landreman, An improved current potential method for fast computation of stellarator coil shapes, *Nuclear Fusion* **57**, 046003 (2017).
- [14] C. Zhu, K. Hammond, T. Brown, D. Gates, M. Zarnstorff, K. Corrigan, M. Sibilica, and E. Feibush, Topology optimization of permanent magnets for stellarators, *Nuclear Fusion* **60**, 106002 (2020).
- [15] C. Zhu, M. Zarnstorff, D. Gates, and A. Brooks, Designing stellarators using perpendicular permanent magnets, *Nuclear Fusion* **60**, 076016 (2020).
- [16] M. Landreman and C. Zhu, Calculation of permanent magnet arrangements for stellarators: a linear least-squares method, *Plasma Physics and Controlled Fusion* **63**, 035001 (2021).
- [17] G. Xu, Z. Lu, D. Chen, L. Chen, X. Zhang, X. Wu, M. Ye, and B. Wan, Design of quasi-axisymmetric stellarators with varying-thickness permanent magnets based on Fourier and surface magnetic charges method, *Nuclear Fusion* **61**, 026025 (2021).
- [18] Z. Lu, G. Xu, D. Chen, X. Zhang, L. Chen, M. Ye, H. Guo, and B. Wan, Development of advanced stellarator with identical permanent magnet blocks, *Cell Reports Physical Science* **3**, 100709 (2022).
- [19] T. Qian, M. C. Zarnstorff, D. Bishop, A. Chambliss, A. Dominguez, C. Pagano, D. Patch, and C. Zhu, Simpler optimized stellarators using permanent magnets, *Nuclear Fusion* (2022).
- [20] C. Z. Cooley, M. W. Haskell, S. F. Cauley, C. Sappo, C. D. Lapierre, C. G. Ha, J. P. Stockmann, and L. L. Wald, Design of sparse Halbach magnet arrays for portable MRI using a genetic algorithm, *IEEE transactions on magnetics* **54**, 1 (2017).
- [21] Z. H. Ren, W. C. Mu, and S. Y. 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 (2018).
- [22] M. Landreman and E. Paul, Magnetic fields with precise quasisymmetry for plasma confinement, *Physical Review Letters* **128**, 035001 (2022).
- [23] F. Wechsung, M. Landreman, A. Giuliani, A. Cerfon, and G. Stadler, Precise stellarator quasi-symmetry can be achieved with electromagnetic coils, *Proceedings of the National Academy of Sciences* **119**, e2202084119 (2022).
- [24] R. Jorge, W. Sengupta, and M. Landreman, Construction of quasisymmetric stellarators using a direct coordinate approach, *Nuclear Fusion* **60**, 076021 (2020).
- [25] R. Jorge, W. Sengupta, and M. Landreman, Near-axis expansion of stellarator equilibrium at arbitrary order in the distance to the axis, *Journal of Plasma Physics* **86** (2020).
- [26] R. Jorge and M. Landreman, The use of near-axis magnetic fields for stellarator turbulence simulations, *Plasma Physics and Controlled Fusion* **63**, 014001 (2020).
- [27] M. Landreman and R. Jorge, Magnetic well and mercier stability of stellarators near the magnetic axis, *Journal of Plasma Physics* **86** (2020).
- [28] E. Rodriguez and A. Bhattacharjee, Solving the problem of overdetermination of quasisymmetric equilibrium solutions by near-axis expansions. I. Generalized force balance, *Physics of Plasmas* **28**, 012508 (2021).
- [29] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proceedings of the National Academy of Sciences* **113**, 3932 (2016).
- [30] K. Champion, P. Zheng, A. Y. Aravkin, S. L. Brunton, and J. N. Kutz, A unified sparse optimization framework to learn parsimonious physics-informed models from data, *IEEE Access* **8**, 169259 (2020).
- [31] K. Kaheman, J. N. Kutz, and S. L. Brunton, SINDy-PI: A robust algorithm for parallel implicit sparse identification of nonlinear dynamics, *Proceedings of the Royal Society A* **476**, 20200279 (2020).
- [32] Y. C. Eldar and G. Kutyniok, *Compressed sensing: theory and applications* (Cambridge university press, 2012).
- [33] A. M. Bruckstein, D. L. Donoho, and M. Elad, From sparse solutions of systems of equations to sparse modeling of signals and images, *SIAM review* **51**, 34 (2009).
- [34] P. Zheng and A. Aravkin, Relax-and-split method for nonconvex inverse problems, *Inverse Problems* **36**, 095013 (2020).
- [35] B. de Silva, K. Champion, M. Quade, J.-C. Loiseau, J. Kutz, and S. Brunton, PySINDy: A Python package for the sparse identification of nonlinear dynamical systems from data, *Journal of Open Source Software* **5**, 2104 (2020).
- [36] A. A. Kaptanoglu, B. M. de Silva, U. Fasel, K. Kaheman, A. J. Goldschmidt, J. Callahan, C. B. Delahunt, Z. G. Nicolaou, K. Champion, J.-C. Loiseau, J. N. Kutz, and S. L. Brunton, PySINDy: A comprehensive Python package for robust sparse system identification, *Journal of Open Source Software* **7**, 3994 (2022).
- [37] A. A. Kaptanoglu, K. D. Morgan, C. J. Hansen, and S. L. Brunton, Physics-constrained, low-dimensional models for magnetohydrodynamics: First-principles and data-driven approaches, *Physical Review E* **104**, 015206 (2021).
- [38] A. A. Kaptanoglu, J. L. Callahan, A. Aravkin, C. J. Hansen, and S. L. Brunton, Promoting global stability in data-driven models of quadratic nonlinear dynamics, *Physical Review Fluids* **6**, 094401 (2021).
- [39] J. M. Hokanson, G. Iaccarino, and A. Doostan, Simultaneous identification and denoising of dynamical systems, *arXiv preprint arXiv:2203.13837* (2022).
- [40] P. Zheng, R. Barber, R. J. Sorensen, C. J. Murray, and A. Y. Aravkin, Trimmed constrained mixed effects models: formulations and algorithms, *Journal of Computational and Graphical Statistics* **30**, 544 (2021).
- [41] K. Maass, M. Kim, and A. Aravkin, A nonconvex optimization approach to IMRT planning with dose–volume constraints, *INFORMS Journal on Computing* (2022).

- [42] A. Sholokhov, J. V. Burke, D. F. Santomauro, P. Zheng, and A. Aravkin, A relaxation approach to feature selection for linear mixed effects models, arXiv preprint arXiv:2205.06925 (2022).
- [43] D. Bertsimas and B. Van Parys, Sparse high-dimensional regression: Exact scalable algorithms and phase transitions, *The Annals of Statistics* **48**, 300 (2020).
- [44] H. Hazimeh, R. Mazumder, and A. Saab, Sparse regression at scale: Branch-and-bound rooted in first-order optimization, *Mathematical Programming*, 1 (2021).
- [45] D. Bertsimas and V. Digalakis, The backbone method for ultra-high dimensional sparse machine learning, *Machine Learning*, 1 (2022).
- [46] K. Hammond, C. Zhu, T. Brown, K. Corrigan, D. Gates, and M. Sibilina, Geometric concepts for stellarator permanent magnet arrays, *Nuclear Fusion* **60**, 106010 (2020).
- [47] H. An and L. Thi, An efficient algorithm for globally minimizing a quadratic function under convex quadratic constraints, *Mathematical programming* **87**, 401 (2000).
- [48] Z. Dostál and L. Pospíšil, Optimal iterative QP and QPQC algorithms, *Annals of Operations Research* **243**, 5 (2016).
- [49] J. Bouchala, Z. Dostál, T. Kozubek, L. Pospíšil, and P. Vodstrčil, On the solution of convex QPQC problems with elliptic and other separable constraints with strong curvature, *Applied Mathematics and Computation* **247**, 848 (2014).
- [50] L. Pospíšil and Z. Dostál, The projected Barzilai–Borwein method with fall-back for strictly convex QCQP problems with separable constraints, *Mathematics and Computers in Simulation* **145**, 79 (2018).
- [51] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, *Mathematical programming* **45**, 503 (1989).
- [52] R. Matsuoka, S. Kyochi, S. Ono, and M. Okuda, Joint sparsity and order optimization based on ADMM with non-uniform group hard thresholding, *IEEE Transactions on Circuits and Systems I: Regular Papers* **65**, 1602 (2017).
- [53] T. Van Leeuwen and A. Y. Aravkin, Variable projection for non-smooth problems, *SIAM journal on scientific computing* **43**, S249 (2021).
- [54] F. Wechsung, A. Giuliani, M. Landreman, A. J. Cerfon, and G. Stadler, Single-stage gradient-based stellarator coil design: stochastic optimization, *Nuclear Fusion* (2021).