

Discrete Approximation of Permanent Magnet Arrays for Mechatronic Applications

He Kai Lim¹, Will Flanagan¹, Cameron Taylor¹, and Tyler R Clites¹

¹Department of Mechanical and Aerospace Engineering, University of California Los Angeles, Los Angeles, CA 90095 USA

November 05, 2025

Abstract

Application of force through controllable magnetic fields holds enormous potential for actuation with little-to-no physical contact in mechatronic systems. Permanent magnet arrays have gained traction as an efficient means of shaping and controlling magnet fields in mechatronic systems, creating new incentives for engineers to design, build, and test new magnet array architectures that unlock new actuation capabilities. Unfortunately, these arrays are both expensive to build and difficult to assemble. In this manuscript, we present a general approach to building complex magnet arrays using low-cost and tileable cube magnets. This method implements practical voxelization and discretization of any magnet array into tileable cube elements, such that assembly of these discrete arrays is straightforward and scalable to any geometry or array design. We illustrate this approach in developing a robot end effector that utilizes a discretized controllable axial-flux Halbach array (D-CAHA). Performance of our discretized prototype in simple, easy-to-model experiments was consistent with model predictions. Our results also enriched the model by providing important empirical insights with implications for mechatronic design. Finally, we discuss the paradigm of discrete approximation in building magnet arrays, describe its novel capabilities in the context of developing mechatronic systems, and highlight open problems in generalizable discretization.

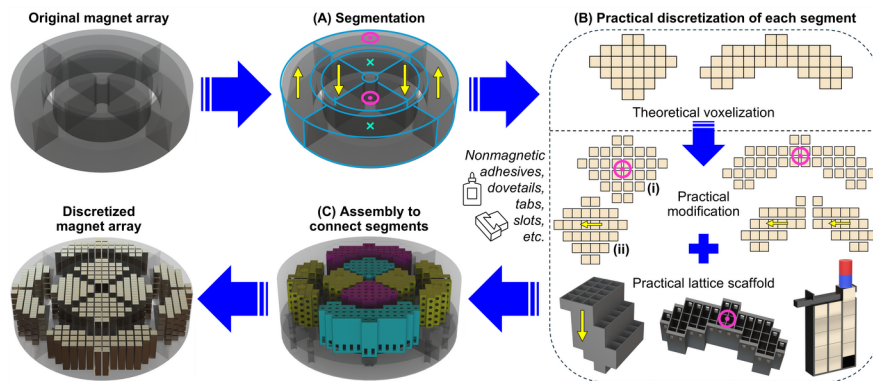


Figure 1: Practical discrete approximation of a permanent magnet array




Figure 2: Key concepts in taking (a) an array of permanent magnets and (b) applying theoretical voxelization, with illustrations on (c) a segment with uniform magnetization and (d) continuously varying magnetization.

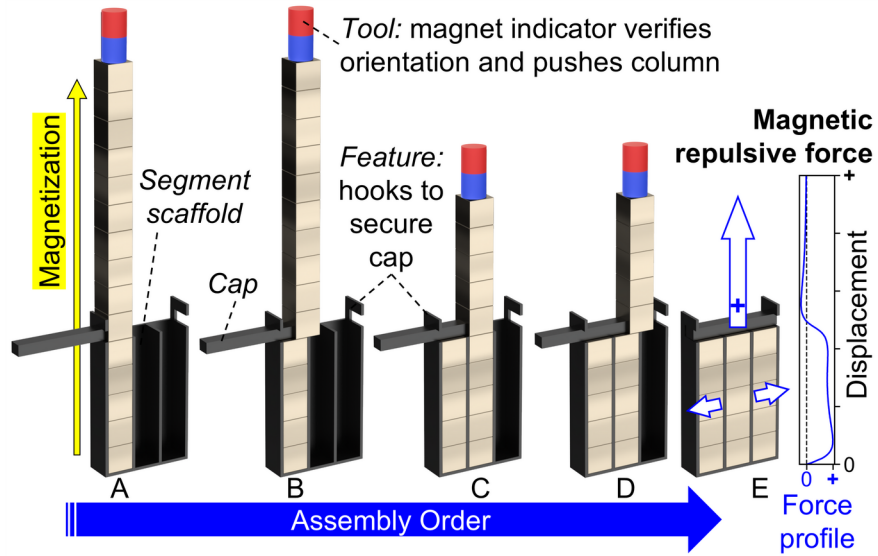


Figure 3: Cross-section view of a practical method to install magnet columns into a scaffold structure. (A) Column of magnets is first inserted into the scaffold until it bottoms out. (B) Cap is used to push away the excess magnets and align them with the adjacent scaffold column. (C) Next column of magnets is inserted until it bottoms out, while the cap ensures the previous column of magnets do not “shoot out” of the scaffold due to magnetic force instability. (D) This process is repeated until (E) all adjacent columns are installed. During installation, each column experiences magnetic repulsion due to the magnetic field created by the previously installed columns; an example of this force profile is shown.

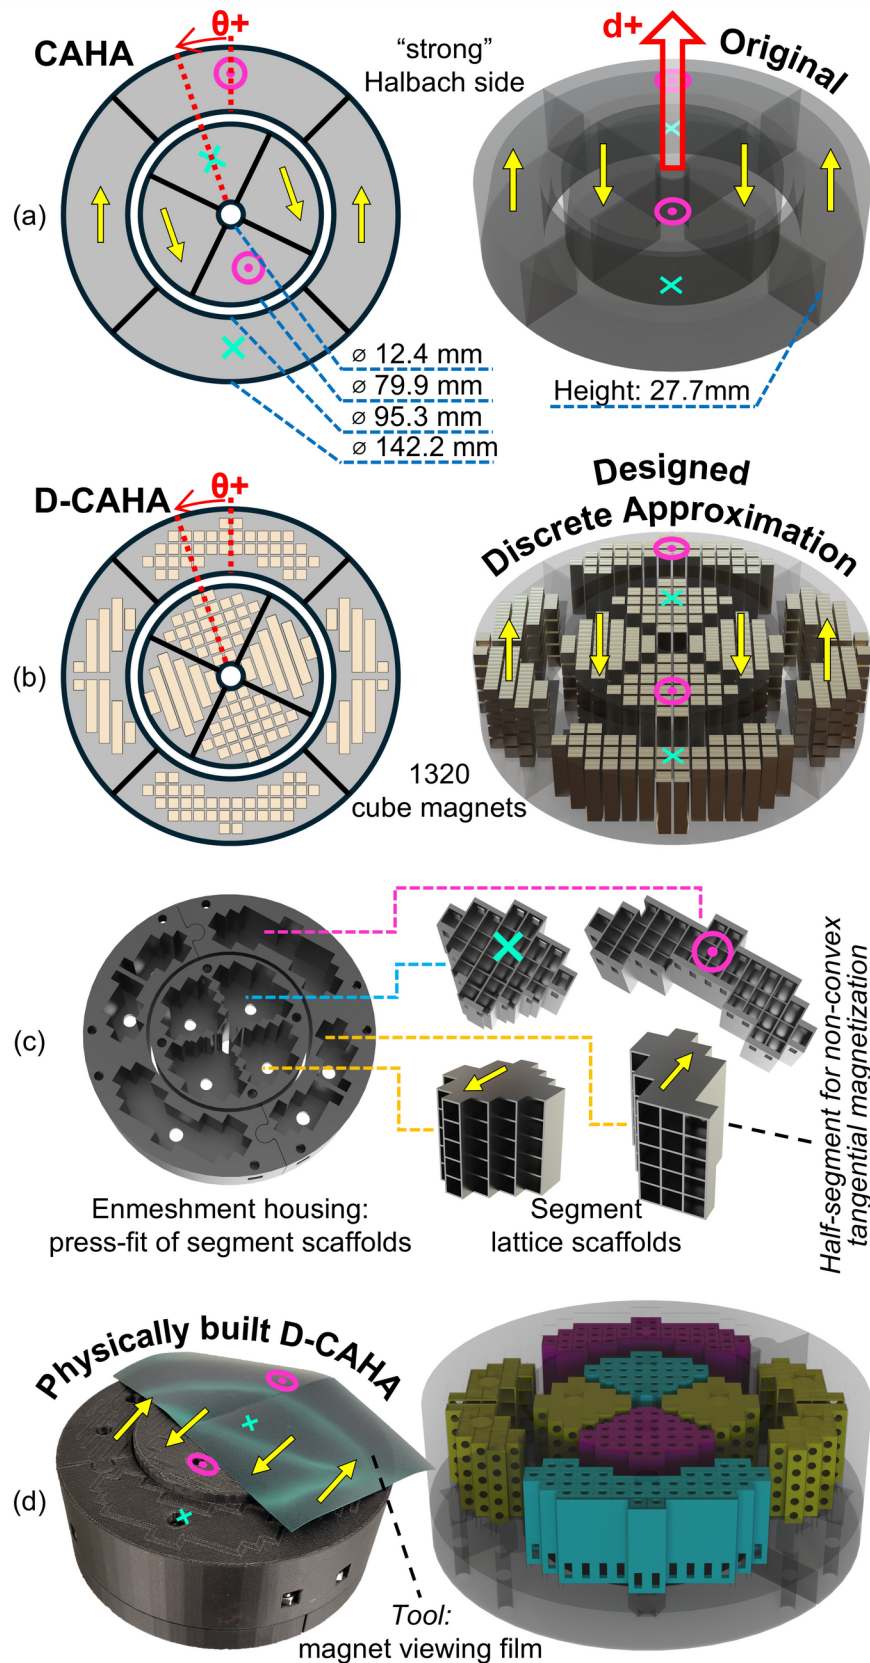


Figure 4: D-CAHA. (a) Original CAHA geometry showing relevant magnetization orientations of each arc segment, in the “OFF” or low-force state. The associated convention for positive angular actuation position ϑ and positive displacement above the strong surface d is shown. (b) D-CAHA geometry showing the superposed 1320 cube magnets over the nominal CAHA geometry, from practical discretization. Each cube magnet is a 3/16 inch (4.7625 mm) Grade N52 NdFeB Block Magnet. All D-CAHA cube elements fit exactly within the illustrated CAHA geometry. (c) After installing cube magnets into their respective lattice scaffolds (via the method shown in Fig. 3), the segments are installed into an unmagnetized housing. (d) Physically built D-CAHA showing the magnet viewing film.




Figure 5: Testing infrastructure for the D-CAHA.

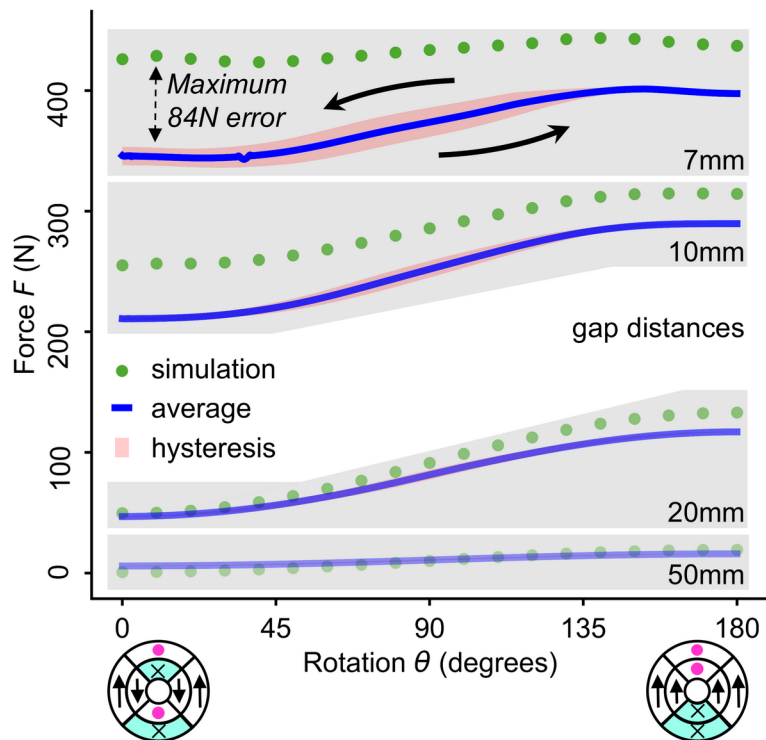


Figure 6: D-CAHA force F on the plate at selected gap distances. Note the hysteresis in pink where F is dependent on both the ϑ position and the direction of rotation (indicated by arrows for 7 mm gap distance). The ϑ -averaged force values are computed and shown. For comparison, simulated force values are shown at $\vartheta = 10^\circ$ increments as light green dots

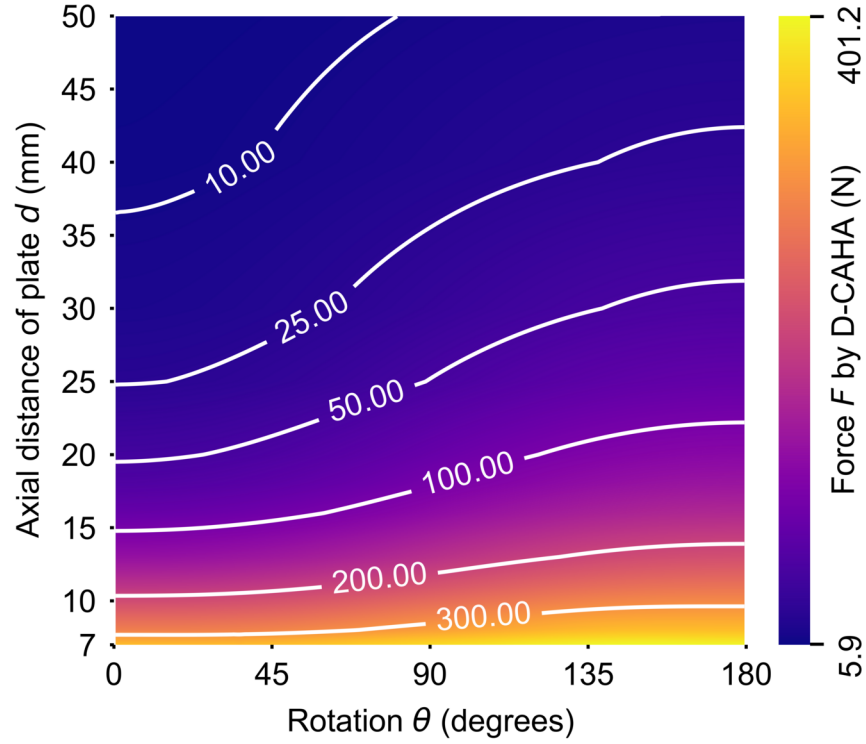


Figure 7: Experimentally measured D-CAHA forces F on the plate. Contour lines of constant force highlight trends on the heatmap. Note the monotonic trends: force decreases with distance d , and force increases with ϑ . While not shown, we report that all numerical partial derivatives on the heatmap are $F/\vartheta = [2.9, -2.33]$ (N/°), $F/d = [-0.17, -67.43]$ (N/mm) and $F/d = [-10, -67.43]$ (N/mm) | $d < 18$. The negative values of F/ϑ are erroneous edge-case points at $\vartheta = 0^\circ$ and $\vartheta = 180^\circ$ for $d = 7$ mm.




Figure 8: Quasistatic torque TIR on the IR across the d - ϑ landscape is shown. Positive torque values (solid contours) indicate torque that seeks to increase ϑ , and negative torque values (dashed contours) indicate that torque that seeks to decrease ϑ . While not shown, we report that all numerical partial derivatives on the heatmap are $TIR/\vartheta = [0.07, -0.07]$ (Nm/°), $TIR/d = [0.02, -0.05]$ (Nm/mm).




Figure 9: Quasistatic force FIR on the IR across the d - θ landscape is shown. Positive force values (solid contours) indicate force that attracts the IR towards the plate, and negative force values (dotted contours) indicate force that repels the IR away from the plate. While not explicitly shown, we report that all numerical partial derivatives on the heatmap are $FIR/\theta = [0.79, -0.87]$ (N/°), $FIR/d = [-0.07, -42.85]$ (N/mm), and $FIR/d = [-6, -42.85]$ (N/mm) | $d < 18$.

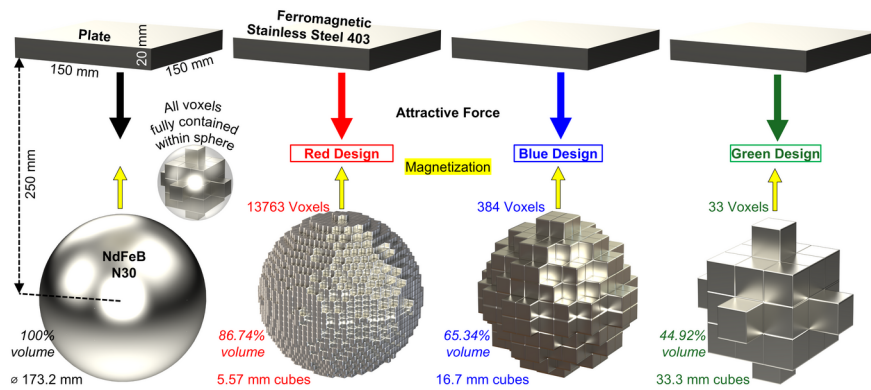


Figure 10: **Fig. S1.** Theoretical case study in voxelization of permanent magnet segments

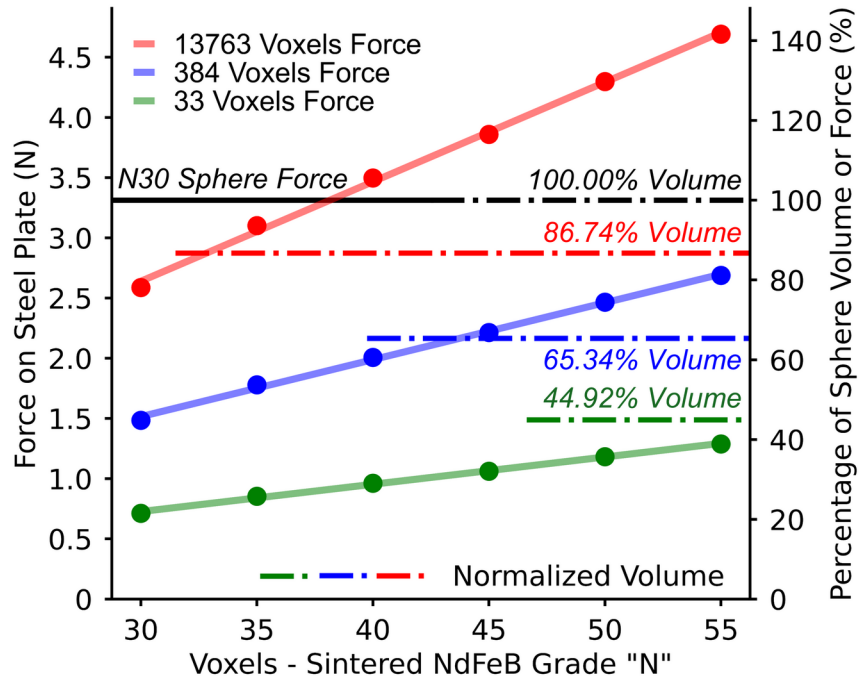


Figure 11: **Fig. S2.** Simulated case study results for each voxelized design, across different sintered NdFeB magnet grades (“N”). The range of N-grades here are common for off-the-shelf magnets. The nominal force exerted by the original N30 sphere is shown as a constant 3.31 N from the left axis, and it is matched against the right axis at 100% to represent 100% nominal force. Each datapoint can thus be read for its empirical value (left axis), and its relative value (right axis). For each dataset, a simple trendline is fitted for visual clarity, with $R^2 > 0.997$. The volume fraction of each design is further illustrated by dashed-dot lines from the right axis.

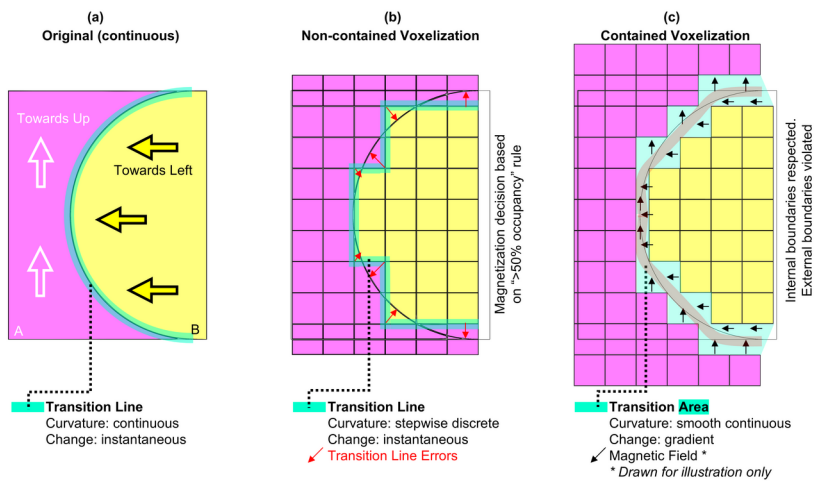


Figure 12: **Fig. S3.** Examining containment requirements on internal and external boundaries for voxelized permanent magnet arrays.

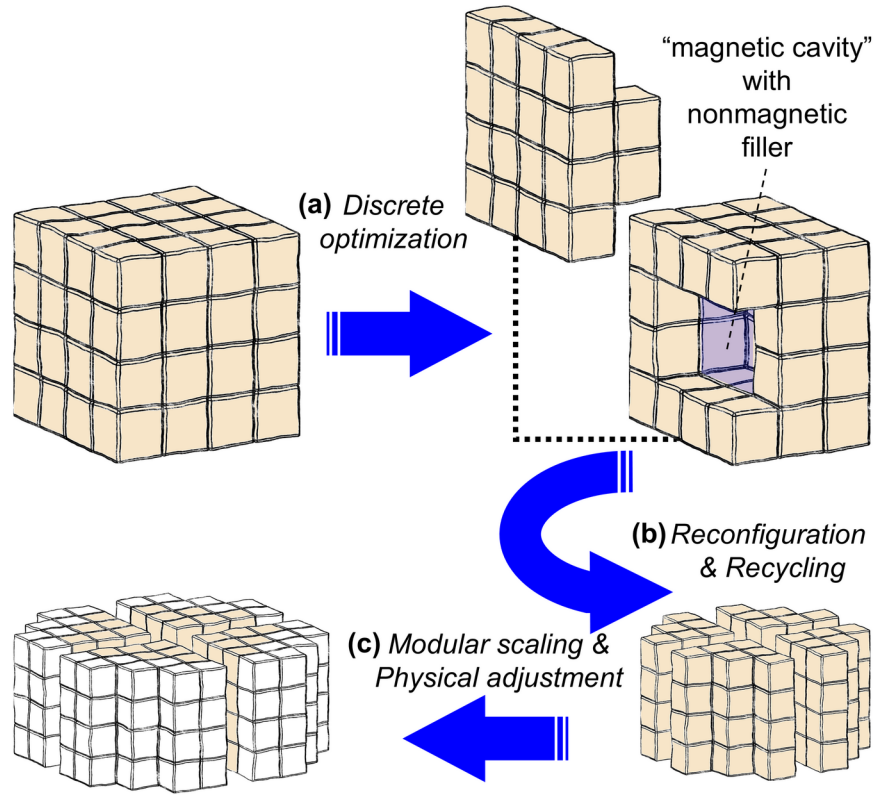


Figure 13: **Fig. S4.** Concepts of novel capabilities in magnet array prototyping with discrete approximation. (a) Geometric optimization of magnets is now feasible with magnetic undercut and internal cavity features. (b) After prototyping and testing, magnets can be recycled and reconfigured into new geometries for new systems and new applications. (c) Magnet array mechanisms can be modular and adjusted on the spot for applications using the same system, without fabricating wholly new permanent magnets every time.

Discrete Approximation of Permanent Magnet Arrays for Mechatronic Applications

He Kai Lim, Will Flanagan, Cameron R. Taylor, and Tyler R. Clites

Abstract—Application of force through controllable magnetic fields holds enormous potential for actuation with little-to-no physical contact in mechatronic systems. Permanent magnet arrays have gained traction as an efficient means of shaping and controlling magnet fields in mechatronic systems, creating new incentives for engineers to design, build, and test new magnet array architectures that unlock new actuation capabilities. Unfortunately, these arrays are both expensive to build and difficult to assemble. In this manuscript, we present a general approach to building complex magnet arrays using low-cost and tileable cube magnets. This method implements practical voxelization and discretization of any magnet array into tileable cube elements, such that assembly of these discrete arrays is straightforward and scalable to any geometry or array design. We illustrate this approach in developing a robot end effector that utilizes a discretized controllable axial-flux Halbach array (D-CAHA). Performance of our discretized prototype in simple, easy-to-model experiments was consistent with model predictions. Our results also enriched the model by providing important empirical insights with implications for mechatronic design. Finally, we discuss the paradigm of discrete approximation in building magnet arrays, describe its novel capabilities in the context of developing mechatronic systems, and highlight open problems in generalizable discretization.

Index Terms—Discrete magnet approximation, magneto-mechanical devices, mechanism design, permanent magnet array, robot actuation.

I. INTRODUCTION

PERMANENT magnet arrays lie at the heart of many mechatronic systems, because they enable energy conversion between mechanical and electrical forms without direct contact between components (e.g. a brushless DC motor). By simply reorienting these magnet arrays or changing a nearby electric field, energy can be remotely stored in the magnetic field, manifesting as a magnetic force with the potential to release energy for actuation and sensing. As mechatronic systems adopt to a wider range of tasks and applications, especially in robot design, the use of specialized magnet arrays to impart force in contact-limited situations has

become increasingly common. For example, microrobots are actuated by magnetic fields because they have limited contact area [1], [2], [3], [4], magnet arrays are used to harvest vibration energy [5], [6], magnetic grippers are used by climbing robots on unstructured or smooth surfaces [7], [8], [9], magnetic levitation systems are implemented on devices in hazardous environments [10], [11], and medical devices utilize magnetic fields for precise control of continuum devices and remote imaging [1], [12], [13], [14]. Across all these examples, magnetic mechanisms are created with an array of strong permanent magnets that change the magnetic field from a simple magnetic dipole approximation into some complex multipole expansion [1], so that it may be optimized for a given application. The canonical Halbach array is one such architected permanent magnet array that is well-suited for mechatronic system design [15], [16], wherein the field is shaped by setting strong permanent magnets in different (i.e., heterogeneous) orientations.

Currently, there is a technological gap between virtually designed magnet arrays and physically built mechanisms for robot actuation. Through computational modelling with finite element analysis (FEA), designers have embarked on deep explorations of complex heterogeneous magnet array designs that achieve optimality and unlock new capabilities [17], [18], [19]. Despite the exciting research into complex magnet arrays, the arrays that are implemented and tested in mechatronic systems are constrained by practicality, and therefore tend to comprise a limited number of magnets, avoid dense packing, and retain a straightforward geometry [2], [4], [5], [6], [7], [8], [9], [10], [11], [15], [16]. This practicality is, foremost, a financial problem, because fabrication of strong NdFeB permanent magnets in unique geometries (whether sintered or bonded) requires extensive tooling, manufacturing controls, and safety equipment [20], [21], [22], which drives up costs [23], [24] amidst the scarcity of necessary rare earth elements [25]. Furthermore, even with simple permanent magnet geometries, there is significant assembly complexity: a magnet array is inherently unstable due to the internal magnetic forces that act between each permanent magnet to push, pull, and twist one another. This means that an increase in the number of magnets in an array will also increase instability of the system [26], and thereby increase the difficulty of assembly into a mechanism. Altogether, the state of the art in building magnet array prototypes is not suited for low-cost iteration, limiting the adoption of task-optimized magnet arrays in mechatronic systems with robotic applications.

This work was supported by the National Science Foundation (Award Number 2144015). (Corresponding author: He Kai Lim).

He Kai Lim, Will Flanagan, and Cameron R. Taylor are with the Department of Mechanical and Aerospace Engineering, University of California Los Angeles, Los Angeles, CA 90095 USA (e-mail: limhekai@ucla.edu, wflan@ucla.edu, ctaylor7@unc.edu).

Tyler R. Clites is with the Department of Mechanical and Aerospace Engineering and Orthopaedic Surgery, University of California Los Angeles, Los Angeles, CA 90095 USA (e-mail: clites@ucla.edu).

This article has supplementary material provided by the authors and color versions of one or more figures available online at <http://ieeexplore.ieee.org>.

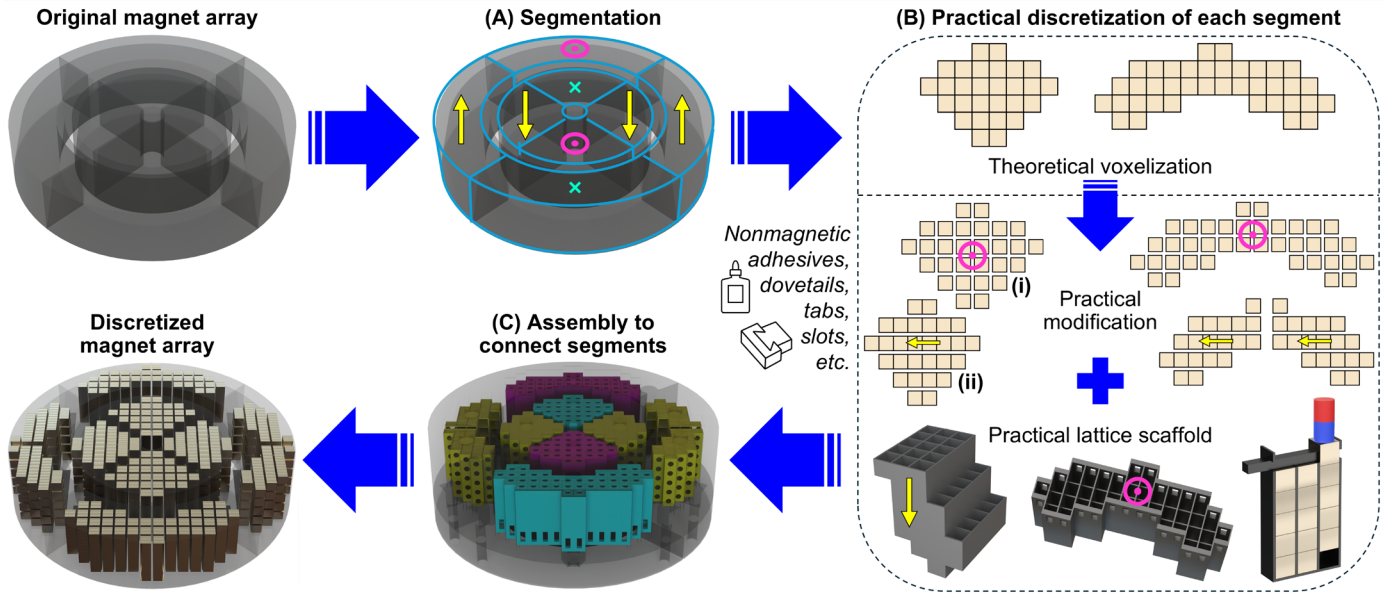


Fig. 1. Practical discrete approximation of a permanent magnet array.

To facilitate experimental iteration and adoption of magnet arrays into mechatronic systems, we developed a strategy for low-cost development of heterogeneous magnet arrays, by way of practical, scalable, and generalizable discrete approximation. We demonstrate this approach by implementing the design-build-test cycle on a discretized Controllable Axial-flux Halbach Array (CAHA) [27] end effector device (comprising an assembly of 1320 cube magnets), and we illustrate the utility of this discrete prototype in de-risking magnetic mechanisms. We also discuss the implications of discrete approximation and highlight future research directions to fuel further development of novel permanent magnet arrays for new mechatronic systems.

II. DISCRETIZATION OF PERMANENT MAGNET ARRAYS

The goal of discrete approximation is to build a low-cost version of any arbitrary magnet array that is otherwise complex and expensive to fabricate. The process of building this *discretized* array encompasses three steps (Fig. 1): (A) divide the array into contiguous segments, each with uniform (parallel) magnetization, (B) theoretically voxelize a segment into low-cost cubic elements, practically modify the voxel grid for assembly into lattice scaffolds, and (C) assemble segments into a complete discretized prototype. Notably, this discrete array creates a *close approximation* to (rather than an exact replica of) the magnetic field of the original, continuous array, which is sufficient for most prototype mechatronic applications. In this section, we describe in detail the key concepts of the discretization process.

A. Voxelization of a Monolithic Permanent Magnet Segment

In the most general case, a theoretical permanent magnet can be of any manifold geometry, so long as it exhibits the desired and designed field for a magnetic mechanism. Naturally, some manifold geometries are costly or

complicated to fabricate as a monolithic magnet because they have features like enclosed internal cavities or inaccessible undercuts. An intuitive approach to scalable and low-cost approximation of these geometries is to modify the LEGO construction problem (“Given any 3D body, how can it be built from LEGO bricks?”) [28], [29], [30] so that it is applicable to magnets. Cube-shaped NdFeB magnets with face-normal magnetization are an excellent choice for the brick element in this approximation because they are low cost, commonly fabricated in large quantities, and tileable in 3D space. Given an original *continuous* manifold geometry (Fig. 2a), we can thus apply *voxelization* [30], [31] and discretize the shape into n cube-shaped “magnet voxels” (Fig. 2c), with a constraint of aligning the voxel grid with the magnetization orientation of the original magnet (Fig. 2c). Along surfaces and edges of adjacent segments (i.e. “internal” boundaries), there is a further constraint of using a voxelization algorithm with strict containment [32], such that voxels do not pass the segment boundary and collide into the voxel of a neighboring segment (Fig. 2b, upper left quadrant). Detailed discussion on voxelization containment for all boundaries is provided in Supplemental S2. With voxelization, the accuracy of any manifold *geometric* approximation increases as the magnet voxel sizes decrease, because it reduces the volume of empty voxel regions along the interior boundary of each segment.

Besides gross geometry, a permanent magnet is defined by the strength and orientation of its magnetization. This is represented in cartesian space (x, y, z) by a magnetization vector field $\vec{M}(x, y, z)$ within the original, continuous magnet. To apply magnet voxelization, the magnetization vector field must be binned into each voxel, analogous to how photographic pixels capture a complex image projection. Just as each pixel encodes a scalar (grayscale) or vector (color) information packet, we assert that each magnet voxel i must encode a “voxel vector” $\vec{V}_i, i \in \{1, \dots, n\}$, that describes the strength and orientation of its cube magnet element. In




Fig. 2. Key concepts in taking (a) an array of permanent magnets and (b) applying theoretical voxelization, with illustrations on (c) a segment with uniform magnetization and (d) continuously varying magnetization.

designing the system for practical assembly, these voxel vectors must be orthogonal to the cube faces of each magnet voxel (Fig. 2b). Importantly, the voxel vector magnitude may be zero to indicate the absence of any magnet element (an empty voxel), but the magnetic field itself might be nonzero due to the exuded field from neighboring magnetized voxels (Fig. 2c).

Care is needed in the voxelization of magnet designs that have continuously *varying* magnetization within a single body (in contrast to the more prevalent *uniform* magnetization). An example of such varying magnetization is a ring magnet with radial magnetization, which has the “south” pole on its inner diameter, and the “north” pole on its outer diameter (or vice versa), which may be represented as

$$\vec{M}(x, y, z) = m_{\text{original}} \begin{bmatrix} x \\ \sqrt{x^2 + y^2} \\ y \\ \sqrt{x^2 + y^2} \\ 0 \end{bmatrix}, \quad (1)$$

where m_{original} is the magnitude of magnetization. Equation (1) also describes the magnetization vector field of the Fig. 2d ring magnet segment. In this scenario, the cubic voxelization grid cannot be truthfully aligned with the magnetization vector field of the original magnet because the vector field varies continuously in space. Each voxel (and its voxel vector) may thus be oriented along the average magnetization vector of the

field it occupies (Fig. 2d), as

$$\frac{\vec{V}_i}{|\vec{V}_i|} = \frac{1}{|XYZ_i|} \iiint_{XYZ_i} \vec{M}(x, y, z) dx dy dz, \quad (2)$$

where XYZ_i is the region occupied by the cubic voxel i . Explicitly, (2) shows equivalency between the unit voxel vector, and the unit average magnetization vector.

Perhaps intuitively, this approach will decrease the voxel density of a region (compare Fig. 2c with Fig. 2d), because it no longer capitalizes on the tileable nature of cubic voxel shapes. This results in unvoxelized regions inside the segment (Fig. 2d), in addition to the pre-existing empty voxel regions along each segment boundary (Fig. 2a). Together, we call these “unmagnetized” regions (Fig. 2b).

Conventional literature on general voxelization suggests that these unvoxelized regions are not desirable because it reduces geometric accuracy, and so it is common to adopt geometric strategies that minimize these regions [32], for example, by using smaller voxels to improve fit, or to optimize voxel geometric position and orientation with half-length perturbations in each cardinal direction. We note, however, that in permanent magnet arrays, the intent is to replicate *function*, not *form*. The magnetic effect of these unvoxelized and unmagnetized regions can be compensated for by using voxels with stronger magnetization (i.e. increasing $|\vec{V}_v|$) (see Supplemental S1 for a case study). These regions can also provide a physical space for the magnetic field to smooth out and more accurately replicate the original magnetic field (see Supplemental S2 for an in-depth discussion).

Given the infinite applications and infinite design space for a magnet array, it is important to use FEA simulation to quantify the true magnetic effects of these unmagnetized regions in choosing a final voxelized design (voxel size, position, and magnetization) from an original and continuous magnet array. If these parameters are insufficient in finding a suitable voxelized design for a given magneto-mechatronic application, other strategies may be explored, such as using non-cubic voxels with better-suited tileability (e.g. arc-segment voxels to replace Fig. 2b), changing voxel containment requirements, or even hybrid combinations of partially-voxelized and partially-continuous arrays.

B. Practical Discretization of a Segment with Parallel Magnetization

When implemented physically, a voxelization grid of cube magnets is almost always mechanically *unstable*, because each cube magnet seeks to push, pull, and twist one another with magnetic force and magnetic torque [26]. This magnetic interaction potential is minimized when the field lines around a magnet form a closed loop in the shortest possible path length. For example, two cube magnets will naturally align their magnetization axes such that the north pole of one magnet points directly at the south pole another, thereby closing its field line in the shortest distance possible. In a counter example, two cube magnets will experience the greatest repulsive force if their like poles are adjacent to one

TMECH-02-2025-20230

another, because their field lines travel across the full magnet length to close the loop. With these two fundamental examples, we can analyze the behavior of a voxelization grid of cube magnets with parallel magnetization: there will be columns of attraction along the axis of magnetization, and there will be mutual repulsion in the orthogonal axes (because like-poles are adjacent to one another). Notably, these same stresses exist within all monolithic solid permanent magnets (i.e. magnetostriction), but they are made to hold themselves together by their engineered material structure. In contrast, practical magnet discretization must account for these destabilizing forces in feasible and scalable assembly.

To create stability in a discrete system, a voxelized magnet must have a physical scaffold between the cube elements to hold the cube magnets in place. Specifically, we represent the voxel grid with a physical three-dimensional cube lattice structure with non-zero-thickness walls that are strong enough to hold all the interspaced cube magnets and enforce stability. In the case where all cubes have parallel magnetization, the lattice structure only needs walls between each stable column (Fig. 3), because intra-column magnets are stable but there is no inter-column stability. The 3D lattice structure is thus simplified to a 2D lattice with no intra-column walls required. Importantly, this means that each modified magnet voxel is now spread apart with non-zero spacing between each column. The use of this lattice structure makes assembly of the discrete magnet feasible, because each column of magnets can be inserted into the lattice cell from an open end and secured in place with a cap structure (Fig. 3).

In keeping with the low-cost objective of this work, the lattice structure can be fabricated with additive manufacturing technology such as fused deposition modelling (FDM), which is inexpensive and allows for rapid production. As a fundamental test, we used SOLIDWORKS Simulation (Dassault Systèmes, Waltham, MA, USA) and found Onyx filament (micro carbon fiber filled nylon) in 0.40 mm thickness walls to be sufficient in strength to hold two N52 grade cube magnets of 4.7625 mm size. For more complex magnet configurations, lattice strength is best assessed with FEA, or alternatively, simple guess-and-check iteration with an FDM printer to find appropriate wall thicknesses.

We fabricated several lattice structures and installed magnet columns as shown in Fig. 3. After the first column, every subsequent column (B, C, etc.) will experience magnetic force as noted by the plot in Fig. 3. Each column is first attracted towards the previous installed column as opposite poles approach one another, then self-center just before insertion into a cell as unlike poles become exactly adjacently aligned. Inserting the column past this point of alignment creates repulsion. We found that gravity is insufficient to overcome this, and so a small external force is required to slowly push in each column. Upon bottoming out exactly, a column will theoretically experience zero force, but this is unstable and highly sensitive to small fabrication tolerance errors. These errors can cause a column to shoot out with high velocity, and so there is danger in the magnets slamming into the ceiling of the room and potentially fracturing. When

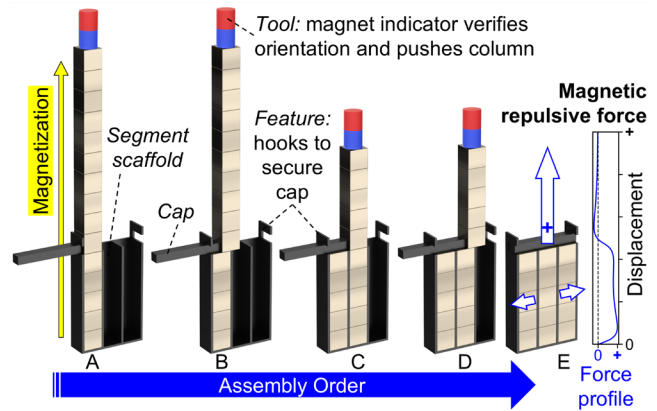


Fig. 3. Cross-section view of a practical method to install magnet columns into a scaffold structure. (a) A column of magnets is first inserted into the scaffold until it bottoms out. (b) The cap is used to push away the excess magnets and align them with the adjacent scaffold column. (c) The next column of magnets is inserted until it bottoms out, while the cap ensures the previous column of magnets do not “shoot out” of the scaffold due to magnetic force instability. (d) This process is repeated until (e) all adjacent columns are installed. During installation, each column experiences magnetic repulsion due to the magnetic field created by the previously installed columns; an example of this force profile is shown.

dealing with columns of varying heights, lateral force effects become more pronounced, and additional fixturing can help align a column above its cell before pushing it in. Finally, we found it helpful to make each lattice cell slightly undersized, to leverage a light press-fit in reducing mechanical stresses and deflections in the cap during the assembly process.

The incorporation of this practical lattice structure with non-zero wall thickness has an important trade-off: the magnetic vector field of the discretized magnet will deviate further from the continuous magnet's magnetic field. Specifically, it changes the discrete magnetic vector field shape and magnitude. In terms of shape, we note that having voxels spread apart means that the field is changed in the regions of space close to the magnet. At further distances, these changes are less apparent, and so the discrete magnet retains more accuracy in approximating its continuous counterpart. Regarding vector field magnitude, the reduction in packing density of voxels means there is simply less magnet material within the same voxel grid, and so the discrete magnetic device would create a weaker magnetic field than a continuous one. This, however, can be qualitatively compensated for by using a stronger grade of NdFeB cube elements than the original continuous magnet (see Supplemental S1 for a case study). As needed, detailed FEA could also be performed for each discretized magnet segment to ascertain appropriate cube magnet grades in each voxel.

C. Forming an Array of Discretized Magnets

A heterogeneous magnet array comprises multiple magnet segments, each arranged in space in various orientations (Fig. 2a). Notably, a single voxelized magnet segment with varying

TMECH-02-2025-20230

magnetization may also be considered as a heterogeneous magnet array (Fig. 2d). All heterogeneous magnet arrays are unstable, because each magnet segment is, once again, pushing and pulling and twisting one another. To connect these segments, we may utilize space within the unmagnetized region (Fig. 1B, Fig. 2d) to incorporate larger mechanical interconnecting features such as adhesives, dovetails, tabs, slots, or press-fits, to hold all these segments in place. Here, the choice of mechanical feature is up to the designer's creativity; our example in the following section demonstrates use of a monolithic housing where all segments are press-fit in place, but it is only one of the myriad possibilities of solutions.

Importantly, assembly of these segments must be done safely and cautiously. Despite its low-cost 3D-printed construction, a permanent magnet array with thousands of cube magnets can hold tremendous internal energy. By forcing each cube magnet into a lattice array of uniform (parallel) magnetization, the internal magnetic potential of the assembly is dramatically increased. If mishandled, the very same magnetic forces that push, pull, and twist each cube can also turn them into projectiles with high kinetic energy, posing a danger to nearby personnel. Proper fixtures for each segment, safety glasses, and plexiglass shields are always recommended in constructing a discretized permanent magnet array.

III. DESIGN AND ASSEMBLY OF A DISCRETIZED CAHA

A CAHA is used to illustrate the utility of permanent magnet array discretization in a mechatronic system. It consists of two nested rings of permanent magnets [27]. Each ring comprises a cylindrical axial-flux Halbach array: an arrangement of arc-segment permanent magnets with axial and tangential magnetization, that seeks to augment field lines on one (strong) side of the ring's flat surface and cancel field lines on the other (weak) side (Fig. 4a). A CAHA further enables cancellation of the field lines on the strong side via adjustment of the alignment between the inner ring (IR) and outer ring (OR), allowing for direct control of the resultant field strength. Effective magnetic attraction on a ferromagnetic object is monotonically correlated with angular position θ within a certain Halbach periodicity (wave number) k [27]; it changes from "off" to "on" with increasing θ (Fig. 4a). The CAHA concept shows promise in robotic applications because it creates controllable axial force at a moderate separation distance and can be easily configured as an end-effector; however, it is costly to build. Anecdotally in the United States, it was quoted at \$11,500 USD to fabricate CAHA magnets for our research group. The low-cost alternative for building a functional prototype was the Discretized CAHA (D-CAHA) in this work, which cost \$658 USD in permanent magnets: we thus achieved a 94% cost savings with our discretization technique.

To assess the viability of using the CAHA in a robot end effector mechanism, given limited financial resources, we applied the discretization process to a CAHA and obtained the D-CAHA (Fig. 4). Segments were identified (Fig. 4a), theoretically voxelized, practically modified (Fig. 4b), then connected to form the completed D-CAHA (Fig. 4c, Fig. 4d).

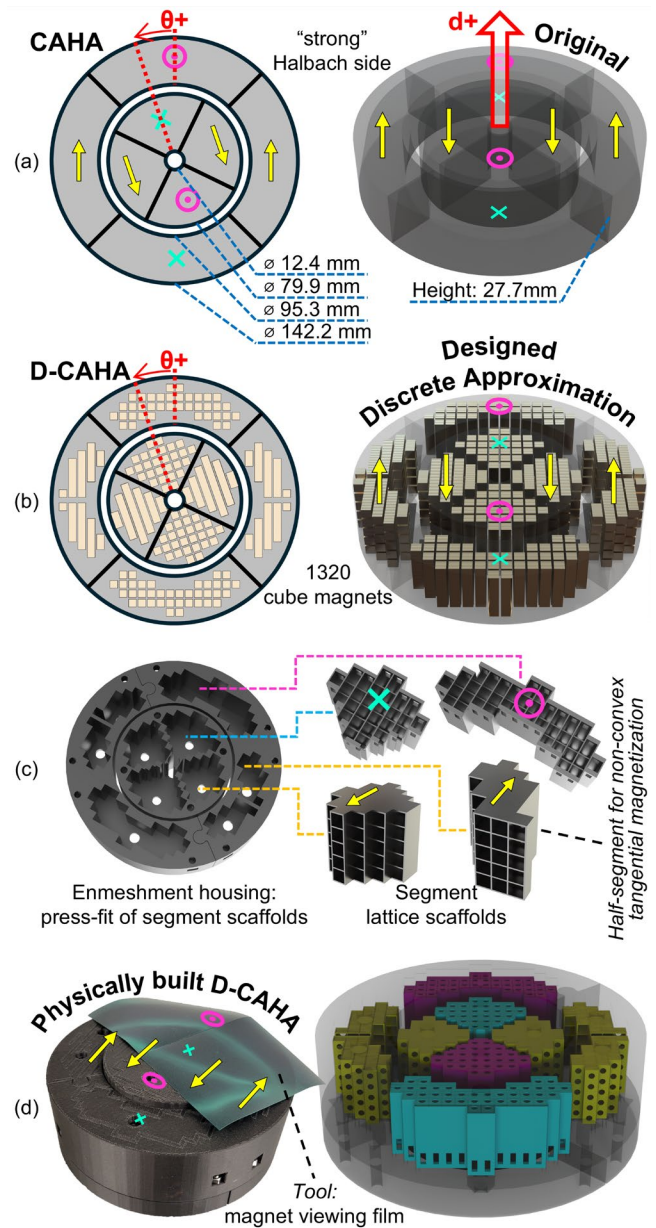


Fig. 4. D-CAHA. (a) The original CAHA geometry showing relevant magnetization orientations of each arc segment, in the "off" or low-force state. The associated convention for positive angular actuation position θ and positive displacement above the strong surface d is shown. (b) D-CAHA geometry showing the superposed 1320 cube magnets over the nominal CAHA geometry, from practical discretization. Each cube magnet is a 3/16 inch (4.7625mm) Grade N52 NdFeB Block Magnet. All D-CAHA cube elements fit exactly within the illustrated CAHA geometry. (c) After installing cube magnets into their respective lattice scaffolds (via the method shown in Fig. 3), the segments are installed into an unmagnetized housing via light press-fit. Note that the nonconvex geometry of the outer segment with tangential magnetization requires a half-segment design to facilitate practical installation of cube magnets. (d) Magnetization of the completed D-CAHA is verified with magnet viewing film, and the rendered translucent model shows the installed internal segments.

Due to limited development time, the discretization process

TMECH-02-2025-20230

was *not* optimized for packing density, as noted in the sizeable gray spaces of Fig. 4b, but this was deemed acceptable in our simulated performance with FEA. Of note is the split of the OR tangential segment (indicated with yellow arrow magnetization) into two halves (Fig. 4c), due to its nonconvex geometry, in order to facilitate practical column insertion using the method of Fig. 3.

We assembled the D-CAHA using 1320 cube magnets, each of size 4.7625 mm (3/16 inch). These were grade N52 NdFeB magnets with Ni-Cu-Ni coating (K&J Magnetics Inc, Pipersville, PA, USA), and their size was selected based on the practical ease of holding and manipulating them by hand without special tooling. We constructed the scaffold of Onyx (carbon fiber filled nylon) on a Mark Two (Markforged, Waltham, MA, USA) 3D printer with 30% infill and 0.4 mm nozzle size. We note that cheaper materials could have been used, but this was the most readily available at the time of implementation. Each D-CAHA segment comprised a lattice scaffold, and we installed the magnets using the method shown in (Fig. 3). We used lattice structure walls that were nominally 0.55 mm thick (smaller dimensions yielded quality defects from the 3D printing process). The structure had cavities that were nominally oversized by 0.42 mm as an FDM XY compensation factor that still achieved suitable press-fit tolerance for the magnet columns. It also had sliding caps to hold the columns in place, similar to the concept shown in Fig. 3. We press-fit all segments into the 3D-printed unmagnetized housing (Fig. 4c), which had a 5 mm wall thickness on the “strong” side to hold all segments securely. Magnetic viewing film and color-coded magnet indicators were used to verify that all cubes, columns, and segments were installed in the correct orientation. Where required, nuts and bolts were only used on the weak-side of the D-CAHA, made of 316 Stainless Steel material (generally non-magnetic).

In total, the assembly contained 1320 cube magnets and 60 individual 3D-printed parts. We included a 1.5 mm radial clearance between the inner and outer rings. Overall, the D-CAHA was $\varnothing 160$ mm in diameter and 50 mm tall; however, all cube magnets were still enclosed within the target CAHA volume (Fig. 4a). Notably, in designing, building, and testing this D-CAHA, a 3D printer was the only customizable fabrication resource.

IV. D-CAHA TESTING METHODOLOGY

The D-CAHA was built as a low-cost method to assess the viability of a CAHA robot end effector mechanism. We thus tested the D-CAHA by actuating it with a robot manipulator (KR210; KUKA Robotics Corp., Clinton Township, MI) while measuring forces and torques in the system (Fig. 5).

The 6DOF load cell (ATI Omega160) on the KUKA tool flange recorded forces experienced by the D-CAHA inner ring (IR). The IR is aligned with the D-CAHA outer ring (OR), which is fixed onto a testbench. By rotating the IR, there is hence a controllable magnetic force by the D-CAHA, across an air gap, onto a ferromagnetic steel plate. Axial force onto the plate was measured (1DOF load cell, Futek LCM300)




Fig. 5. Testing infrastructure for the D-CAHA.

during actuation of the D-CAHA (Fig. 5). Further details on the experimental setup and data analysis for this case study are found in Supplemental S3 and S4.

All simulation results were obtained through FEA with JMAG software (JSOL Corporation, Tokyo, Japan) via transient analysis in steps of $\theta = 10^\circ$ increments per second. A sliding mesh was employed on the D-CAHA inner ring to save on computation time, and mesh element sizes were initialized at 1.5 mm for all solid bodies and air volumes.

V. RESULTS AND CHARACTERIZATION

The primary performance results from our D-CAHA consist of axial forces and torques that the array exerts when interacting with a ferromagnetic plate, which directly relate to the mechanics of actuating the D-CAHA on an end effector. For completeness, since they were also recorded, we note that radial forces and torques on the IR were a nonzero function of d and θ , and in polar coordinates, their magnitudes never exceeded 9.2 N and 1.3 Nm in all tested configurations.

The D-CAHA generated a controllable axial force on the metal plate that varied with both θ and gap distance (Fig. 6). We observed reasonable agreement between simulation predictions and experimental results, with a maximum observed prediction error of 84 N (24% error), occurring at $\theta = 10^\circ$ and the smallest tested gap distance ($d = 7$ mm). Detailed analysis of these discrepancies (i.e. from tolerance stack, mechanical deformation, demagnetization) is provided in the Supplemental S6. While the test infrastructure enabled a theoretical minimum $d > 5$ mm for the experiment, we discarded those boundary values because the forces were strong enough to create untenable direct contact between magnet and the plate. We also observed hysteresis in the attractive force, especially at small values of d , where the force magnitudes are large. This is likely due to some combination of magnetic hysteresis and viscoelastic deformation in the 3D printed mechanical support. To simplify analysis, we report the θ -averaged data hereon. Controllability of the attractive force was maintained across all tested gap distances (Fig. 7).

We observed a two-peak trend in the torque required to rotate the IR relative to the OR as a function of θ (Fig. 7). Peak torques were measured at the largest distances from the plate, with a maximum absolute torque of 1.3 Nm. Across the heatmap, most torques are negative, meaning that the “off”

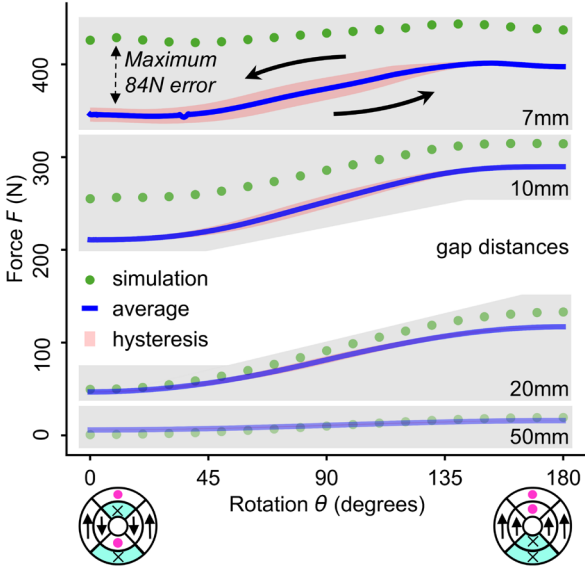


Fig. 6. D-CAHA force F on the plate at selected gap distances. Note the hysteresis in pink where F is dependent on both the θ position and the direction of rotation (indicated by arrows for 7mm gap distance). The θ -averaged force values are computed and shown. For comparison, simulated force values are shown at $\theta = 10^\circ$ increments as light green dots.




Fig. 7. Experimentally measured D-CAHA forces F on the plate. Contour lines of constant force highlight trends on the heatmap. Note the monotonic trends: force decreases with distance d , and force increases with θ . While not shown, we report that all numerical partial derivatives on the heatmap are $\partial F/\partial\theta = [2.9, -2.33]$ ($N/^\circ$), $\partial F/\partial d = [-0.17, -67.43]$ (N/mm) and $\partial F/\partial d = [-10, -67.43]$ (N/mm) $| d < 18$. The negative values of $\partial F/\partial\theta$ are erroneous edge-case points at $\theta = 0^\circ$ and $\theta = 180^\circ$ for $d = 7mm$.

configuration is a preferred state. There are three small regions where this torque is positive ($\theta \approx 0^\circ$, $\theta \approx 180^\circ$, $\theta \approx 90^\circ$ $| d \lesssim 20$) because of the two peak torque trend (Fig. 8). Axial force experienced by the inner ring during rotation is a complex function of θ and d (Fig. 9). Forces generally increased as plate distance decreased, with a maximum 248 N measured at the experimental limit $d = 7$ mm volumes.




Fig. 8. Quasistatic torque T_{IR} on the inner ring across the d - θ landscape is shown. Positive torque values (solid contours) indicate torque that seeks to increase θ , and negative torque values (dashed contours) indicate that torque that seeks to decrease θ . While not shown, we report that all numerical partial derivatives on the heatmap are $\partial T_{IR}/\partial\theta = [0.07, -0.07]$ ($Nm/^\circ$), $\partial T_{IR}/\partial d = [0.02, -0.05]$ (Nm/mm).




Fig. 9. Quasistatic force F_{IR} on the inner ring across the d - θ landscape is shown. Positive force values (solid contours) indicate force that attracts the IR towards the plate, and negative force values (dotted contours) indicate force that repels the IR away from the plate. While not explicitly shown, we report that all numerical partial derivatives on the heatmap are $\partial F_{IR}/\partial\theta = [0.79, -0.87]$ ($N/^\circ$), $\partial F_{IR}/\partial d = [-0.07, -42.85]$ (N/mm), and $\partial F_{IR}/\partial d = [-6, -42.85]$ (N/mm) $| d < 18$.

VI. DISCUSSION AND CONCLUSION

In this manuscript, we proposed a general method of constructing discrete approximations of permanent magnet arrays. We then demonstrated the practicality of this method by designing, building, and testing a D-CAHA robot end effector. In this discussion, we will interpret specific results of our D-CAHA testing, and make broad observations about the paradigm of discrete approximation for mechatronic systems.

A. The D-CAHA end effector mechanism

The D-CAHA behaved grossly as it was intended to, creating a controllable magnetic force on a ferromagnetic object without direct contact. Although its initial conception was to simply emulate the continuous CAHA, we found that it may have stand-alone value as a low-cost alternative in settings without tight space constraints. For a given manipulation task with known requirements of F and d , Fig. 7 can be used to check for feasibility and design an appropriate high-level control system for θ and d parameters.

To implement the D-CAHA as a deployed robotic end-effector, the IR's position θ would need to be controlled by a motor or stepper. We thus require an understanding of axial torques (T_{IR}) and forces (F_{IR}) during actuation, to design an appropriate motor-transmission system. Fig. 8 shows a two-peak torque trend, which is very different from the one-peaked CAHA [27]. The two-peak phenomenon can be understood as a form of “cogging torque”, similarly found in brushless permanent magnet motors, that stems from discrete construction. We also note the intrinsic system instability: for an initial position θ , a disturbance to increase d will decrease F (Fig. 7), make T_{IR} more negative (Fig. 8), reduce θ and further decrease F (Fig. 7), so that d increases unboundedly.

Feedback control on d is required to manage this instability in a robotic system. In exploring methods of sensing and feedback on the state of the system, we examined the internal *torques* of the D-CAHA system and investigated the viability of inferring plate distance based on IR torque; a similar concept to inferring robot manipulator joint torque with motor current. The heatmap of T_{IR} (Fig. 8) shows that every position θ has a unique value of T_{IR} across distance d (the inverse is not true with non-unique values of T_{IR} at a given position d). This means it is possible to infer plate distance with IR torque data, and so inferred feedback control is plausible without direct distance measurement. However, T_{IR} shows low sensitivity to d given the low values of $\partial T_{IR}/\partial d$, and so the signal-to-noise ratio of this feedback data stream is expected to be poor.

We also examined internal *forces* of the D-CAHA system. The heatmap of F_{IR} across d - θ (Fig. 9), shows that F_{IR} is unintuitively *not* monotonic in θ , despite the monotonic trends from F (plate force) (Fig. 7), because of the magnetic interaction between IR and OR (this non-monotonicity in F_{IR} was also replicated in simulation of a CAHA). Notably, Fig. 7 and Fig. 9 together can be used to calculate axial force on the outer ring, since force on the plate is equal to the sum of forces by the inner and outer rings: it is a one-dimensional three-body problem. Both Fig. 7 and Fig. 9 also show uniqueness in F and F_{IR} : for every position θ , there is a unique value of F (or F_{IR}) across distance d . Additionally, we highlight the large magnitudes of $\partial F/\partial d$ and $\partial F_{IR}/\partial d$ for $d < 18$: this lends confidence to the idea of measuring axial force on the D-CAHA (as a whole, or simply from an individual ring) with a load cell, to precisely infer distance d as feedback *for* control in a deployed mechatronic system. Through building and testing this discrete system, we thus discovered that a D-CAHA mechatronic system is conceptually able to both actuate *and* sense its actuation state.

B. A Comparison between the D-CAHA and a CAHA

Although the fundamental performance characteristics of the D-CAHA mimicked those of the CAHA (i.e. controllable strong-side attractive force), discrete approximation theory suggests that they are not interchangeable: for the same weight, the D-CAHA has a larger volume because it has a lower packing density of voxels. Furthermore, while we understand the impact of geometric parameters (e.g. ring diameter ratio, height) on CAHA performance [27], we do not know their impact on the D-CAHA since they relate closely to the voxelization algorithm and discretization resolution. Any simulated comparison between a D-CAHA and CAHA would require extensive study of the voxelization process on CAHA morphology, which is beyond the scope of this manuscript.

Despite the lack of interchangeability between these two systems, the D-CAHA established critical empirical intuition for the CAHA, which is an unintuitive and complex magnetic system, at 6% of the cost. Significant unintuitive takeaways include 1) the sensitivity of performance error to magnet position deviation within the array (Fig. 6), 2) the presence of cogging torque (Fig. 8) in discretized arrays, and 3) the use of axial force as a promising strategy for feedback control (Fig. 9). The D-CAHA test results have enriched our understanding of the CAHA beyond the simulated model and theory [27], thereby enabling research into a more performant robotic device at a fraction of the cost. We summarize that the D-CAHA is a research tool for de-risking unknowns and developing confidence in our strategy towards noncontact magnetic end effectors using permanent magnet arrays, and an effective case study for our discrete approximation approach.

C. The Practical Discrete versus Continuous Paradigm

Practical discrete approximation will never perfectly represent a continuous system, but it can be close. To understand the fundamental shortfall of practical discrete approximation, we examine the idealized and simplest discrete system possible, which comprises two off-the-shelf NdFeB cube magnets with parallel magnetization, in a perfectly aligned voxel grid, with zero lattice wall thickness. It will not have the same magnetic behavior as a continuous (fabricated) rectangular NdFeB magnet of the exact same size, magnetization, and magnetic grade, because the cube magnets will have a thin corrosion-resistant coating (typically Ni-Cu-Ni) on their sandwiched surfaces. The cube magnets will likely have fillets and chamfers on that surface edge, which further modifies their geometry. There might also be partial demagnetization or magnetic hysteresis between the two cubes, depending on their magnetic specification and the ambient temperature. For all these reasons and more, the idealized practical discrete system would never exhibit the exact same magnetic field of its continuous counterpart. Still, the magnetic field of a discrete system can be grossly similar to its continuous counterpart and have roughly the same magnitude; with sufficient error tolerance in an application, a discrete system will be acceptable, for example, in the low-cost prototyping phase of developing a mechatronic system.

In accepting these small errors, discrete approximation

TMECH-02-2025-20230

enables faster development times in constructing a magnetic mechatronic system. In fact, with off-the-shelf components and hobby-grade 3D printers, we developed a miniature D-CAHA prototype, end-to-end, in less than a day. This quick design-build-test cycle enables collecting empirical data to develop intuition and unlock rapid iteration across a design space to find the right magnetic architecture for a mechatronic system. In contrast, it is our experience that fabricating a permanent magnet with custom geometries will take several months and incur very high financial costs. Discrete approximation can save time and save costs.

Discrete approximation is also necessary when its continuous counterpart is physically impossible to fabricate, for example, when a designed $\vec{M}(x, y, z)$ cannot be imbued into a single contiguous permanent magnet with any commercially available fabrication equipment. This viewpoint to magnet design may be obvious because it is the reason why magnet arrays exist in the first place; but scaled discretization enables an even further refined resolution of a desired magnetization vector field $\vec{M}(x, y, z)$ that cannot otherwise be imbued into permanent magnets with feasible magnetization equipment, to obtain a more ideal mechatronic system.

Still, there are limits to discretization. Sources of error between simulation and device need to be analyzed (Supplemental S6). The required discrete design may be too complex for assembly (Fig. S1). Continuous geometry becomes necessary when high magnetic forces are required in small spaces, such as when there is no space for a lattice structure. Continuous geometry may be warranted if field accuracy (e.g. in terms of homogeneity) is necessary at very short distances approaching physical contact, and so we mainly consider *noncontact* or *contact-limited* scope in manipulator mechanisms and devices.

Notably, demagnetization of permanent magnets is of concern in both the continuous and discretized paradigms. If the magnetic material has insufficient coercivity (H_c), the magnetic strength (residual flux density) of each given magnetization voxel can be permanently decreased by the magnetic field contributions at that location from all other magnetization voxels (in addition to external fields), permanently altering the resultant magnetic field of the heterogeneous magnet array. It is thus important to choose permanent magnets with sufficient coercivity to resist the total field at any point from all other voxels, a function of all voxel positions and their orientations. For a CAHA, grade N52 sintered NdFeB magnets with thermal ratings of "M" and above have sufficient coercivity [27]. For the D-CAHA, demagnetization was not critical in our proof-of-concept goal.

D. The Benefits of Discretization

Generally, the precision of discretization increases with the scale of the original object and decreases with voxel size. This is intuitive because these changes will increase the resolution of the voxelization grid, and enable voxelization algorithms with more robust coverage and connectivity [32]. However, there is a practical limit to increasing precision by decreasing voxel size, because the corresponding increase in the number

of lattice walls (with finite and nonzero thickness) will decrease packing density of voxels. Precision is further limited by practical voxel size: as cube magnet sizes decrease, they become more difficult to handle and assemble, and specialized techniques may be required; at the sub-millimeter scale, lithography techniques are better suited to construct magnet voxels [3]. We thus find that discrete approximation seems to be particularly well-suited for mechatronic systems in manipulator robot applications, such as with the D-CAHA, because the scale of these systems seems to be very appropriate for discrete approximation.

Apart from time and cost savings, discrete assembly enables three novel abilities for a magnetic mechatronics designer.

Firstly, it allows for feasible geometric optimization of permanent magnets with hitherto difficult features like internal cavities and undercuts (Fig. S4a). Use of these features is desirable, because magnetized material inside a magnet will contribute the least in shaping its external magnetic field, compared with magnetized material on the surface. This is due to the inverse cube law of (point-like) magnets. It stands to reason that the best method to optimize a permanent magnet design (e.g. to minimize its weight subject to a constrained minimum magnetic field strength for a specific application) would involve material removal from the center of the magnet. Conventionally, fabrication of such a feature is difficult, and even if successful, the magnet may collapse into itself by virtue of buckling from excessive magnetostriction around this hollow and unsupported void. But using discrete approximation, some magnet voxels can be substituted with a non-magnetic, lightweight, and strong material during the assembly process (e.g. a 3D-printed plastic cube) to effectively create these magnetically hollow and undercut geometries (Fig. S4a), without fear of buckling or collapse.

Secondly, cube magnets can be reconfigured and recycled into different applications after each design iteration (Fig. S4b). Holes can be strategically added to the bottom of a lattice scaffold (Fig. 4d) and housing structure (Fig. 4c) to allow a tool to push out cube magnets that were previously press-fit. So long as no glue or adhesive is used, magnet columns in each segment can thus be easily extracted, such as by following Fig. 3 in reverse. Amidst the burgeoning scarcity of rare earth elements [25], this ability to reuse and recycle permanent magnets is sustainable in reducing waste and cost effective for iterating on designs.

Lastly, discrete assembly enables modular scaling and physical adjustment of a permanent magnet array (Fig. S4c). Magnet cubes can be simply added or removed from a system to increase or decrease its field strength and tune its magnetic force for different applications. This is desirable because it improves the *robustness* of a mechatronic system and ensures it can be performant across a wider range of applications. For example, this means that we can adjust the force-distance- θ performance of our D-CAHA (Fig. 7) to meet different application requirements, by simply adding or removing more cube elements, without redesigning or fabricating an entirely new system. The ability to modularize and scale magnetic performance thus creates *customization* potential, which is

TMECH-02-2025-20230

especially useful in applications like wearable prosthetic devices where magnetic forces and torques need to be adjusted for different individuals with different biomechanics [14].

E. Open Problems

We have framed discrete approximation with the intent of taking an optimal continuous design and forming an accurate reproduction of its magnetic field using off-the-shelf magnets. It remains an open problem to improve this discretization process – our illustration with the D-CAHA was largely based on heuristics and intuition during the design process, which resulted in the visible gray spaces of Fig. 4b. Although our efforts were successful, we believe that more can be done to maximize discretization accuracy when compared with the original, continuous design. It is worth exploring the use of *non-uniform* grid spacing to achieve a mix of voxels with different physical sizes, and deeper analysis is warranted in varying voxel vector magnitudes *within* the same discrete design to reduce field errors seen in Supplemental S1 and S2. Hybrid approaches that create partially discretized and partially continuous arrays may also yield many benefits of full discretization. It is finally, very enticing to investigate self-interlocking magnet geometry (instead of cubes) that remove the need for a lattice structure, such as with permanent magnets in LEGO building block shapes [28] that would allow simple pick-and-place assembly. This is not simple, and would involve interdisciplinary research between mathematical topology, magnetics, and material science, to develop magnets with tileable shapes that will not fracture under the magnetic forces that seek to push, pull, and twist each magnet voxel.

We may also invert our framing for discrete approximation. Instead of making a discrete array that approximates an optimal continuous design, one may simply take an *un-optimized* continuous magnet architecture, crudely discretize it, then work to optimize the discrete design for some specific mechatronic application (Fig. S4a). For example, we may take a CAHA design, crudely discretize it, then optimize its force-to-weight ratio by removing voxels in the centroids of each segment. It remains an open problem to formulate a generally robust approach for this process. It may be plausible to integrate magnetic physics into geometric topology optimization methods [30] and generate these optimized designs with empty and filled voxels. Furthermore, since *optimality* in this context is defined by magneto-mechatronic performance (not simply packing density or volume fraction – Supplemental S1), then larger unmagnetized regions may be beneficial (Supplemental S2). This means that optimization may involve choosing the exact placement and orientation for each voxel, and so gradient-descent approaches are worth exploring, for example, to perturb voxel orientations and positions in all directions, *beyond* the cardinal directions.

There is, finally, an open problem to better understand the efficiency and durability of a discrete magnet, when compared to a continuous magnet. While our work in Supplemental S1 and S2 suggest some theoretical parameters, there are other practical considerations such as magnet type (NdFeB, Alnico, Ferrite, etc.) and coating (NiCuNi, parylene, etc.), and their

gross performance in a discrete magnet is worth investigation.

F. Conclusion

In this manuscript, we laid out the key concepts in practical discrete approximation of permanent magnet arrays. We believe that this method to mechanism prototyping will play a vital role in the development of complex magnet arrays that unlock new mechatronic capabilities, especially in the realm of mechatronic systems for robotics. In developing magnetic mechanisms, the main benefits of discretization are cost and time to iterate and develop design intuition. We illustrated this with our design, build, and test of the D-CAHA. Discretization also yields further capabilities of optimization, reconfiguration, and modularity in magnet array development. Ultimately, we hope to fuel further research into mechatronic systems that utilize novel permanent magnet arrays.

ACKNOWLEDGMENT

The authors thank Dean Chen for his academic guidance and the support he provided concerning FEA simulations. The authors HKL, TRC, WF, CRT are named inventors on a provisional patent filing that describes this technology. The provisional patent is owned by UCLA.

REFERENCES

- [1] J. J. Abbott, E. Diller, and A. J. Petruska, "Magnetic Methods in Robotics," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 3, no. Volume 3, 2020, pp. 57–90, May 2020, doi: 10.1146/annurev-control-081219-082713.
- [2] D. Son, M. C. Ugurlu, and M. Sitti, "Permanent magnet array-driven navigation of wireless millirobots inside soft tissues," *Sci. Adv.*, vol. 7, no. 43, p. eabi8932, Oct. 2021, doi: 10.1126/sciadv.abi8932.
- [3] J. Zhang *et al.*, "Voxelated three-dimensional miniature magnetic soft machines via multimaterial heterogeneous assembly," *Sci. Robot.*, vol. 6, no. 53, p. eabf0112, Apr. 2021, doi: 10.1126/scirobotics.abf0112.
- [4] N. Ebrahimi *et al.*, "Magnetic Actuation Methods in Bio/Soft Robotics," *Adv. Funct. Mater.*, vol. 31, no. 11, p. 2005137, 2021, doi: 10.1002/adfm.202005137.
- [5] Y. Wang *et al.*, "An electromagnetic vibration energy harvester using a magnet-array-based vibration-to-rotation conversion mechanism," *Energy Convers. Manag.*, vol. 253, p. 115146, Feb. 2022, doi: 10.1016/j.enconman.2021.115146.
- [6] Q. Zhang and E. S. Kim, "Vibration Energy Harvesting Based on Magnet and Coil Arrays for Watt-Level Handheld Power Source," *Proc. IEEE*, vol. 102, no. 11, pp. 1747–1761, Nov. 2014, doi: 10.1109/JPROC.2014.2358995.
- [7] J. Quan, M. Zhu, and D. Hong, "Re-Examining Climbing Robots: Design and Performance of a Lightweight, Low-Cost Robot with a Highly Extendable Limb," in *2024 6th International Conference on Reconfigurable Mechanisms and Robots (ReMAR)*, Jun. 2024, pp. 409–416. doi: 10.1109/ReMAR61031.2024.10619924.
- [8] D. Roy, "Development of novel magnetic grippers for use in unstructured robotic workspace," *Robot. Comput.-Integr. Manuf.*, vol. 35, pp. 16–41, Oct. 2015, doi: 10.1016/j.rcim.2015.02.003.
- [9] A. Peidr , M. Tavakoli, J. M. Marin, and  . Reinoso, "Design of compact switchable magnetic grippers for the HyReCRo structure-climbing robot," *Mechatronics*, vol. 59, pp. 199–212, May 2019, doi: 10.1016/j.mechatronics.2019.04.007.
- [10] M. B. Khamesee, N. Kato, Y. Nomura, and T. Nakamura, "Design and control of a microrobotic system using magnetic levitation," *IEEE/ASME Trans. Mechatron.*, vol. 7, no. 1, pp. 1–14, Mar. 2002, doi: 10.1109/3516.990882.
- [11] L. Song *et al.*, "Motion Control of Capsule Robot Based on Adaptive Magnetic Levitation Using Electromagnetic Coil," *IEEE Trans. Autom. Sci. Eng.*, vol. 20, no. 4, pp. 2720–2731, Oct. 2023, doi: 10.1109/TASE.2022.3201966.

TMECH-02-2025-20230

[12]G. Ciuti, P. Valdastri, A. Menciassi, and P. Dario, "Robotic magnetic steering and locomotion of capsule endoscope for diagnostic and surgical endoluminal procedures," *Robotica*, vol. 28, no. 2, pp. 199–207, Mar. 2010, doi: 10.1017/S0263574709990361.

[13]H. Su *et al.*, "State of the Art and Future Opportunities in MRI-Guided Robot-Assisted Surgery and Interventions," *Proc. IEEE*, vol. 110, no. 7, pp. 968–992, Jul. 2022, doi: 10.1109/JPROC.2022.3169146.

[14]W. Flanagan *et al.*, "Prosthetic Limb Attachment via Electromagnetic Attraction Through a Closed Skin Envelope," *IEEE Trans. Biomed. Eng.*, vol. 71, no. 5, pp. 1552–1564, May 2024, doi: 10.1109/TBME.2023.3342652.

[15]W. A. V. Stepson, A. D. I. M. Amarasinghe, P. N. R. Fernando, and Y. W. R. Amarasinghe, "Design and development of a mobile crawling robot with novel halbach array based magnetic wheels," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 6561–6566. doi: 10.1109/IROS.2017.8206567.

[16]J. Boisclair, P.-L. Richard, T. Laliberté, and C. Gosselin, "Gravity Compensation of Robotic Manipulators Using Cylindrical Halbach Arrays," *IEEEASME Trans. Mechatron.*, vol. 22, no. 1, pp. 457–464, Feb. 2017, doi: 10.1109/TMECH.2016.2614386.

[17]J. Liu, H. Choi, and M. Walmer, "Design of Permanent Magnet Systems Using Finite Element Analysis," *J. Iron Steel Res. Int.*, vol. 13, pp. 383–387, Jan. 2006, doi: 10.1016/S1006-706X(08)60214-9.

[18]J.-S. Choi and J. Yoo, "Design of a Halbach Magnet Array Based on Optimization Techniques," *IEEE Trans. Magn.*, vol. 44, no. 10, pp. 2361–2366, Oct. 2008, doi: 10.1109/TMAG.2008.2001482.

[19]J.-M. Grenier, R. Pérez, M. Picard, and J. Cros, "Magnetic FEA Direct Optimization of High-Power Density, Halbach Array Permanent Magnet Electric Motors," *Energies*, vol. 14, no. 18, Art. no. 18, Jan. 2021, doi: 10.3390/en14185939.

[20]S. Constantinides, "Manufacture of Modern Permanent Magnet Materials." Arnold Magnetic Technologies Corp, 2014. [Online]. Available: <https://www.arnoldmagnetics.com/wp-content/uploads/2017/10/Manufacture-of-Modern-Permanent-Magnet-Materials-Constantinides-PowderMet-2014-ppr.pdf>

[21]X. Zhou, M. P. Paranthaman, and J. W. Sutherland, "Comparative Techno-economic Assessment of NdFeB Bonded Magnet Production: Injection Molding versus Big-Area Additive Manufacturing," *ACS Sustain. Chem. Eng.*, vol. 11, no. 36, pp. 13274–13281, Sep. 2023, doi: 10.1021/acssuschemeng.3c01942.

[22]J. Ormerod and S. Constantinides, "Bonded permanent magnets: Current status and future opportunities (invited)," *J. Appl. Phys.*, vol. 81, no. 8, pp. 4816–4820, Apr. 1997, doi: 10.1063/1.365471.

[23]D. Brown, B.-M. Ma, and Z. Chen, "Developments in the processing and properties of NdFeb-type permanent magnets," *J. Magn. Magn. Mater.*, vol. 248, no. 3, pp. 432–440, Aug. 2002, doi: 10.1016/S0304-8853(02)00334-7.

[24]Y. Yang *et al.*, "Comparative Life Cycle Assessment and Life Cycle Cost Analysis of Bonded Nd-Fe-B Magnets: Virgin Production versus Recycling," *Sustainability*, vol. 16, no. 19, Art. no. 19, Jan. 2024, doi: 10.3390/su16198599.

[25]H. Jin *et al.*, "Life Cycle Assessment of Neodymium-Iron-Boron Magnet-to-Magnet Recycling for Electric Vehicle Motors," *Environ. Sci. Technol.*, vol. 52, no. 6, pp. 3796–3802, Mar. 2018, doi: 10.1021/acs.est.7b05442.

[26]N. Taberlet, J. Ferrand, and N. Plihon, "Stability Analysis of an Array of Magnets: When Will It Jump?," *Phys. Rev. Lett.*, vol. 120, no. 26, p. 264301, Jun. 2018, doi: 10.1103/PhysRevLett.120.264301.

[27]W. Flanagan, H. K. Lim, C. R. Taylor, and T. R. Clites, "A Controllable Axial-Flux Halbach Array," 2025. doi: 10.2139/ssrn.4793044.

[28]R. A. H. Gower, A. E. Heydtmann, and H. G. Petersen, "LEGO: Automated Model Construction: 32nd European Study Group with Industry," *32nd Eur. Study Group Ind. - Final Rep.*, pp. 81–94, 1998.

[29]Z. Wang, P. Song, and M. Pauly, "State of the Art on Computational Design of Assemblies with Rigid Parts," *Comput. Graph. Forum*, vol. 40, no. 2, pp. 633–657, 2021, doi: 10.1111/cgf.142660.

[30]T. Kollsker, "Mathematical Models and Algorithms for Optimisation of the LEGO Construction Problem," Ph.D. dissertation, Department of Technology, Management and Economics, Technical University of Denmark, Lyngby, Denmark, 2020.

[31]D. Cohen-Or and A. Kaufman, "Fundamentals of Surface Voxelization," *Graph. Models Image Process.*, vol. 57, no. 6, pp. 453–461, Nov. 1995, doi: 10.1006/gmp.1995.1039.

[32]M. Aleksandrov, S. Zlatanova, and D. J. Heslop, "Voxelisation Algorithms and Data Structures: A Review," *Sensors*, vol. 21, no. 24, p. 8241, Dec. 2021, doi: 10.3390/s21248241.



He Kai Lim received a B.S. in mechanical engineering from the University of California, Los Angeles (UCLA), Los Angeles, CA, USA in 2023.

He is currently a doctoral candidate in the department of mechanical and aerospace engineering at the same institution. His research interests concern the development and application of novel permanent magnet devices towards biomechatronic applications, especially in addressing multidisciplinary clinical problems.



Will Flanagan received a B.S. in mechanical engineering with a minor in physiology from the Georgia Institute of Technology, Atlanta, GA, USA in 2020.

He is currently working towards the Ph.D. degree in the Department of Mechanical and Aerospace Engineering at the University of California, Los Angeles, Los Angeles, CA, USA.



Cameron R. Taylor received the B.S. degree in electrical engineering from Brigham Young University, Provo, UT, USA, in 2014 and the M.S. and Ph.D. degrees in media arts and sciences from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2016 and 2020.

He held a postdoctoral scholar appointment with the Department of Mechanical and Aerospace Engineering at the University of California, Los Angeles, Los Angeles, CA, USA from 2023–2025. He is starting a position as an Assistant Professor of Biomedical Engineering at the University of North Carolina at Chapel Hill and NC State University, Chapel Hill, NC, USA, in 2025. His research interests include electromagnetics (particularly magnetoquasistatics), algorithm design, and bioinstrumentation.



Tyler R. Clites received the B.S. degree in engineering sciences from Harvard College, Cambridge, MA, USA in 2014, and the Ph.D. degree in medical engineering and medical physics from the Health Sciences and Technology program, Massachusetts Institute of Technology, Cambridge, MA, USA, and Harvard Medical School, Boston, MA.

He is currently an Assistant Professor of Mechanical and Aerospace Engineering, Bioengineering, and Orthopaedic Surgery with the University of California, Los Angeles, CA, USA. His research interests include the field of bionics, and more specifically anatomics, which he defines as co-engineering of body and machine in pursuit of improved human function.

SUPPLEMENTAL SECTIONS

S1. Case Study in Theoretical Voxelization of a Magnet

Existing literature on voxelization of 3D manifold objects is typically concerned with volumetric and geometric accuracy, especially in applications of computer vision and computer graphics [32]. In voxelization of permanent magnet segments, we are similarly constrained by volumetric boundary surfaces of an original 3D manifold object, for example, if the voxelized assembly must fit within a mechanical housing, or if adjacent segment magnetizations must be respected. However, in mechatronic applications the intent of a discretely approximated permanent magnet array is to use the magnetic field in performing a mechatronic function across a distance (e.g. to apply a force or to *sense* the state of a system). To this end, conventional metrics for quantifying geometric voxelization accuracy may be misleading, because the practical goal here is to replicate *function*, not *form*.

To illustrate the complexity of both *magnetics* and *geometry* in theoretical voxelization of permanent magnet segments, we formulated a simple, arbitrary, permanent magnet system (Fig. S1), and compared the simulated performance of two voxelization parameters: voxel vector strength (magnetization N-grade), and volume fraction (i.e. the ratio of summed volume from all voxels, to the volume of the original continuous geometry).

In the continuous domain, this arbitrary system uses a spherical, N30 sintered NdFeB magnet, to apply attractive force onto a Ferromagnetic SS403 plate at a fixed distance of 250 mm from the center of the sphere (Fig. S1). The sphere was set to 173.2 mm diameter (which is exactly sufficient to contain a 100 mm cube), and this yielded 3.31 N force onto the plate. We then applied theoretical voxelization onto this spherical magnet and obtained three designs: “Red” (13763

voxels, 86.74% volume fraction), “Blue” (384 voxels, 65.34% volume fraction), and “Green” (33 voxels, 44.92% volume fraction) (Fig. S1). In these designs, we make no claim that they maximize volume fraction for a given voxel size, but we do assert that they are always fully contained within the geometry of the original sphere. Finally, we used JMAG software (JSOL Corporation, Tokyo, Japan) to run FEA simulations on each of these designs and determine their attractive force on the steel plate, for different magnetization strength (“grade” of sintered NdFeB magnet) in all voxels.

Our simulated results are, foremost, intuitive. As volume fraction decreases (Fig. S1 left to right), force decreases (Fig. S2 top to bottom). As the magnetization strength (i.e. N-grade) increases, force increases (Fig. S2 left to right). Interestingly, if all voxels were set at N30, the same as the original sphere magnet, then their *mechatronic performance* (i.e. the force) is *less* than their volume fraction would suggest, as shown by the datapoints that are below their respective dashed-dot volume fraction lines (Fig. S2). This stems from the inverse cube law of magnetics, which means that it is more important to know *where* the un-voxelized regions are in relation to the mechatronic task, than to know the total volume of voxelized regions. Naturally, this introduces great complexity in constructing any generalizable voxelization algorithm, because existing algorithms prioritize geometric accuracy and computational efficiency, not mechatronic performance. Conceivably, a form of voxelized magnetic topology optimization would be useful to solve this algorithmic problem [30], but it is beyond the scope of this manuscript.

Instead of voxelization algorithm optimization, there is a simpler heuristic to improving mechatronic performance: utilizing stronger N-grade voxels. Fig. S2 suggests that higher volume fraction designs are more sensitive to voxel N-grade, indicated by the increasing trendline gradients for increasing

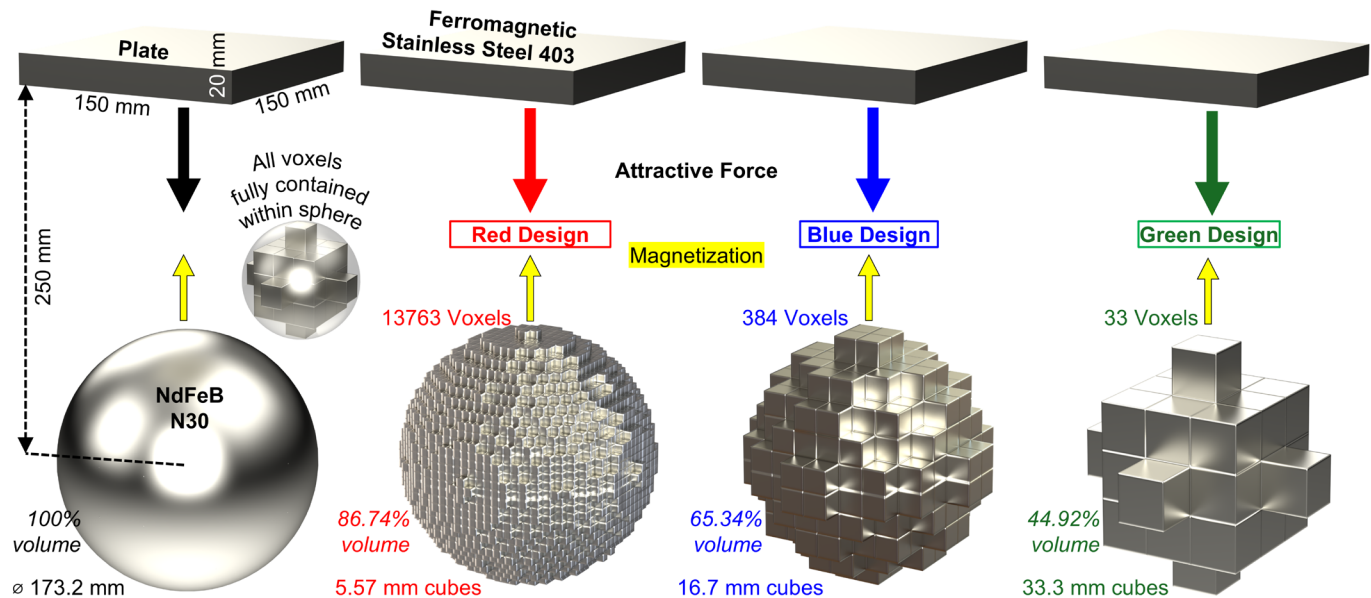


Fig. S1. Theoretical case study in voxelization of permanent magnet segments

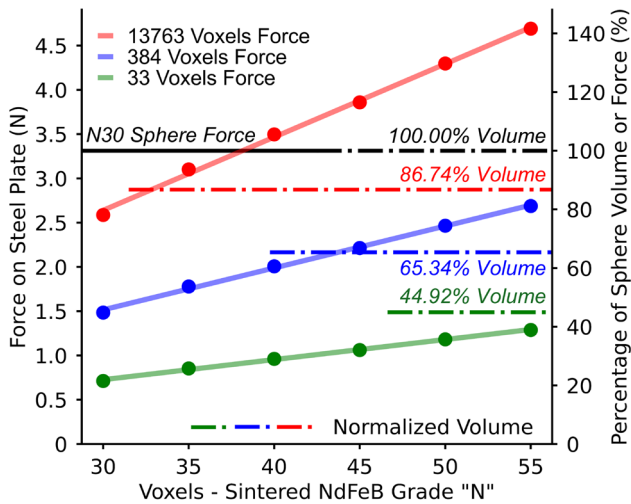


Fig. S2. Simulated case study results for each voxelized design, across different sintered NdFeB magnet grades (“N”). The range of N-grades here are common for off-the-shelf magnets. The nominal force exerted by the original N30 sphere is shown as a constant 3.31 N from the left axis, and it is matched against the right axis at 100% to represent 100% nominal force. Each datapoint can thus be read for its empirical value (left axis), and its relative value (right axis). For each dataset, a simple trendline is fitted for visual clarity, with $R^2 > 0.997$. The volume fraction of each design is further illustrated by dashed-dot lines from the right axis.

volume fraction. It is noteworthy that the Red design can be made functionally equivalent to the original sphere by utilizing magnet grades between N35 and N40 (Fig. S2). It is also noteworthy that a Blue-N55 design was approximately equal to a Red-N30 design in performance (Fig. S2), albeit

with much less assembly complexity. From a financial perspective for the mechatronic engineer, Fig. S2 thus provides food for thought: given equal performance, is it more cost effective to purchase 13763 small and weak magnets, 384 larger and very strong magnets, or a single, large, customized spherical magnet? In practice, other considerations such as stability, assemble-ability, tolerance, robustness, and sustainability may come into play when choosing amongst the illustrated 18 possible configurations (3 designs x 6 possible magnet grades).

Finally, we note that performance of the Green design never matched its volume fraction (Fig. S2) (in our experience, it is very rare to find off-the-shelf magnets stronger than N55), and this underscores the importance of carefully choosing the right metric to quantify voxelization performance. Although these voxelization concepts are heavily inspired by computer graphics and vision systems, this case study highlights the importance of evaluating voxelized magnet segments based on their magnetic performance for a specific mechatronic application.

S2. Voxelization Containment on Internal and External Boundaries of Permanent Magnet Arrays

In permanent magnet arrays, we define “internal” boundaries as boundaries that are shared between two (or more) neighboring segments. Recall that neighboring segments are defined to have different magnetizations, with the intent of creating a non-homogenous magnetic source, to produce an architected magnetic field for a desired magnetic mechanism. On these internal boundaries, there thus exists a transition line for the magnetization vectors (Fig. S3a), which

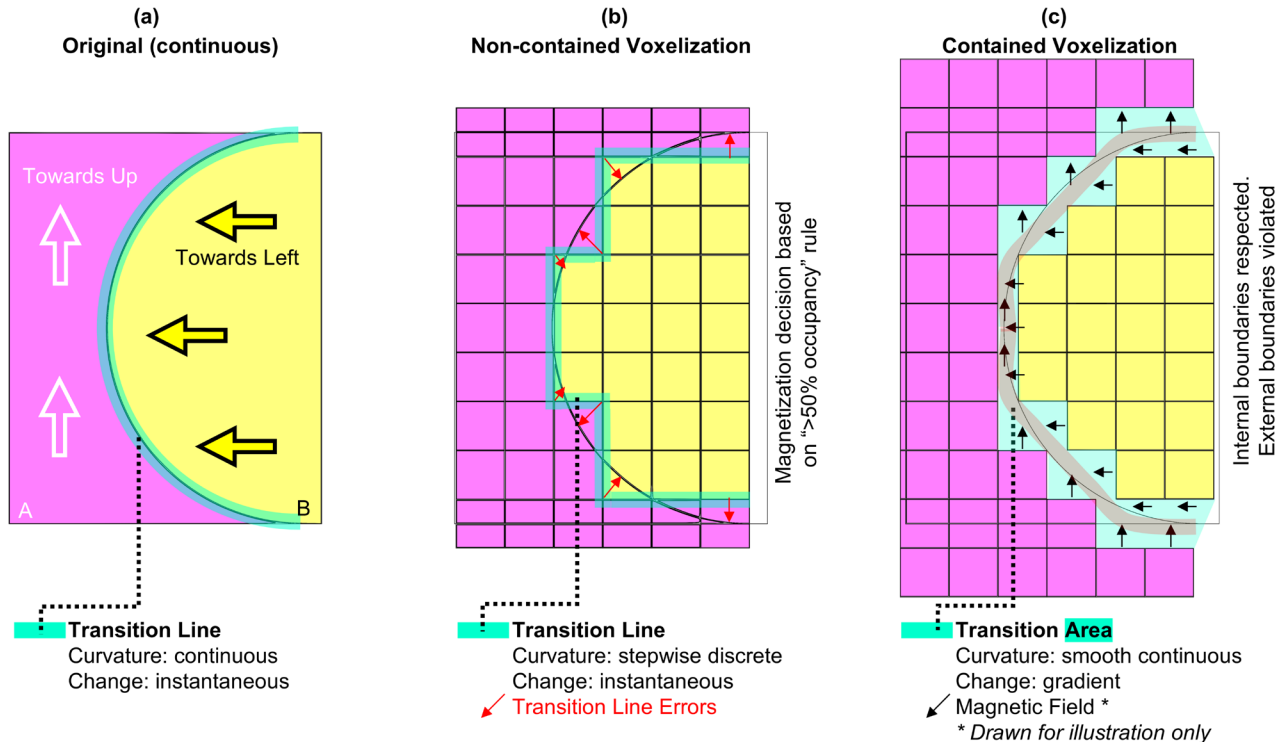


Fig. S3. Examining containment requirements on internal and external boundaries for voxelized permanent magnet arrays.

will have some geometrically defined shape. Notably, we describe this as a transition “line” for clarity in illustration with Fig. S3, but more accurately, it is a 2D surface in \mathbb{R}^3 , with nonmanifold junctions in the intersection of three or more magnet segments. In the most general case for permanent magnet arrays, we may consider these transition lines to have arbitrary continuous geometric curvature (Fig. S3a). This transition line is important, because it describes the unique physical geometry of the original, continuous permanent magnet array.

When permanent magnet arrays are discretized, each segment “A” and “B” is assigned a voxel grid. Along the smooth internal boundaries of the array, there is thus a choice in whether the respective segment’s voxels are strictly contained within the boundary (Fig. S3c), or not (Fig. S3b). This is known as a “containment condition” in present literature for voxelization in computer vision and computer graphics applications [32]. It is also sometimes referred to as “conservative voxelization”, though definitions may vary widely based on specific contexts.

With a non-contained voxelization algorithm, the transition line between segments is stepwise discrete (Fig. S3b). From this, we may consider a geometric discretization error between the voxelized transition line, and the original transition line (Fig. S3b, red arrows). Although this geometric error is *inside* the magnet array, we note that this error will project into the magnetic field outside the magnet array. By definition, these errors occur at the exact locations where the magnet design warranted a unique field created by the unique geometries of each original and adjacent segment. These field errors will dilute and decrease with distance from the magnet array due to superposition effects, and so for a specified magneto-mechatronic application, there will be a minimum working distance beyond which these errors create negligible effects. Some error-minimization is possible by perturbing the voxel grid in half-voxel lengths along each cardinal direction, to determine a “best fit” orientation. Still, for a given continuous curvature, there will always be a fundamental amount of error for a given voxel size, and this fundamental error will only approach zero in the limit of an infinitesimally small voxel.

By contrast, a constrained voxelization algorithm will create a transition “area” with nonzero volume (Fig. S3c). This volume, however, is *not empty*, because it manifests the exuded magnetic field from its adjacent voxels, between the two voxelized segments. This means that internally, within the magnet array, the magnitude of magnetic field errors from discretization will already have had space to dilute and decrease (conceptually annotated with the dark line in the transition area of Fig. S3c). This further implies that the farther magnetic fields *outside* the magnet will manifest less error, exactly where it is important to have less error. Again, error minimization is possible by perturbing each voxel grid in cardinal directions, and simulating the vector field in the unvoxelized region. Additionally, the magnetic strength of each voxel can be varied as well (see Supplemental Section S1) to minimize this error. We note, however, that this minimization problem is expected to be highly nonlinear given

the dependence on nonlinear magnetic physics, and the convexity of this optimization problem may not be guaranteed.

In contrast with internal boundaries, “external” boundaries have no adjacent magnet segment, which means that there is no transition line to be defined; as such, this external boundary need not be respected in designing the discrete array. In fact, we have found it useful in voxelization to apply intentional non-containment along these external boundaries; Fig. S3b and Fig. S3c show the utility in allowing non-containment along these external boundaries, to minimize the transition line error.

Still, from a purely volumetric perspective, containment may be desirable on external boundaries. In fact, the designed, built, and tested D-CAHA in this manuscript adheres to this containment property on its external boundaries (Fig. 4), even though it was not illustrated in Fig. 2b of the main text. Using Fig. S3, the reader is encouraged to consider the cases of external boundary containment, and its effect on the transition lines of Fig. S3b and Fig. S3c. The volumetric need to implement partial coverage is highly application-dependent: some applications may require strict enforcement (e.g. to fit within a housing), and others may have flexibility to extend beyond these external boundaries (e.g. if the housing can be slightly modified to accommodate the non-containment). As such, we do not prescribe any strict containment requirement on external boundaries.

In sum, we believe that voxel containment is highly desirable along internal boundaries, because it provides space for discretization errors to dilute, with respect to the original magnet array. This is important to create the most-accurate approximation of an original permanent magnet array.

S3. D-CAHA Testing: Experimental Setup

The 6DOF load cell (ATI Omega160) on the KUKA tool flange was attached to the weak side of the D-CAHA inner ring, with a 107.5 mm offset to mitigate any residual magnetic field effects on the load cell. The D-CAHA outer ring was mounted onto a test bench with 3D printed support structure. The D-CAHA rings were then assembled to create a controllable magnetic force across an air gap onto a ferromagnetic plate of 1018 grade low-carbon steel (McMaster 7786T241, 139.7 mm in diameter, 12.7 mm thick). Support structures on the plate were used to maintain an air gap with the D-CAHA. Low-cost linear bearings were 3d printed to keep all components axially aligned. These linear bearings also corrected for any off-axis loads, such that only axial force onto the plate was measured (1DOF load cell, Futek LCM300). The linear bearings were further connected with more support structures and mounted onto a manually controlled, nonmagnetic 1DOF linear stage (Velmex UniSlide A6), which controlled the magnetic gap distance d from the D-CAHA strong surface. Note that the smallest magnetic gap possible is $d = 5$ mm, even when there is zero air gap, due to the internal unmagnetized housing of the D-CAHA. All gap distances d are hence measured from the true magnet surface, illustrating the magnetic gap distance. We set the plate at various d and actuated the D-CAHA by rotating the IR

forwards and backwards ($\theta = 0^\circ$ to 180° to 0°) quasi-statically at $6^\circ/\text{sec}$ with the KUKA. During this actuation, we collected 1DOF force (F) data on the plate to characterize the D-CAHA's performance at various distances d . Simultaneously, we collected 6DOF force (F_{IR}) and torque (T_{IR}) on the IR to understand the internal forces of the system, as they relate to changes in distance d and position θ . Data from the 6DOF load cell was collected at 250 Hz directly through the KUKA controller, while data from the 1DOF load cell was collected at 500 Hz with the FUTEK USB225 signal conditioner and time-synchronized with the KUKA.

Rotation beyond $[0^\circ, 180^\circ]$ was not performed because of physical interference between the load cell and the test bench. Nonetheless, the magnetic forces and torques are expected to be axisymmetric given the 180-degree periodicity of the D-CAHA [27]. Explicitly, measured forces in the range of $\theta = [0^\circ, 180^\circ]$ will be theoretically equal to the range of $\theta = [0^\circ, -180^\circ]$, barring tolerance errors in FDM fabrication or lattice structure deformation. Torques in the rotation axis follow the convention of signed displacement from $\theta = 0^\circ$ (a stable equilibrium with zero torque), because a positive torque seeks to increase θ angle, and a negative torque decreases the θ angle. It thus follows that the measured torques in the range of $\theta = [0^\circ, 180^\circ]$ will be exactly opposite in direction (and sign) from $\theta = [0^\circ, -180^\circ]$.

S4. D-CAHA Testing: Data Analysis

We applied quasistatic gravity compensation to remove mass effects of the D-CAHA IR from the 6DOF load cell data. This consisted of load-cell recorded gravity data of the IR with no OR and no plate, such that there was no other magnetic material present during D-CAHA actuation. These IR gravity recordings were organized into a gravity-compensation lookup table based on the position and orientation of the KUKA end effector. Data from each experiment were processed by removing these gravity forces through interpolation from the lookup table. We approached gravity compensation in this manner, instead of using KUKA's built-in gravity compensation tool, to illustrate the most general approach for low-cost testing with no advanced robotics software available.

Our raw data consist of a time series of forces and torques collected in forward-and-backwards θ rotation at various distances d . We generate θ -series data by reordering into ascending θ values, for each controlled distance d . Depending on the results being generated, the θ -series and time-series data were individually processed with a third order Savitsky-Golay filter to address noise while maintaining fidelity in complex trends. The corresponding filter windows are unique to each dataset based on the signal-to-noise ratio, because different configurations of the D-CAHA have different magnitudes of dynamics. Heatmaps were generated by linear gridded interpolation with SciPy 1.14.1 in Python 3.10.11.

S5. Concepts on New Capabilities with Discrete Approximation

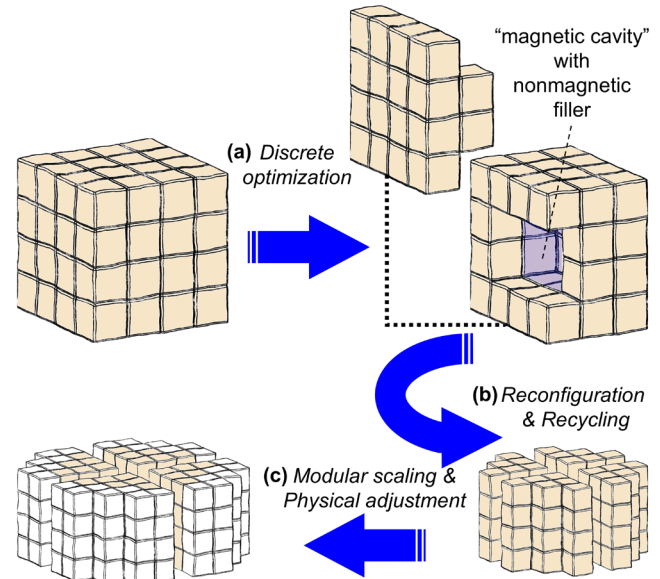


Fig. S4. Concepts of novel capabilities in magnet array prototyping with discrete approximation. (a) Geometric optimization of magnets is now feasible with magnetic undercut and internal cavity features. (b) After prototyping and testing, magnets can be recycled and reconfigured into new geometries for new systems and new applications. (c) Magnet array mechanisms can be modular and adjusted on the spot for applications using the same system, without fabricating wholly new permanent magnets every time.

S6. D-CAHA Results Analysis: Sources of Error and Mitigation Strategies

From Fig. 6 of the main manuscript, while we observed reasonable agreement between simulation predictions and experimental results, there remains a maximum observed prediction error of 84 N (24% error), occurring at occurring at $\theta = 10^\circ$ and the smallest tested gap distance ($d = 7$ mm). Our investigation suggests there are three main sources of error that contribute to this discrepancy, and we discuss these herein.

The first source of error stems from geometric tolerances associated with each individual cube magnet. In large quantities, these tolerances compound into significant discrepancies in magnetic performance between the physical magnet array, and the simulated array based on nominal dimensions. Furthermore, this does not account for the fillet radius around each edge of the cube (not specified by the manufacturer), and the additional plating thickness (NiCuNi) around each cube magnet (for corrosion mitigation) that is included in the nominal cube dimension, both of which further reduce the true volume of magnetized NdFeB material within each magnet. In the example of the D-CAHA, we estimate up to a 6.37% reduction (9088.86 mm^3) in magnetic material compared to our simulated model, based on the least material condition (LMC) of the manufacturer's official tolerance specifications (Table S1). As we illustrate in Supplemental S1

however, volume-based metrics are a poor indicator of true magneto-mechatronic performance, and so the true effects by geometric tolerance on the errors shown in Fig. 6 can only be understood with highly detailed simulation or experimentation. In the original simulation of Fig. 6 (that utilized nominal geometry) the computed force at $d = 7\text{mm}$ for $\theta = 0^\circ$ was 426.0 N, whereas our highly detailed LMC simulation with fillet geometry and minimum material condition resulted in 404.1 N (Table S1). This LMC simulation thus shows a 5.1% reduction (21.9 N) in magnetic force, simply due to tolerance stacking effects from nominal geometries. To mitigate this, better quality control can be sought from manufacturers of cube magnets, and alternatively, the precise dimensions of all cube magnets can be measured and inspected (though this may prove not feasible for large quantities of magnets). Still, these tolerance effects are not sufficient to account for the large prediction errors we observed.

We also note that realistically, for high-volume production components like these cube magnets, a root-sum-squared (RSS) tolerance stack analysis approach is more appropriate than our LMC approach. The manufacturer's specifications do not make it clear whether their tolerances are net values or RSS values, and so we adopt the worst-case LMC approach for present analysis. Furthermore, this singular datapoint of our LMC simulation, having several-orders denser of a mesh refinement (to achieve convergence with cube fillet geometries), required almost 2.5 hours of computation time on a relatively high-end computer (Intel Xeon W-2245 @ 3.9 GHz, configured with 8 degrees of parallelism in JMAG). There was also significant additional time spent to configure the simulation parameters of the more complex model. Altogether, the resource-intensive nature of this LMC simulation approach strongly indicates that it is preferable to simply use a nominal simulation to obtain ballpark values, before directly building and testing a physical discrete device to obtain true magnetic measurements of the system.

Table S1. Table of geometric properties of the D-CAHA, given the manufacturer's listed tolerance of +/- 0.1016 mm (0.004in) on the cube length. The fillet radius is not officially provided but we estimate it to be 0.1mm.

	Nominal	Least Material Condition (LMC) w/ Fillets on Cube Edges
Cube Length	4.7625 mm (3/16 in)	4.6609 mm
Cube Volume	108.0202 mm ³	101.1347 mm ³
D-CAHA Magnetic Volume	142586.664 mm ³	133497.738 mm ³
Percentage Decrease in Magnetic Material Volume	0%	6.37%
Magnetic Force on Plate at $d=7\text{mm}$, $\theta = 0^\circ$	426.0 N	404.1 N

A second likely source of error is plastic deformation in the 3D printed material, which is not directly observable. Specifically, the large forces at $d = 7\text{mm}$ are likely to create plastic deformation in the 3D printed material *within* the structure. While our experiment strived to ensure that the gap distance of $d = 7\text{mm}$ from the surface of the magnets was held accurately, we cannot say for sure whether the position of each cube within the array was held in the same relative positions. This effect is further compounded by the aforementioned geometric tolerance errors that exist in each cube magnet, which leads to further true position inaccuracy of each cube magnet. Considering large forces like this, a possible mitigating solution is to use stronger 3D printed material to construct the lattice scaffolds and housing (e.g. polycarbonate or polyphthalamide nylon), to reduce the magnitude of plastic deformation and minimize the shifting of cube positions within the array. Furthermore, the percentage of 3D printed infill can be increased to provide additional strength, alongside exploration of alternative 3D print orientations that maximize rigidity for the given force profiles. Where resources are available, metal 3D printing or CNC machining approaches can also be used to construct the lattice and housing with suitably strong (and non-magnetic) material. Altogether, while positional error induced by plastic deformation is certainly worrisome, we note that the key benefit of discrete approximation does not lie in high-force applications with field homogeneity, but rather, in being able to construct architected heterogenous magnet arrays that meet some mechatronic objective (like the ability to control axial flux in a permanent magnet system, in the D-CAHA). As such, we interpret Fig. 6 as further evidence as to the value of quickly designing, building, and testing discrete approximations of magnet arrays.

The third likely source of error is the potential demagnetization phenomenon that can occur with permanent magnets of insufficient coercivity, under the influence of strong external magnetic fields. This is discussed in depth in Section VI-C of the main manuscript. In the case of large forces (e.g. for $d < 20\text{mm}$ in Fig. 6), there are certainly very concentrated and strong fields present, and given the lack of a strong thermal rating for the magnets used in this experiment (thermal ratings are correlated with coercivity in off-the-shelf magnets), it is likely the demagnetization occurred and reduced the magnetic force in this experiment. Given that the authors lack access to appropriate test equipment, we cannot empirically ascertain the extent to which demagnetization occurred in the D-CAHA. Still, future use of magnets with higher coercivity will minimize this phenomenon. Nonetheless, through our case study with the D-CAHA, these various sources of error highlight the critical importance of rapidly designing, building, and testing discrete approximations to obtain empirical data, which enriches both the model and the overall understanding of the system's performance.