Certified architecture transfer from polyhedral packing catalogues to finite-cluster magnetic response

Alejandro Zarzuelo Urdiales
alejandrozarzuelo@gmail.com

27 October 2026
Version 1.1

Abstract

Dense packing is a geometric objective, while a material response also depends on arrangement and interactions. We study the transfer of a credited, certified 28-cube packing witness into idealized magnetite-like magnetic architectures. Robust exclusion proofs and a bounded Lean geometry layer are kept separate from the uncalibrated, finite-cluster equilibrium macrospin model. A new surface-charge interaction implementation is checked against independently reduced integrals and adaptive quadrature. Forty-eight fresh, independently seeded runs then validate three previously selected comparisons at 298 K and 10 mT. A fully aligned catalogue derivative yields 32.94 kA m^{-1} per cavity volume, versus 20.47 kA m^{-1} for a matched-count, matched-cavity, matched-body-frame directional grid: a 60.94% increase, with approximate 95% Monte Carlo interval 60.64% to 61.25%. At identical centres and density, source rotations yield 8.52% more response than half alignment, with interval 8.31% to 8.73%. This intervention jointly changes finite-body tensors and registered crystal anisotropy. At the same nominal operating point, compacting the source changes the response by -2.85%, with approximate 95% Monte Carlo interval $[-3.56, -2.13]\%$, despite increased loading. This finite-run estimate does not imply that compression generally lowers response. Intervals condition on the nominal numerical Hamiltonian and exclude physical-model uncertainty. The contribution is a traceable catalogue-transfer benchmark with controlled comparisons and numerical validation. Geometry-dependent magnetism and geometric seeding of functional design are established precedents. No global design superiority, fabrication route, experimental performance or rigorous continuum-response certificate is claimed.

1 Introduction

The arrangement of constituent particles is a design variable in particulate materials. Density constrains how much active material fits in a device, while relative positions and orientations determine which interactions act between those constituents. Geometric packing is therefore relevant to materials design without being identical to it. Dense-polyhedron searches, for example, provide explicit configurations and structural hypotheses [1, 2]. Experiments on polyhedral silver nanocrystals connect shape and packing to assembled structures [3]. These precedents do not imply that a densest configuration optimizes a magnetic, optical or mechanical observable, nor that an arbitrary finite packing will self-assemble.

Magnetic nanocubes provide a concrete setting in which this distinction can be tested. Rounded magnetite cubes have been studied under confinement [4]; field-directed assembly of iron-oxide cubes depends on particle and process conditions [5]. Directional coupling between cubic bodies and magnetic orientations is already established in nanocube simulations [6]. Finite-size corrections and competing magnetic and surface interactions have also been examined explicitly [7]. The present work does not claim that architecture-dependent magnetism is a new physical principle.

Closer precedents connect particle arrangement and crystal-frame texture to collective properties. Wetterskog et al. link dynamic-susceptibility anisotropy in nanocube assemblies to dipolar coupling and partial texture [8]. Håkonsen et al. combine macrospin Monte Carlo and magnetic superstructures to examine crystal anisotropy, cohesion and mechanical response [9]. Reduced equilibrium magnetization can also decrease as an interacting nanoparticle assembly is made denser [10]. Our fixed-body, finite-field comparison differs from those studies, but does not establish a new general relation between density, orientation and magnetism.

Geometric candidate transfer is also established outside magnetism. Point patterns have seeded photonic-band-gap design [11], and polyhedral colloid calculations connect structural changes with photonic response [12]. A silver-polyhedron dual-structure supercrystal gives greater surface-enhanced Raman scattering response than the denser Minkowski comparison [13]. Thus neither geometric seeding in general nor the distinction between density and function is claimed as new here.

Instead, we ask a narrower question: can an existing, independently checkable packing catalogue supply useful candidate architectures even after its original orientations are simplified? The source is the polyhedral-packing programme described in Zarzuelo's deposited methods and patterns papers [14, 15], its reproducibility archive [16], and a later, specifically pinned community witness [17]. The soft-to-rigid search lineage is acknowledged separately [18]. The chosen witness is a feasible finite configuration, not a proven global optimum.

We construct a hierarchy of controls: a simple grid, an expanded source at the same loading, a compact source, rotation-only families at fixed centres, and orientation-adjusted source derivatives with centres moved by a geometric linear programme. A stronger directional grid tests whether a positive result depends only on the initial grid choice. These structures are compared using the same equilibrium magnetic Hamiltonian. The paper retains all 64 historical conditions, including unfavourable adaptations and parameter dependence, and adds six reference-integral validation conditions. Its contribution is a reproducible instance of architecture transfer from a credited, certified finite packing to controlled response comparisons. The numerical validation tests the selected comparisons; it does not establish a general advantage of catalogue seeding over direct property optimization.

This edition makes the transfer mathematically explicit. It separates a proof about ideal exclusion envelopes, a specified equilibrium probability model, numerical evidence for that model, and a future material-level validation. The propositions below expose assumptions and establish consequences used to interpret the comparisons. They are derived here for completeness; no priority claim is made for support-function separation, Gibbs differentiation, detailed balance or covariance bounds. A separate Lean layer checks bounded geometry and normalization results, with the source-to-certificate interface and the unformalized physics identified in Appendix C. This edition adds an independent interaction-integral reference and a fixed-length, independent-seed validation of the three selected contrasts. The earlier exploratory screen remains separately labelled.

2 Geometric source, adaptations and controls

2.1 Provenance and dimensional interpretation

The retained input contains $N = 28$ congruent cubes and is pinned to repository commit `1abc63218db34cea95802aeb3b6d4a6b6fe1c61b`. Its byte-level SHA-256 is recorded in Appendix D. The geometry and original certificate are credited to the contributor `konoha-ily2`; `PingYou Ltd` supplied the recorded padding refinement. The author's catalogue, framework and generalized verification work is distinct from discovery of this particular layout. The deposited archive predates the pinned witness, so the two references are not treated as the same snapshot.

Each body has a nominal exclusion-envelope side $a_e = 24.1\,\mathrm{nm}$ and a magnetic core modelled as a full cube of side $a_c = 22.7\,\mathrm{nm}$. These dimensions are inspired by the rounded-cube system of Wang et al. [4]. The ideal full cube has a different volume and surface from the experimental rounded core; we do not identify the two specimens. The cavity is cubic and open-boundary magnetic interactions are used, whereas the cited confinement experiment is not a cubic-vessel validation.

All candidate geometries use the same declared tolerance check: effective envelope side up to $25.305\,\mathrm{nm}$ (5% above nominal) and simultaneous independent centre-coordinate errors within $\pm 0.1\,\mathrm{nm}$. Orientations are held fixed in this guarantee. Exact rational support inequalities check wall containment and one separating plane for every pair. This certifies the stated exclusion geometry; it does not certify a fabrication process, ligand mechanics or the magnetic response under those perturbations.

2.2 Architecture and certificate objects

Definition 1 (Frozen architecture)**.** *An architecture is* $\mathcal{A} = (N, L, \{\mathbf{c}_i, R_i\}_{i=1}^{N}, a_e, a_c)$*, where* N *is a positive integer,* $L > 0$*,* $0 < a_c < a_e$*,* $R_i \in SO(3)$*, and the nominal exclusion body is* $\mathbf{c}_i + R_i[-a_e/2, a_e/2]^3$*. The magnetic core replaces* a_e *by* a_c*. The particle centres and body frames are fixed inputs, not random variables in the response model. A specification* $\theta = (T, M_\mathrm{s}, K_1, \mathbf{B})$ *with* $T > 0$ *and* $M_\mathrm{s} > 0$ *supplies the material and operating parameters.*

An architecture is admissible when its nominal exclusion cubes are contained in the cavity and have pairwise disjoint interiors. The strict inequality $a_c < a_e$ then separates every pair of magnetic cores by a positive gap. The architecture deliberately separates the exclusion envelope, magnetic core and enclosing cavity. The core filling fraction $\phi_c = N a_c^3/L^3$ and exclusion-envelope fraction $\phi_e = N a_e^3/L^3$ are bookkeeping quantities, not additional Hamiltonian terms. Different certified architectures may have equal ϕ_c but different interactions. The box separating-axis method is classical [19]; the present certificate records a stated robust sufficient condition for the retained numerical designs.

A robust geometry certificate stores $a_{\max} \geq a_e$ and $\epsilon \geq 0$, exact rational centres and rotations, six wall tests per vertex and a nonzero signed separating normal for each unordered pair. All rational rotations must first satisfy $R_i^{\mathsf{T}} R_i = I$ and $\det R_i = 1$. For the normalized-quaternion construction used here this holds algebraically, rather than by accepting a floating-point orthogonality tolerance. The inequalities of Appendix A then have a precise universal meaning.

Proposition 1 (Robust exclusion certificate)**.** *Suppose the wall inequalities and pair inequalities in Appendix A hold exactly at* $a_{\max}$ *and* ϵ*. For every choice of independent sides* $0 < a_i \leq a_{\max}$ *and centre perturbations* $\boldsymbol{\delta}_i \in [-\epsilon, \epsilon]^3$*, all perturbed exclusion cubes are contained in the cavity and have pairwise disjoint interiors. Boundary contact is permitted. Rotations are unchanged.*

Proof. The support of a rotated cube in direction $\mathbf{n}$ is $h_i(\mathbf{n}; a_i)$, and is at most $h_i(\mathbf{n}; a_{\max})$. A centre perturbation changes a coordinate support by at most ϵ, which proves containment from the six signed wall inequalities. In a pair direction, $\mathbf{n} \cdot (\boldsymbol{\delta}_j - \boldsymbol{\delta}_i) \geq -2\epsilon\|\mathbf{n}\|_1$. The certified inequality therefore leaves the upper support of cube i no larger than the lower support of cube j. A separating plane proves interior disjointness. Repeat independently for every pair. The fixed normals make this a sufficient certificate, not a complete description of all feasible configurations. □

For $N = 28$, the retained implementation evaluates $28 \times 8 \times 6 = 1{,}344$ vertex-wall inequalities and $\binom{28}{2} = 378$ pair-plane inequalities. The vertex representation is equivalent to the support test for cubes. This is an exact-arithmetic geometric statement about the encoded data. It is not an error bound for physical shape, crystal registration or magnetic response.

2.3 Baseline hierarchy

The original grid, G, is a radial 28-site cutout of a simple cubic lattice in a cavity of side $L_G = 105.22\,\mathrm{nm}$. The expanded catalogue, C_e, scales the source centres uniformly into the same cavity while preserving its rotations. Hence G and C_e have identical particle count and global loading, although both their centre patterns and body rotations differ. The compact catalogue, C_c, retains the source topology and rotations in a cavity of side $L_c = 96.4340743\,\mathrm{nm}$. Comparing C_c with C_e changes spacing and loading at fixed rotations; it does not isolate loading from the interaction changes caused by compression.

Two additional grids strengthen the baseline. A tighter version of G removes its original extra pitch allowance. A seven-column directional grid, G_z, places four aligned cubes along each of seven parallel z columns. Its transverse sites are the centre, four axial neighbours and two opposite corners. Both tighter grids fit $L = 102.0201020\,\mathrm{nm}$ under the same tolerance criterion. Neither is asserted to optimize the magnetic objective.

2.4 Orientation interpolation and centre adjustment

For body alignment, each source rotation is moved continuously toward its nearest common cubic frame, accounting for the 24 proper rotational symmetries of a cube. The interpolation fraction t takes values $0, 0.25, 0.5, 0.75, 1$ for geometric construction. The mean deviation from the target frame decreases from 15.5133° to 7.7566° at $t = 0.5$, and to zero at $t = 1$. This is deviation from a common frame, not a claim that every neighbour-pair misorientation decreases monotonically.

Rotating bodies can invalidate a packing. The adjusted-centre family therefore retains a selected directed separating plane per pair and solves a restricted geometric linear programme: first minimize the enclosing side, then minimize ℓ_1 centre displacement from the source at effectively that side. A small inflation and independent exact checks protect against rounded optimizer output. This is a geometric projection into a feasible family, not magnetic-energy minimization or a physical relaxation process. Appendix A specifies the support constraints.

The resulting sides for $t = 0, 0.25, 0.5, 0.75, 1$ are respectively 93.6129, 94.9064, 97.6757, 101.4660, and 102.0201 nm. The half- and fully aligned designs move centres by mean distances of 1.4592 and 2.5328 nm; the largest fully aligned displacement is 9.6672 nm. The $t = 0$ LP control is essential: it also moves centres and removes the source's extra pitch allowance. The adaptations are therefore not attributed to rotation alone.

2.5 Strict rotation controls and easy-axis family

To test body rotations independently of density and centre positions, source, quarter-, half- and fully aligned rotations are assigned exactly the same centres in a common 107.1016505 nm cavity. All layouts independently pass the exclusion checks. Within this family, only the body rotations change. Because crystal axes are assumed registered to the cube body, this control changes both the finite-shape interaction tensors and the anisotropy-frame directions.

A distinct family rotates each cube's nearest $\langle 111 \rangle$ crystal direction toward laboratory z, motivated by negative cubic anisotropy. Complete easy-axis alignment is not the same as a common cube-face frame. A separate strict-control cavity of side 120.741167 nm holds common centres for the source and fully easy-axis-aligned variants. The half-rotation geometry exists but was not magnetically simulated in that strict family. Centre-adjusted easy-axis variants at quarter, half and full rotation were simulated; their cavity costs are reported with their responses.

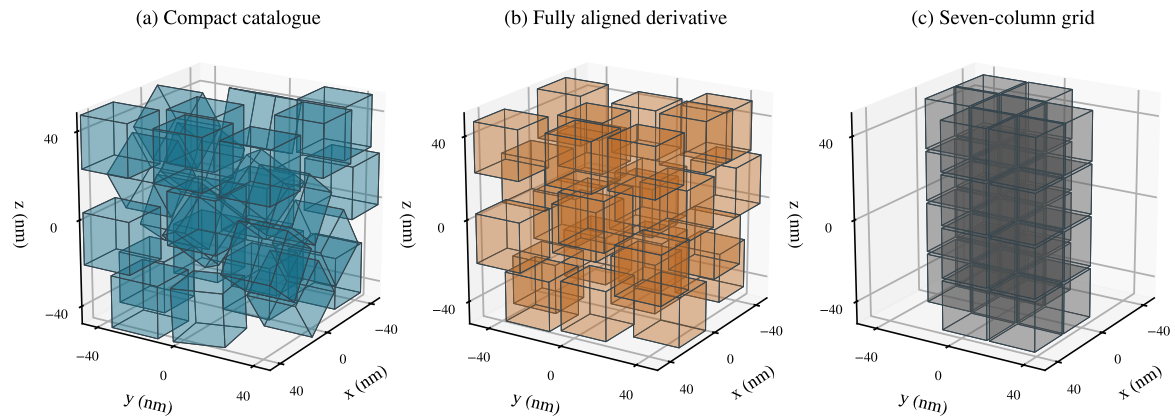

Figure 1: Source and derived architectures. Nominal exclusion envelopes, rather than magnetic core surfaces, are shown. The fully aligned derivative preserves a source-derived centre pattern after constrained adjustment. Its cavity side matches the seven-column grid to approximately 10^{-6} nm; both have the same physical cube orientation up to cubic symmetry. The compact source uses a smaller cavity. These are proposed fixed architectures, not observed assemblies.

3 Magnetic model and numerical methods

3.1 Fixed-architecture equilibrium Hamiltonian

Each core carries a uniform moment $\boldsymbol{\mu}_i = m\mathbf{s}_i$, with $m = M_\mathrm{s}V_c$, $V_c = a_c^3$ and $\|\mathbf{s}_i\| = 1$. Particle centres and body rotations R_i are fixed; moment directions are thermally sampled. With $\mathbf{a}_i = R_i^\mathsf{T}\mathbf{s}_i$, the energy is

$$E = K_1 V_c \sum_{i=1}^{N} (a_{ix}^2 a_{iy}^2 + a_{iy}^2 a_{iz}^2 + a_{iz}^2 a_{ix}^2) - m\sum_{i=1}^{N} \mathbf{B}\cdot\mathbf{s}_i + \frac{\mu_0 m^2}{4\pi}\sum_{i<j} \mathbf{s}_i^\mathsf{T}\mathcal{T}_{ij}\mathbf{s}_j. \tag{1}$$

Here $\mathbf{B}$ is the prescribed uniform external induction entering the Zeeman term. Demagnetizing interactions between particles are already included in $\mathcal{T}$; no additional empirical macroscopic demagnetizing correction is imposed on the finite cluster. Perfect-cube self-demagnetizing energy is orientation-independent for a spatially uniform macrospin and is omitted as a constant.

The nominal parameters are $T = 298\,\mathrm{K}$, $B = 10\,\mathrm{mT}$, $M_\mathrm{s} = 460\,\mathrm{kA\,m^{-1}}$ and $K_1 = -11.5\,\mathrm{kJ\,m^{-3}}$, using the literature-inspired magnetic values in Moreno et al. [20]. Negative K_1 favours $\langle 111\rangle$ axes. These parameters are not fitted measurements of the dimensional-reference specimen. We also calculate selected conditions with $M_\mathrm{s} = 250\,\mathrm{kA\,m^{-1}}$ and, separately, $K_1 = 0$. Neither change is an experimental uncertainty distribution.

No interparticle exchange, surface anisotropy, oxidation layer, ligand force, internal magnetic texture or translational motion is included. Uniform-core moment rotation is an explicit approximation. Faceting, surface and shape effects in a real magnetite particle can require more detailed modelling [20].

3.2 Equilibrium probability space and exact consequences

Let $\Omega = (S^2)^N$ carry the product of normalized sphere-area measures ν. Define $U_{\mathcal{A},\theta} = E/(k_\mathrm{B}T)$ from Eq. (1), and

$$Z_{\mathcal{A},\theta} = \int_\Omega e^{-U_{\mathcal{A},\theta}(s)}\, d\nu(s), \qquad d\pi_{\mathcal{A},\theta}(s) = Z_{\mathcal{A},\theta}^{-1} e^{-U_{\mathcal{A},\theta}(s)}\, d\nu(s). \tag{2}$$

Strictly separated compact cores have finite interaction tensors. Thus U is continuous and bounded on compact Ω, so $0 < Z < \infty$ and every bounded observable has a well-defined expectation. A geometric certificate establishes admissibility of its envelopes; the Gibbs model separately defines the response.

Proposition 2 (Field symmetry and equilibrium response)**.** *Fix an admissible architecture at $T > 0$ and all parameters except the signed field magnitude B along a fixed unit direction $\mathbf{h}$. Let $S(s) = \sum_i \mathbf{s}_i \cdot \mathbf{h}$ and $q(B) = \mathbb{E}_B[S]/N$. Then $q(0) = 0$, $q(-B) = -q(B)$, and*

$$\frac{dq}{dB} = \frac{m}{Nk_{\mathrm{B}}T}\,\mathrm{Var}_B(S) \geq 0, \qquad \frac{dM_{\parallel}}{dB} = \frac{m^2}{L^3 k_{\mathrm{B}}T}\,\mathrm{Var}_B(S) \geq 0. \tag{3}$$

Proof. Global spin inversion $s \mapsto -s$ preserves the pair and cubic-anisotropy energies, sphere measure and architecture, while reversing S. Substitution in Eq. (2) gives the symmetry statements. Compactness permits differentiating the bounded integrand at finite B; the logarithmic derivative of the Gibbs density is $m(S - \mathbb{E}_B S)/(k_{\mathrm{B}}T)$. Multiplication by S/N gives the variance identity. The second identity uses $M_{\parallel} = m\mathbb{E}_B S/L^3$. □

These identities are model properties, not additional measured or simulated results. The retained campaign samples only 10 mT and does not estimate the derivative in Eq. (3). At finite N and positive T, smooth equilibrium response also does not establish a dynamic switching or hysteresis law.

3.3 Finite-volume interaction tensors

For nonoverlapping magnetic core volumes V_i and V_j,

$$\mathcal{T}_{ij} = \frac{1}{V_c^2}\int_{V_i}\int_{V_j}\left(\frac{I}{r^3} - \frac{3\mathbf{r}\mathbf{r}^{\mathsf{T}}}{r^5}\right) dV_j\, dV_i, \qquad \mathbf{r} = \mathbf{x}_j - \mathbf{x}_i. \tag{4}$$

The historical screen computes the tensor in laboratory coordinates with normalized tensor-product Gauss–Legendre volume quadrature. Order p uses p^3 nodes per core and p^6 node pairs per particle pair. All 378 distinct particle pairs are retained. The tensors have SI units m^{-3} in Eq. (1); the saved numerical arrays use nm^{-3} with an explicit 10^{27} conversion.

Finite-body magnetic interactions are standard in demagnetizing-tensor methods [21], and departures from central point dipoles are relevant in magnetic-cube assembly models [7]. Our numerical double-volume integral is not an implementation of an exact analytic cuboid formula. The broad screen uses $p = 7$; nine selected nominal z conditions are rerun with $p = 9$ and fresh seeds. Validation checks cover pair reciprocity, tracelessness, far-field behaviour, global-rotation covariance and pair counting. In the selected original geometries, replacing central point dipoles by finite cores changes interaction tensors appreciably; finite-volume integration is therefore part of the model definition, not a cosmetic correction.

For the new reference we independently apply the divergence theorem twice to the separated cores:

$$\mathcal{T}_{ij,ab} = \frac{1}{V_c^2}\int_{\partial V_i}\int_{\partial V_j}\frac{n_{i,a}(\mathbf{x})\,n_{j,b}(\mathbf{y})}{|\mathbf{y} - \mathbf{x}|}\, dS_y\, dS_x. \tag{5}$$

Indeed, $\partial_{x_a}\partial_{y_b}|\mathbf{y} - \mathbf{x}|^{-1}$ equals the kernel in Eq. (4). For uniform magnetization, each cuboid face carries constant surface charge. The inner rectangular-face Coulomb potential is integrated analytically, leaving only two-dimensional Gauss–Legendre integration over each receiver face. Arbitrary body rotations are handled in source and laboratory frames. All 378 pairs are still included. This is an independently implemented standard magnetostatic representation, not a new interaction law or a claim of a new integration method.

Semi-analytic interaction reductions and polyhedral field calculations are established [22, 23]. Closed-form demagnetizing formulas require attention to cancellation as well as analytic correctness [24]. We therefore compare outer orders 16, 24, 32 and 48 and cross-check selected pairs by a separate analytic-source-field/three-dimensional-target-volume implementation and adaptive integration. Pre-symmetrization reciprocity, tracelessness, rigid-rotation covariance, near-face/edge/corner cases and the far-field sign and units are checked separately. Refinement differences and high-accuracy cross-checks are numerical evidence, not outward-rounded interval certificates. Section 4.1 reports the observed errors and the Hamiltonian actually used for the new trajectories.

3.4 A bound linking tensor error to equilibrium response

Proposition 3 (Conditional response-error bound)**.** *On the same spin space, let U and U' be two bounded measurable dimensionless Hamiltonians. If $\sup_s |U'(s) - U(s)| \le \eta$, then for any measurable observable $f \in [-1, 1]$,*

$$|\mathbb{E}_{\pi'} f - \mathbb{E}_{\pi} f| \le \min(2, \eta). \tag{6}$$

For a change only in the dimensionless interaction tensors, a valid choice is

$$\eta = \sum_{i<j} \|D'_{ij} - D_{ij}\|_{\mathrm{op}}, \qquad D_{ij} = \frac{\mu_0 m^2}{4\pi k_{\mathrm{B}} T} \mathcal{T}_{ij}. \tag{7}$$

Any known upper bound on the right side is therefore admissible in Eq. (6). For $f = S/N$, multiply the result by Nm/L^3 to bound cavity-response error when normalization is unchanged.

Proof. Interpolate $U_t = U + t\Delta U$ for $0 \le t \le 1$. Differentiation under the integral gives $\frac{d}{dt}\mathbb{E}_t f = -\operatorname{Cov}_t(f, \Delta U)$. A variable in $[-1, 1]$ has variance at most one; one in $[-\eta, \eta]$ has variance at most η^2. Cauchy–Schwarz bounds the covariance by η. Integration proves the stated bound, also bounded by the total range two. For unit spins, $|\mathbf{s}_i^{\mathsf{T}} \Delta D_{ij} \mathbf{s}_j| \le \|\Delta D_{ij}\|_{\mathrm{op}}$; summing proves the tensor estimate. □

This proposition supplies a route to a response certificate if a rigorous integration or tensor-error bound is obtained. The observed difference between quadrature orders 7 and 9 bounds neither order's error relative to the continuum integral. It can compare the two discretized Hamiltonians if evaluated as a certified matrix bound, but it cannot by itself certify their distance to the physical material or to the exact integral. An independent binary64 evaluation of the right side of Eq. (7) for the two retained quadrature orders gives approximately 1.134 for the fully aligned derivative and 24.452 for the seven-column grid. Those loose estimates cannot certify the reported response contrast; they are also not interval-certified norm bounds. These historical estimates remain in the supplement. The new surface-reference calculation substantially reduces the observed integration discrepancy and supports fresh simulations, but it also remains a floating-point numerical reference rather than a completed interval-certified continuum-response bound.

3.5 Replica-exchange sampling and diagnostics

For each historical condition, four runs use distinct seeds and overdispersed starts: random directions, field-aligned moments, reversed moments, and another random start. Each run contains 24 replicas with relative inverse-temperature factors geometrically spaced from 0.03 to 1. These auxiliary factors temper the reference Hamiltonian for sampling [25]; the calculation does not claim temperature-dependent physical material parameters along that ladder.

Each sweep visits particles in randomized order. A moment is proposed using the von Mises–Fisher conditional for the linear Zeeman-plus-pair field, then accepted using a Metropolis correction for cubic anisotropy. The implementation uses a uniform proposal at an extremely small local field, as specified in Appendix B. A 15% mixture of involutive reflections about that local field supplies additional proposals preserving the linear conditional. Adjacent replica exchanges alternate even and odd pairs. Cached fields and energies are recomputed every 1,000 sweeps to monitor numerical drift. The same protocol uses 4,000 burn-in and 16,000 production sweeps, retaining every tenth cold-replica sample, hence 1,600 samples per run. Monte Carlo sweeps are not physical time.

We monitor conventional split-$\widehat{R}$ and autocorrelation-based effective sample size (ESS) for the field projection and total energy. The declared exploratory gate is $\widehat{R} \leq 1.05$ and ESS$\geq$ 400 for both observables. It is supplemented by start-specific means, production-half drift, exchange acceptance and whole-ladder round trips. Conventional diagnostics have known limitations [26, 27]; rank-normalized, folded and local diagnostics were not used as the screen's gate. The earlier formal edition added an independent post-hoc rank-normalized and folded split-$\widehat{R}$ audit of the retained chains. The present edition also adds the separate reference-integral trajectories described below. The new validation campaign additionally computes bulk and tail ESS and local-spin diagnostics; those do not retroactively alter the original screen gate. Passage does not prove equilibration of every collective mode. The positive paired-sequence ESS estimate and diagnostics are explicit in the retained code, rather than presented as an independent mathematical convergence proof.

3.6 Observables, normalization and uncertainty

The dimensionless per-core field projection and cavity-normalized magnetization are

$$q = \frac{1}{N}\sum_i \langle \mathbf{s}_i \cdot \widehat{\mathbf{B}} \rangle, \qquad M_\parallel = \frac{N M_\mathrm{s} V_c}{L^3} q. \tag{8}$$

Comparing q removes the explicit loading multiplier, but does not remove spacing-dependent interaction effects. Comparing $M_\parallel$ evaluates moment per chosen cubic cavity volume, not intrinsic saturation magnetization. In particular,

$$\frac{M_\parallel^{(A)}}{M_\parallel^{(B)}} = \left(\frac{L_B}{L_A}\right)^3 \frac{q_A}{q_B} \tag{9}$$

for equal particles and particle count. The common intrinsic M_s does not improve because the architecture changes. A single finite-field response does not establish susceptibility, permeability, coercivity or heating.

Historical tables report per-condition Monte Carlo standard errors (MCSEs); historical plot error bars are 1.96 MCSE. Several historical conditions reuse seed numbers. Their cross-condition covariance was not estimated, so the old ratios remain point estimates. The new campaign uses disjoint seed identities for every run and estimates uncertainty for each selected ratio directly.

Each new condition has eight runs, with 8,000 burn-in and 32,000 production sweeps and thinning by ten. The 48 seeds, six geometries, three contrasts, run length and analysis protocol were frozen before observing the new trajectories. These are numerical validations of previously selected exploratory results, not a preregistered test of catalogue superiority. The update kernels and random-number operations are unchanged; a recording-only modification additionally retains cold-replica spins.

For each new condition, we compare MCSE estimates from independent run means, split-chain autocorrelation spectra and batch means at 50, 100 and 200 retained draws [28]. The primary

ratio interval propagates the largest of these estimates on the log scale, using independent-condition variance addition and a Student critical value with seven degrees of freedom. Replicate-only Welch and Fieller calculations provide sensitivity checks [29]. These are approximate 95% Monte Carlo intervals under Markov-chain central-limit and finite-run mixing assumptions. The variance envelope is a transparent caution, not a theorem of conservative finite-sample coverage. A shared condition or an overlapping interval does not imply independent contrast estimates. All intervals condition on the fixed numerical Hamiltonian, architecture and nominal parameters. They exclude specimen uncertainty, fabrication variation, model discrepancy and an uncomputed rigorous quadrature remainder.

3.7 What each contrast identifies

The response functional is $\mathcal{F}(\mathcal{A}, \theta) = \phi_c M_s \, \mathbb{E}_\pi[S/N]$. A density change modifies the prefactor and generally the distribution π through pair spacing. To isolate a body-rotation change, centres, cavity, core dimensions and magnetic parameters must be fixed. With body-registered crystal axes, that intervention still changes both crystal anisotropy and finite-shape tensors. Removing K_1 removes the former, not the latter.

Contrast	Held fixed	Interpretable change
Grid versus expanded source	N, L, a_c, θ	Centre and body architecture jointly
Shared-centre source versus half aligned	Centres, N, L, a_c, θ	Body/crystal-frame rotations jointly
Fully aligned derivative versus G_z	N, L, a_c, θ, physical body frame	Centre architecture, to stated numerical equality of L
Expanded versus compact source	N, a_c, θ, body rotations	Loading and spacing-dependent interactions jointly

Table 1: Controlled meaning of the principal comparisons. These are interventions in the declared model, not causal estimates for an unmeasured specimen.

For equal particles and N, Eq. (9) is exact whenever the denominator response is nonzero. It gives no monotonicity with density because q changes. Scaling only all centre coordinates is not a uniform rescaling of both magnetic core volumes and separations, and no simple λ^{-3} tensor rule is assumed for that operation.

4 Results

4.1 Independent integral and sampling validation

The new campaign uses surface-order-48 tensors, with every particle pair retained, for six nominal architectures. All seven audited geometries, including the original grid, were evaluated at surface orders 16, 24, 32 and 48. Relative to the new reference, the historical Q9 global Frobenius discrepancy is 0.204154% for the directional grid and 0.021728% for the fully aligned derivative. The corresponding strict-rotation controls are about 0.0146%; the compact and expanded sources give 0.01956% and 0.0003923%. This is a real near-neighbour quadrature issue, not a sign, units or normalization defect.

The largest pair-relative change from surface order 32 to 48 is 1.54×10^{-11} across all 2,646 audited unordered pairs. A separately implemented source-field/target-volume calculation agrees on 23 cases; actual-geometry cases agree to at most 3.29×10^{-14} relative. Adaptive outer integration agrees on seven difficult actual pairs to at most 7.74×10^{-15}. Far-field cancellation is

Architecture	q	MCSE envelope(q)	$M_{\parallel}$ (kA/m)
Expanded catalogue	0.32704	0.00081	42.296
Compact catalogue	0.24459	0.00046	41.090
Fixed centres, source rotations	0.28908	0.00018	35.451
Fixed centres, half aligned	0.26639	0.00014	32.668
Seven-column field grid	0.14426	0.00006	20.469
Fully aligned LP derivative	0.23218	0.00016	32.943

Table 2: Fresh reference-integral estimates at the nominal z-field operating point. Each row contains eight independent runs and 25,600 retained draws. The MCSE is the largest of replicate, spectral and specified batch estimates; it is not experimental uncertainty.

Contrast	Old Q9 gain	New gain	Approx. 95% MC interval
Aligned derivative / field grid	61.15%	60.94%	$[60.64, 61.25]\%$
Source / half rotations, fixed centres	8.34%	8.52%	$[8.31, 8.73]\%$
Compact / expanded source	-4.01%	-2.85%	$[-3.56, -2.13]\%$

Table 3: Selected response gains, $100(M_{\parallel}^{(A)}/M_{\parallel}^{(B)} - 1)$. Old Q9 estimates are historical point estimates with unestimated cross-condition covariance. New intervals use the fixed sample-budget protocol in Section 3 and condition on the reference numerical Hamiltonian at the nominal parameters. In particular, the compression row is a finite-run, condition-specific estimate. These are approximate Monte Carlo intervals, not finite-sample guarantees or physical-model uncertainty intervals.

separately visible at larger separation. These are cross-implementation convergence observations, not certified error enclosures.

Small tensor-relative errors alone would not justify the old response percentages. For Q9 versus surface 48, the empirical right side of Eq. (7) is about 9.106 for the directional grid and 0.280 for the aligned derivative. These are too loose to certify that contrast through Proposition 3. The same empirical quantity for surface 32 versus 48 is at most 4.38×10^{-9}. We therefore use actual new trajectories for the headline comparisons, rather than interpret tensor convergence as a rigorous response bound.

For the two finite numerical Hamiltonians, a narrower exact certificate is available. Treat the actual binary64 dimensionless tensor entries after the sampler's coefficient multiplication as exact rationals. For each architecture, sum all $378 \times 9 = 3{,}402$ absolute entry differences between orders 32 and 48. Since each spin component has magnitude at most one, this entrywise sum bounds the uniform energy difference. An independent exact-arithmetic replay verifies the sums and hashes. Across the six validation architectures the largest bound is 7.801×10^{-9}. Proposition 3 therefore bounds the change in exact Gibbs projection by that number and the cavity response by $0.001107\,\mathrm{A\,m^{-1}}$, using each recorded loading constant. This certifies stability between those two explicitly encoded Hamiltonians only. It does not bound either one's distance to the continuum integral, rounding in the model constants, sampling error or physical-model discrepancy.

The selected architecture and rotation advantages persist in the fresh campaign (Table 3). The architecture gain at matched loading and common physical frame is 60.94%. The source-rotation gain at fixed centres is 8.52%. At $T = 298\,\mathrm{K}$, $B = 10\,\mathrm{mT}$ along z, $M_{\mathrm{s}} = 460\,\mathrm{kA\,m^{-1}}$ and $K_1 = -11.5\,\mathrm{kJ\,m^{-3}}$, the fresh compression estimate is -2.85%, with approximate 95% Monte Carlo interval $[-3.56, -2.13]\%$. Additional loading therefore does not compensate for the lower moment projection in this specified nominal comparison. The interval conditions on the numerical Hamiltonian, not uncertainty in a physical specimen. The rotation intervention still changes both finite shape and crystal-frame anisotropy.

Matched-architecture identification. For the aligned derivative A and directional grid G_z, the retained body/crystal frames are all exactly I_3. Particle count, core dimensions, moment, field, temperature and single-particle anisotropy also match. On the same 28-spin space, Eq. (1) therefore gives the pointwise identity

$$E_A(s) - E_{G_z}(s) = \frac{\mu_0 m^2}{4\pi} \sum_{i<j} \mathbf{s}_i^{\mathsf{T}} (\mathcal{T}_{ij}^{A} - \mathcal{T}_{ij}^{G_z}) \mathbf{s}_j. \tag{10}$$

With all pair tensors set to zero, the two Gibbs distributions, and hence their per-core responses q, are identical. This follows algebraically from the common single-particle terms and measure; it is not a new simulation. Thus the model term changed by the centre intervention is identified independently of a proposed microscopic pathway. The aligned layout's loading is actually 0.000003% lower. Replaying the normalization changes a 60.94128888% projection gain to a 60.94128405% cavity-response gain, only −0.00000483 percentage points. Higher loading or a more favourable single-particle easy-axis frame cannot account for this contrast.

The retained spin means describe how the responses differ. Write $q_i = \langle s_{iz} \rangle$, $q_+ = N^{-1} \sum_i \max(q_i, 0)$ and $q_- = N^{-1} \sum_i \min(q_i, 0)$, so $q = q_+ + q_-$. Table 4 sums all 28 sites. The opposing contribution becomes less negative by 0.11204114, more than offsetting a 0.02412458 decrease in the positive contribution; their difference is the observed $\Delta q = 0.08791657$. This is an exact decomposition of the retained mean estimates, not an allocation to individual couplings.

Architecture	q_+	q_-	q	c_{all}
Directional grid G_z	0.439095	−0.294831	0.144264	−0.014811
Aligned derivative A	0.414971	−0.182790	0.232181	+0.021239

Table 4: Post-hoc organization summaries from the same eight retained runs per architecture. q_- is signed. The raw all-pair correlation $c_{\mathrm{all}} = 378^{-1} \sum_{i<j} \langle \mathbf{s}_i \cdot \mathbf{s}_j \rangle$ uses every unordered pair in both layouts, without a nearest-neighbour selection or tie rule. It is a total-moment second-moment summary, not a connected correlation or independent confirmation of the gain.

All four sites in each of three grid columns have negative mean z projection, while four columns have positive means. The grid's raw longitudinal products average +0.54515 over its 42 within-column pairs and −0.06695 over its 336 between-column pairs. Independent all-pair energy replay of the 51,200 retained configurations for these two architectures agrees with stored energies within $6.24 \times 10^{-11} k_{\mathrm{B}} T$. These checks document the computed organization and its consistency with the specified interactions. They do not prove finite-time equilibration or identify a unique bond family, reduced frustration, assembly pathway or general catalogue mechanism.

Paired self-normalized importance reweighting of the new cold-spin trajectories estimates the change back to the historical Q9 Hamiltonians. Q9 minus reference shifts the architecture, rotation and compression gains by approximately +0.1337, +0.00346 and −0.02141 percentage points, respectively. The minimum Kish weight fraction is 99.759%; correlated-chain influence estimates and paired uncertainty are retained. Thus the estimated Q9-to-reference effect is small for these contrasts but was not assumed to be zero.

Compression: retained claim and cross-campaign check. The fresh nominal result has a direct normalization explanation. Equation (9) gives

$$\frac{M_{\parallel}^{(C_c)}}{M_{\parallel}^{(C_e)}} = \underbrace{1.29898263}_{\text{loading ratio}} \underbrace{0.74787904}_{\text{projection ratio}} = 0.97148188. \tag{11}$$

Loading rises 29.90%, but per-core projection falls 25.21%. Break-even would require a projection ratio of 0.769833; re-expressing the existing approximate 95% Monte Carlo interval gives $[0.742393, 0.753405]$, wholly below that threshold. This is the same interval in different units, not an independent test. It supports the stated -2.85% nominal-condition estimate under the sampling assumptions, without a general density penalty.

The historical gain is -4.00559%, compared with -2.85181% fresh, a 1.15378-percentage-point shift. The historical projection MCSE envelopes are 0.0005794 for the compact source and 0.0009825 for the expanded source. Varying their unknown cross-condition error correlation over $[-1, 1]$ in first-order log-ratio propagation gives historical gain MCSEs from 0.058 to 0.516 percentage points; the fresh gain MCSE is 0.302 points. Four reused seed pairs provide too little information to treat a particular covariance assumption as settled. The shift is therefore assessed against sampling uncertainty rather than assumed negligible.

For a same-Hamiltonian comparison, importance reweighting of the fresh retained trajectories to the historical Q9 tensors predicts -2.87322%, with approximate Monte Carlo interval $[-3.58596, -2.15521]\%$. The paired Q9-minus-reference shift is only -0.02141 points, with approximate interval $[-0.02505, -0.01777]$ points. This uses the same saved spins and is not an independent Q9 rerun. A 1.13237-point gap from the historical Q9 estimate remains. The estimated tensor refinement does not explain the gap, and no covariance assumption is treated as a resolution. Its full source remains unresolved. The fresh campaign has eight runs per condition with 32,000 production sweeps per run, versus four with 16,000 historically. The current quantitative claim uses the fresh estimate and interval only. The historical magnitude is retained as history, and the lower-moment reversal prevents extending this sign to other parameters.

All 48 fixed-length runs completed, adding 153,600 retained draws without replacing the original 256 chains. Post-hoc diagnostic checks cover projection, energy, 28 local projections, 28 local cubic quantities and a neighbour-correlation statistic, 354 condition-observable combinations. Across these new checks, the maximum rank-normalized or folded split-$\widehat{R}$ is 1.008915, the minimum bulk ESS is 1423 and the minimum tail ESS is 3578. Start-specific means, quarter/half drift, exchange statistics and energy checks are retained in the supplement. The 1.01 rank diagnostic and 800-ESS thresholds were applied post hoc; they were not the campaign's prespecified acceptance gate. These diagnostics support the finite-run calculation without proving that every collective slow mode has mixed. The recording-only code change passes an exact fixed-seed equivalence test for the original observables and does not change the update or random-number sequence.

The original Q7/Q9 figures and tables below remain an exploratory directional and sensitivity record. They use different approximate Hamiltonians and should not be mistaken for the new six-condition campaign. In particular, the lower-moment and zero-anisotropy panels were not all rerun with the surface reference. Physical-model and fabrication uncertainty have no measured probability law here and are not assigned an invented confidence interval.

4.2 Historical exploratory screen at equal loading

Figure 2 shows the four-direction order-7 screen. The finer order-9 z values appear in Table 5. The expanded catalogue produces $q = 0.32854$ and $M_{\parallel} = 42.490\,\mathrm{kA\,m^{-1}}$, versus $q = 0.05930$ and $7.669\,\mathrm{kA\,m^{-1}}$ for the original grid. Since both have $N = 28$ and $L = 105.22\,\mathrm{nm}$, the approximately 5.54-fold difference cannot arise from global loading. It is an architecture effect combining positions and body rotations, not a pure-angle measurement. Its magnitude is specific to this baseline, field, geometry and model; the stronger grid comparison below provides a more demanding reference.

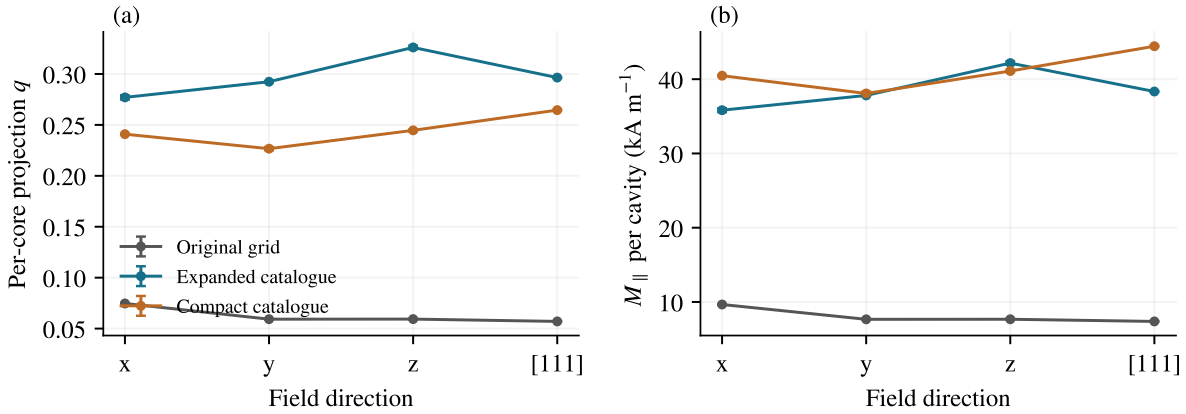

Figure 2: Historical order-7 directional screen at the nominal parameters. (a) Per-core moment projection; (b) response per cavity volume. The grid and expanded catalogue share the same cavity. Compact and expanded catalogue layouts share the source rotation pattern but differ in centre spacing and loading. Bars show approximate 95% Monte Carlo intervals for each mean, not model uncertainty. Directions are separate calculations on fixed bodies, not a hysteresis protocol.

Architecture	L (nm)	q	MCSE(q)	$M_{\parallel}$ (kA/m)
Original grid	105.220	0.05930	0.00019	7.669
Expanded catalogue	105.220	0.32854	0.00071	42.490
Compact catalogue	96.434	0.24279	0.00046	40.788
Tighter original grid	102.020	0.05313	0.00018	7.539
Fixed centers, source angles	107.102	0.28893	0.00034	35.432
Fixed centers, half aligned	107.102	0.26670	0.00032	32.706
Half-aligned LP derivative	97.676	0.24831	0.00023	40.144
Seven-column field grid	102.020	0.14423	0.00012	20.463
Fully aligned LP derivative	102.020	0.23242	0.00027	32.977

Table 5: Historical order-9 results at $T = 298\,\mathrm{K}$, $B = 10\,\mathrm{mT}$ along z, $M_s = 460\,\mathrm{kA\,m^{-1}}$ and $K_1 = -11.5\,\mathrm{kJ\,m^{-3}}$. Each row pools four runs. MCSE is sampling uncertainty for that row; uncertainty in the physical model and quadrature is not included.

4.3 Historical compression comparison

Compacting C_e to C_c increases number loading by a factor 1.29898. Nevertheless, the finer z-field response decreases from 42.490 to 40.788 kA m^{-1}, a 4.01% loss. The per-core projection decreases from 0.32854 to 0.24279, enough to offset the loading multiplier in Eq. (9). The order-7 screen gives a smaller, same-sign z loss (2.53%), while its [111] response increases by 15.9%. These historical comparisons show direction dependence and nonmonotonicity within the tested candidates. The fresh nominal estimate and the unresolved change in magnitude are assessed in Section 4.1; the historical 4.01% loss is not the current quantitative conclusion. Neither campaign identifies the full response landscape or proves that the expanded source is optimal.

4.4 Historical body-angle controls

In the strict common-centre family, the finer source-rotation response is $q = 0.28893$, compared with 0.26670 for half alignment. Thus the original mixed rotations yield 8.34% more z response with exactly the same particle positions and cavity. In the broader order-7 screen, half alignment reduces q by 7.37% along z and 8.23% along [111]; full alignment also reduces both responses (Fig. 3). This historical comparison tests whether orientations matter beyond density: they do in this registered-crystal, finite-cube model, but imposing a common body frame is unfavourable

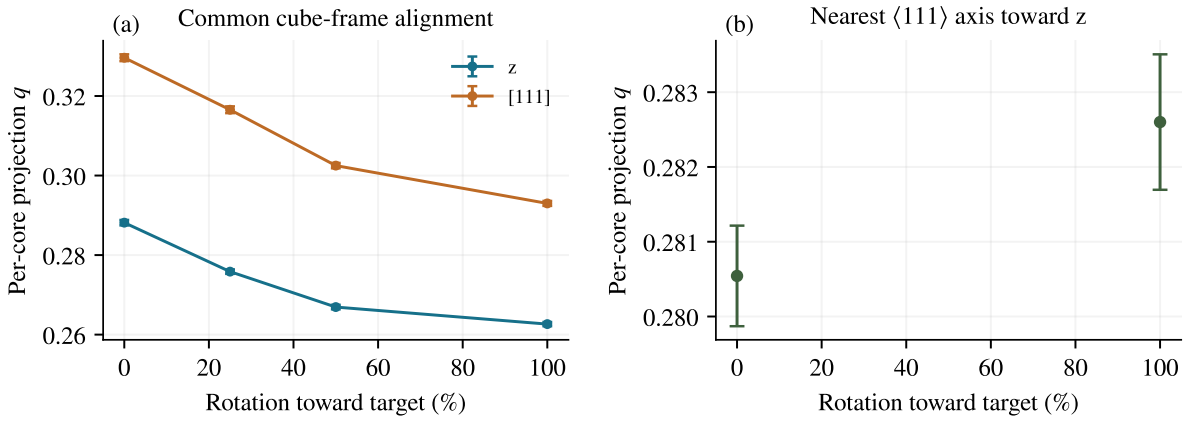

Figure 3: Historical order-7 strict rotation controls at identical centres and cavity within each family (order-7). Common body-frame alignment (a) and easy-axis alignment (b) use different common cavities and should not be compared as a single fixed-density family. The source and half-body-aligned z pair is independently repeated at order 9 in Table 5. No line is fitted to the two easy-axis points.

in these controls.

Complete rotation of the nearest $\langle 111 \rangle$ axis toward z produces only a 0.73% increase in its separate, larger-cavity order-7 strict control. That small effect is not used as the central positive result. It has not received the same order-9 confirmation as the principal comparisons and is small compared with unresolved physical-model uncertainty.

4.5 Historical orientation-adjusted candidates

The half-aligned, centre-adjusted candidate has side 97.6757 nm and finer $M_{\parallel} = 40.144\,\mathrm{kA\,m^{-1}}$, close to the compact source's $40.788\,\mathrm{kA\,m^{-1}}$ while halving mean deviation from the common frame. This is a candidate tradeoff, not an improvement over the source. The order-7 LP family also shows that the unaligned centre-adjusted control exceeds the half-aligned derivative in response per volume. Removing the extra pitch and moving centres must not be credited to alignment.

The cleaner historical positive comparison is fully aligned LP against G_z. Their cavity sides differ by approximately 10^{-6} nm, their counts and cores match, and every cube has the same physical body and cubic-anisotropy orientation up to cubic symmetry. The finer responses are 32.977 and $20.463\,\mathrm{kA\,m^{-1}}$, a ratio of 1.61152. This 61.15% increase therefore reflects the centre arrangement and its interactions, rather than extra loading or residual body tilts. It demonstrates that one source-derived centre pattern remains useful after the original orientations have been removed. It does not show superiority to every grid or to a magnetic-objective optimizer.

4.6 Historical parameter sensitivity and numerical refinement

The fully aligned LP advantage over G_z persists with $M_s = 250\,\mathrm{kA\,m^{-1}}$: its order-7 ratio is 1.591, compared with 1.622 nominally (Table 6). This is a useful parameter check, not proof of robustness over a calibrated uncertainty range. Other comparisons change more strongly: at the lower moment, the historical compact/expanded z response ratio is 1.192, so compression increases rather than decreases the estimated response. This exploratory lower-moment calculation was not repeated with the surface reference; it nevertheless prevents interpreting the fresh nominal −2.85% result as a general compression penalty. The $K_1 = 0$ ablation gives half/source strict-body ratios of 1.008 along z and 0.903 along [111]. Removing crystal anisotropy does not

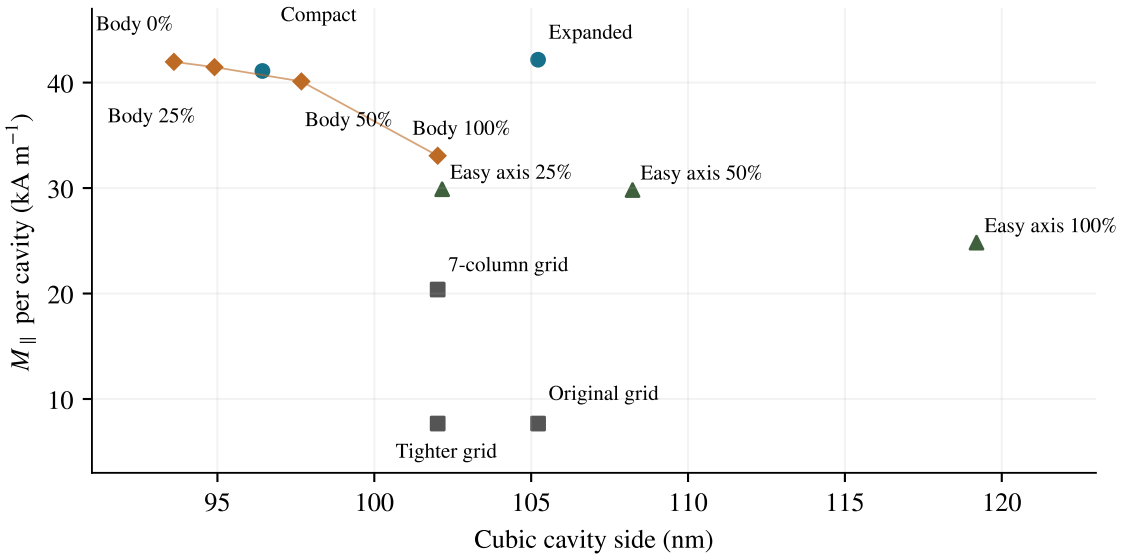

Figure 4: Historical screen: packing and finite-field response are different objectives. Order-7 z response is plotted against cavity side for the baselines, body-aligned LP family and easy-axis-aligned LP family. Smaller side means greater loading. The line connects constructed body-alignment variants only; it is not a continuous optimum frontier. Easy-axis rotations require larger cavities and do not supply a volumetric improvement over the unmodified catalogue at this operating point.

Order-7 response ratio	Nominal model	$M_\mathrm{s} = 250\,\mathrm{kA\,m^{-1}}$
Expanded / original grid, q_z	5.497	2.124
Compact / expanded, $M_\parallel^{(z)}$	0.975	1.192
Fully aligned LP / G_z, $M_\parallel^{(z)}$	1.622	1.591
Half / source body rotations, fixed centres, q_z	0.926	0.975

Table 6: Historical sensitivity ratios. These are point estimates, not independent-run confidence intervals. Changing M_s changes both the Zeeman scale and pair coupling; it is not a density intervention. Full per-condition MCSEs are retained in the accompanying CSV.

remove shape orientation from finite-core interaction tensors; it is not a universal no-orientation control.

All 64 configured conditions completed, comprising 256 run-level chains and 409,600 retained cold-replica samples. The largest recorded conventional split-$\widehat{R}$ across the gated projection and energy statistics is 1.00947, and the minimum ESS is approximately 504. The chosen diagnostic gates therefore pass. The independent post-hoc audit rehashed all 256 raw arrays and recomputed every projection and cavity normalization. Across all 128 condition-observable combinations, the larger of rank-normalized and folded split-$\widehat{R}$ remains below 1.01, with maximum 1.009684. This stronger check supports consistency of the retained chains without supplying a mixing-time or tail-ESS guarantee. Separate checks of the sampler agree with noninteracting quadrature and a Langevin limit, and local energy changes agree with full recomputation. Geometric validation is independent of magnetic sampling.

The order-9 repetitions preserve the headline directions and the large aligned-LP/grid difference, but not every order-7 value agrees within MCSE: selected projection differences reach about 3.9 combined MCSEs. The refined aligned/grid ratio is 1.6115, versus 1.6224 at order 7. These order-9 values are retained as historical checks; the present headline estimates use the independent surface-reference campaign in Section 4.1. This is evidence of numerical stability of the main comparison, not a rigorous bound on quadrature bias or proof that sampling errors

exhaust numerical uncertainty.

5 Discussion: what the catalogue contributes

The most defensible use of the packing catalogue is as a constrained generator of architectures. It contributes explicit poses, compact centre topology, contributor provenance and checkable feasibility. A physical objective can then decide whether to retain, rotate, expand or locally adjust those poses. The calculation provides two distinct reasons this may be useful. First, the original orientations affect response at fixed density and positions. Second, even when all those orientations are simplified to a common frame, a derived centre arrangement still outperforms a separately constructed directional grid in the declared model.

This is a bridge between discrete geometry and a materials calculation, not a claim that the geometrically best packing is the physically best material. Equation (9) explains why loading and per-core response must be assessed together. Magnetic interactions may favour collective arrangements that a density objective neither encodes nor rewards. Prior nanocube and superstructure studies already establish the importance of geometry and collective interactions [7, 30]. The controlled comparisons and independent numerical validation make the narrower catalogue-transfer argument testable. Their new contribution is the credited finite witness, geometrically checked adaptation family, explicit control hierarchy and reproducible numerical benchmark, not a new general inverse-design principle.

Equation (10) identifies the altered model term: centre-dependent pair interactions. An uncoupled control has identical per-core response algebraically, and the tiny loading mismatch slightly reduces the measured gain. The retained first-moment decomposition then locates the larger response in reduced opposing projection, while the common all-pair summaries describe the accompanying spin organization. The model-term identification is stronger than a visual correlation between packing and magnetization. Attribution to a particular geometric motif, bond family or reduction of frustration remains a hypothesis. The result concerns the tested directional grid and nominal finite-cluster model.

The compression conclusion is narrower. At the nominal parameters, increased loading fails to compensate for the lower projection: Eq. (11) gives a fresh gain of -2.85% with approximate 95% Monte Carlo interval $[-3.56, -2.13]\%$. The same-trajectory Q9 bridge preserves a similar decrease but leaves the historical magnitude unexplained. The current claim therefore concerns this fresh conditional comparison, not replication of a 4.01% loss. Its sign does not extend to the lower-moment case, which reverses the ordering.

The same distinction applies to manufacture. Partial alignment may be easier to target than a heterogeneous pose pattern, but ease of fabrication was not measured. A packing certificate only rules out stated exclusion violations; it does not predict energy minima, mechanical support, structural lifetime, attainable assembly yield or process cost. Earlier oriented-superlattice experiments show why process and crystallographic registration need separate treatment [31]. No claim about a practical magnetic device follows from the present finite-field number alone.

6 Limitations and next tests

The principal physical limitations are the full-cube rather than rounded core, assumed uniform magnetization, ideal body-to-crystal registration and uncalibrated magnetic parameters. Surface disorder, oxidation, polydispersity and ligand spacing may change both moments and interactions. Thermal rotation of whole particles, interparticle exchange and a surrounding matrix are absent. The finite open-boundary $N = 28$ cluster is not a periodic bulk composite. Cavity normalization is meaningful for comparing these proposed finite cells, but does not establish a

macroscopic constitutive law.

Geometric tolerances and magnetic uncertainty are separate. Exact checks and 1,240 independently replayed worst-size translation samples test exclusion geometry, with deliberate invalid controls rejected. The response calculation uses each nominal architecture once per condition; it does not sample the $\pm 5\%$ size and $\pm 0.1\,\mathrm{nm}$ position error space. A response-robustness claim would require that additional propagation. Rotational fabrication error is outside even the geometric guarantee.

The candidate family and operating point are restricted. Only one field magnitude and temperature were screened. The grids are explicit baselines, not optimized magnetic designs; the directional grid was added after initial results as a stronger comparison. The selected LP planes restrict reachable configurations. The campaign is exploratory, not a preregistered confirmatory study, and the reported large gains should not be interpreted as expected gains under equal-budget inverse design. Some future gains may be reduced by stronger baselines or micromagnetic corrections.

A stronger claim would require: calibrating core shape, active magnetic volume and anisotropy to one measured specimen; testing fields and temperatures; propagating fabrication variation; comparing macrospin and spatially resolved magnetic calculations; establishing an assembly or placement route; and benchmarking catalogue-seeded against direct magnetic optimization using the same evaluation budget. Any proposed customer or application metric should then be chosen separately. Equilibrium response alone does not establish heating efficiency, sensor noise, switching, memory retention or commercial value.

7 A falsifiable validation and application-transfer protocol

The exploratory screen and its subsequent numerical validation establish a computational case study. A confirmatory continuation should freeze its objective, architecture family, parameter range, baseline budget and acceptance threshold before producing new evaluation data. The following stages separate mathematical checking from material relevance.

1. *Geometry and data identity.* Retain exact pose and envelope certificates, source credits, immutable hashes and units. Add rotational uncertainty and realistic truncated or rounded-core geometry if a specimen requires them. Reject infeasible candidates before any property calculation.

2. *Numerical error.* The independent surface reference and fresh fixed-length comparisons complete a numerical cross-check for the selected nominal contrasts. A rigorous continuation would bound the remaining tensor error with interval or analytic remainder control and use Proposition 3 where informative. Freeze an allowable error relative to a future physically meaningful objective before testing new candidate families.

3. *Physical adequacy.* Fit M_{s}, anisotropy, active volume and crystal registration to one specified sample. Compare the uniform-core approximation with a resolved micromagnetic reference where its exchange length, size or surface structure warrants it. A reference solver should first pass a recognized micromagnetic benchmark, rather than only agree with this code [32]. Propagate a measured joint fabrication distribution through geometry and response. Neither exact geometry nor more Monte Carlo samples can remove model discrepancy.

4. *Fair design comparison.* Compare catalogue seeding, random feasible seeds, grids and direct property optimization under the same constraints and total physics-evaluation budget. Use a held-out parameter set for final ranking. Report the response gain, feasibility yield, computational cost and uncertainty. This tests usefulness of a catalogue as a search prior, rather than merely finding one favourable comparison.

5. *Material and device evidence.* Specify an assembly or placement route, an experimental observable, a reference sample and a pass/fail threshold. Measure the field or angular dependence needed by the application. A material-level improvement requires data beyond the present equilibrium finite-cluster point.

7.1 Transfer to other materials applications

The reusable object is the architecture-to-constraint-to-property workflow. It can propose shape and orientation families, certify stated exclusion constraints, and organize a fair downstream solver comparison. Each application requires its own material phases, operating conditions, dimensional scale, manufacturing constraints and objective. Depending on that specification, future calculations could concern thermal or elastic tensors, transport through porous structures, or admissible microstructure under stress. These are prospective models, not results of this study. The magnetite macrospin calculation does not validate plasma confinement, radiation resistance or fusion-component performance.

8 Conclusion

Within an explicitly idealized equilibrium model, a polyhedral-packing catalogue supplies useful magnetic-architecture candidates beyond the benefit of increased loading. The strict controls identify joint body/crystal-frame orientation effects, while a fully aligned source derivative retains a 60.94% response advantage over the tested matched-density, matched-orientation directional grid. The pointwise Hamiltonian contrast identifies changed pair interactions, and retained site means decompose the gain into opposing and positive contributions without assigning a unique microscopic pathway. At the stated nominal parameters, compression has a fresh finite-run gain estimate of -2.85%, with approximate 95% Monte Carlo interval $[-3.56, -2.13]\%$. Reweighting the fresh trajectories to the historical Q9 tensors leaves the old-to-fresh magnitude discrepancy unresolved, and an exploratory lower-moment case reverses the ordering. No general compression penalty or universal benefit of stronger alignment follows. These findings support a workflow in which certified geometric configurations seed subsequent physical design. The present evidence is computational and finite; physical realization and calibrated performance remain open.

Data, code, attribution and preparation

The original version 1.0 computational record, including `materials-supplement-v1.0.zip` and `materials-latex-source-v1.0.zip`, remains available at `https://doi.org/10.5281/zenodo.23629684`. Version 1.1 revises interpretation and declarations and adds reproducible post-hoc summaries of retained trajectories. It changes no raw trajectory, headline estimate, existing figure or numerical table, geometric certificate or proof artifact. Editable sources are supplied as `materials-joatp-latex-source-v1.1.zip`; the additional analysis scripts, derived summaries and provenance are supplied separately as `materials-revision-audit-v1.1.zip`. The supplement contains the frozen witness and attribution, derived geometries, rational feasibility checks, historical and reference interaction tensors, 304 retained run outputs, analysis scripts and checks, the bounded Lean development, environment specifications and file manifests. It distinguishes the 64 historical conditions from the six new reference-integral conditions. The original version 1.0 manuscript-source archive contains the earlier manuscript, bibliography, vector figures, tables and compact numerical CSV files. Reproduction instructions distinguish analysis of saved data from optional fresh simulations. Appendix D describes the retained record.

The author's earlier packing work is cited through its deposited DOI records [14–16]. The specific community witness and method lineage retain their separate attribution [17, 18]. No contributor executable was used to run the magnetic screen. Code, numerical analysis, literature retrieval and manuscript preparation used AI assistance under the author's scientific direction. The author reviewed the manuscript and is responsible for its scientific interpretation. Separate code implementations and computational audits support numerical checks; they are not external peer review or experimental validation. This version retains the historical results and geometry-verification layer and includes the independent interaction reference, 48 additional runs and ratio-level Monte Carlo uncertainty. No experimental data are reported.

Funding

This study was supported by the author's personal resources. It received no external grant or client funding.

Competing interests

The author is a cofounder of Superradiant, whose commercial research interests include mathematical verification and materials design and overlap the subject of this paper. This professional interest is disclosed separately from the study's personal funding. No other relevant competing interests are declared.

The manuscript and study data are provided under the Creative Commons Attribution 4.0 International licence. New study software and Lean sources are provided under the MIT licence. Applicable upstream notices and contributor credits are retained separately in the supplement. These statements do not relicense third-party dependencies.

A Support constraints and the restricted centre LP

Let Q_i be a cube with centre $\mathbf{c}_i$, rotation R_i and side a. For a direction $\mathbf{n}$, its centred support radius is

$$h_i(\mathbf{n}; a) = \frac{a}{2} \sum_{k=1}^{3} |\mathbf{n} \cdot R_i \mathbf{e}_k|. \tag{12}$$

With independent centre error $\boldsymbol{\delta}_i \in [-\epsilon, \epsilon]^3$, wall containment in a centred side-L cavity is guaranteed by

$$\pm c_{i\ell} + h_i(\mathbf{e}_\ell; a_{\max}) + \epsilon \le L/2 \quad (\ell = x, y, z). \tag{13}$$

For a directed pair normal $\mathbf{n}_{ij}$, robust separation follows from

$$\mathbf{n}_{ij} \cdot (\mathbf{c}_j - \mathbf{c}_i) \ge h_i(\mathbf{n}_{ij}; a_{\max}) + h_j(\mathbf{n}_{ij}; a_{\max}) + 2\epsilon \|\mathbf{n}_{ij}\|_1. \tag{14}$$

The pair normal is fixed during a particular LP, so the inequalities are linear in centres and L. Their feasibility is sufficient for nonoverlap, but the retained normals restrict the search and cannot establish a global packing optimum. The second-stage ℓ_1 objective uses standard positive auxiliary variables for coordinate displacement while allowing at most 10^{-8} relative increase over the first-stage side. Centres are rounded to 12 decimal places in nm, followed by uniform inflation by 1.000001 and exact rechecking. Body interpolation uses shortest rotation-vector scaling after minimization over the 24 proper cubic symmetries.

Quaternions stored as binary64 values are interpreted as exact rationals and normalized algebraically to define rational orthogonal rotation matrices. Final wall and pair checks use these exact objects and a small protective inflation of rounded LP solutions. There are 1,344 vertex-wall inequalities and 378 pair-plane checks for each design. A separate floating-point

separating-axis replay recomputes candidate axes and tests bounded translations, rather than trusting optimizer status. Appendix C specifies the narrower geometry lemmas and two finite-layout certificates checked by Lean in this edition; this does not cover every adapted layout or the magnetic model.

B Sampler conditional and diagnostics

Writing $E/(k_{\mathrm{B}}T)$ as $\kappa \sum_i f(R_i^{\mathsf{T}}\mathbf{s}_i) - \sum_i \mathbf{b}\cdot\mathbf{s}_i + \sum_{i<j}\mathbf{s}_i^{\mathsf{T}} D_{ij}\mathbf{s}_j$, define

$$\kappa = \frac{K_1 V_c}{k_{\mathrm{B}}T}, \quad \mathbf{b} = \frac{m\mathbf{B}}{k_{\mathrm{B}}T}, \quad D_{ij} = \frac{\mu_0 m^2}{4\pi k_{\mathrm{B}}T}\mathcal{T}_{ij}. \tag{15}$$

For a single spin, the linear conditional field is $\mathbf{h}_i = \mathbf{b} - \sum_j D_{ij}\mathbf{s}_j$. At ladder factor β, a proposal proportional to $\exp(\beta\mathbf{h}_i\cdot\mathbf{s}_i')$ cancels the linear energy terms in the Metropolis–Hastings ratio. The remaining acceptance probability is $\min\{1, \exp[-\beta\kappa(f_i' - f_i)]\}$. The reflection proposal is involutive and preserves $\mathbf{h}_i\cdot\mathbf{s}_i$, hence uses the same anisotropy correction. Replica swaps use $\min\{1, \exp[(\beta_r - \beta_{r+1})(e_r - e_{r+1})]\}$ with reference-temperature dimensionless energies.

Proposition 4 (Invariant target of the idealized update)**.** *At a fixed positive ladder factor β, exact von Mises–Fisher linear-conditional proposals with the stated anisotropy Metropolis–Hastings correction preserve $\pi_\beta \propto e^{-\beta U}$. A measure-preserving, involutive spin reflection preserving the linear conditional, corrected by the same anisotropy ratio, also preserves π_β. Composing such single-spin kernels and the stated accepted replica swaps preserves the product target across replicas. This specializes classical acceptance-ratio and exchange methods [25, 33, 34].*

Proof. For the independent linear-conditional proposal density $g(s) \propto e^{\beta\mathbf{h}_i\cdot s}$, the proposal-density ratio cancels the linear-energy contribution to the target ratio. The remaining ratio is exactly the anisotropy factor displayed above, proving detailed balance for that single-spin kernel. An orthogonal involution preserves sphere measure and is its own reverse proposal, so its linear-energy-preserving move has the same remaining target ratio. The swap acceptance is the ratio of the two exchanged product densities. Each kernel therefore preserves its target, as do deterministic compositions and random mixtures of the kernels. □

For $\mathbf{h}_i \neq 0$, the implementation's reflection is $s \mapsto 2(\mathbf{h}_i\cdot s)\mathbf{h}_i/\|\mathbf{h}_i\|^2 - s$, an orthogonal involution. A zero-field branch must use an independently valid update; at zero field the uniform-sphere proposal is the exact linear conditional. The retained implementation also uses that proposal when $\beta\|\mathbf{h}_i\| < 10^{-8}$, a small-field approximation rather than an exact conditional at nonzero field. Proposition 4 is an algebraic statement conditional on exact updates. Floating-point arithmetic, that approximation, finite burn-in and possible slow collective modes require the separate implementation checks and diagnostics. A composition can preserve the target without itself being reversible; no finite-time error or mixing-time bound is asserted.

The sampler is checked against direct energy differences and independent noninteracting spin quadrature. These checks test algebra and implementation on controlled cases; they do not eliminate possible slow collective modes of interacting architectures. The historical gate uses ordinary split-chain variances. The separate audit of historical chains follows rank normalization and folding [27]; it is not a retroactive preregistration. The new reference campaign also includes tail and local diagnostics, retained separately from the historical audit. Raw projection and energy series, start means, swap statistics and round trips remain available for a stronger downstream audit.

C Verification inventory and trust boundaries

The Lean development contains 15 generic theorem endpoints for three-dimensional projection algebra, cube and error-box support, relative coordinate error, robust separation, strict-halfspace disjointness, wall containment, scale and translation transport, and loading-response algebra. A second module checks 54 finite fact blocks containing 3,444 rational inequalities for two layouts: the original simple grid and the compact catalogue, 1,722 inequalities each. These are 2,688 wall inequalities and 756 pair inequalities in total. Both modules compile with Lean 4.28.0 and the pinned Mathlib revision. The 69-endpoint axiom audit reports only standard foundational axioms, with no admitted proofs or native-decision proof shortcut. The certificate layer is a new verification artifact, not a new separating-axis theorem.

The finite source interface remains explicit: an independent Python implementation reconstructs rational geometry from the raw quaternion and pose data and compares all 756 saved pair residuals. Lean checks the emitted rational nonnegativity facts and their counts. It does not parse the original JSON, verify file hashes, or prove that Python's emitted residuals equal the intended geometry formulas. The generic support lemmas and the finite rational facts therefore have separate evidence paths, joined by a checked computational bridge rather than a fully formalized input pipeline. Compiler versions, dependency pins, theorem names, source hashes and axiom output are retained in the supplement's verification records.

Layer	Checked evidence	Remaining boundary
Geometry algebra	Lean kernel proofs of 15 endpoints	No encoded physical-shape assertion
Two finite certificates	Lean nonnegativity of 3,444 rational inequalities	Python JSON-to-residual derivation; not all adaptations
Other geometry variants	Exact Python checks and separate numeric replay	No Lean replay of all 31 prepared geometries
Gibbs and sampler propositions	Analytic proofs in this paper	No Lean formalization of integration or kernels
Magnetic response	Historical trajectories plus fresh reference-integral runs and contrast intervals	Non-certified integration remainder, finite-run and model discrepancy
Material or fusion use	Explicit future validation protocol	No experiment or fusion-property calculation

Table 7: Verification is a layered claim. The presence of Lean proofs does not make the full magnetic study formally verified.

D Run inventory and reproducibility

The frozen witness is the file

```
records/community/cubincub/
cubincub_n28_pingyou_20261013.json
```

at the commit cited in Section 2. Its SHA-256 is

```
2d9a692351ee680ed0a5397b0b90172f77e70b2af18724e8abe083ab5b7fb319
```

The original version 1.0 source bundle's historical `data/ALL-CONDITIONS.csv` records every condition's geometry, direction, parameters, quadrature model, cavity side, projection, MCSE, cavity-normalized response and diagnostics. `data/selected-order9.csv` supplies Table 5. The

Panel	Conditions	Run-level chains
Main directional / adaptation screen	30	120
Lower-moment and anisotropy sensitivity	16	64
Easy-axis and directional-grid extension	7	28
Fresh-seed order-9 confirmation	9	36
Lower-moment aligned/grid confirmation	2	8
Historical total	64	256
New reference-integral validation	6	48
Combined retained record	70	304

Table 8: All historical and new configured conditions completed; no condition was omitted for an unfavourable mean or failed diagnostic. The 31 prepared geometries include unrun controls; 18 distinct architecture names were simulated. Pilot/self-test runs are excluded from these totals.

accompanying full computational package retains input files, content-addressed tensor caches, the sampler, geometry scripts and all raw arrays. The new numerical supplement additionally records the frozen 48-run protocol, six reference conditions, three contrast intervals, local diagnostics and interaction-refinement errors. Paired-Hamiltonian reweighting estimates integration sensitivity from retained cold spins, with weight diagnostics and approximate paired Monte Carlo uncertainty. Package manifests provide file-level hashes; they are integrity records, not scientific validation certificates.

The intended reproduction environment is Python 3.12 with NumPy 2.3.5, Numba 0.68.0, llvmlite 0.50.0, SciPy 1.18.1 and Matplotlib 3.11.2. The package relocation helper maps retained input paths to an extracted archive and directs fresh outputs to a separate reproduction folder, preserving original chains. Numerical runs do not execute catalogue-contributor code. Recompiling this manuscript requires standard LaTeX packages, BibTeX and the supplied vector figures; the release build also records a clean source-archive compilation check.

References

[1] Salvatore Torquato and Yang Jiao. Dense packings of polyhedra: Platonic and Archimedean solids. *Physical Review E*, 80:041104, 2009. doi: 10.1103/PhysRevE.80.041104. URL https://journals.aps.org/pre/abstract/10.1103/PhysRevE.80.041104. Corrected article; errata published in 2010.

[2] Pablo F. Damasceno, Michael Engel, and Sharon C. Glotzer. Predictive self-assembly of polyhedra into complex structures. *Science*, 337(6093):453–457, 2012. doi: 10.1126/science.1220869. URL https://doi.org/10.1126/science.1220869.

[3] Joel Henzie, Michael Grünwald, Asaph Widmer-Cooper, Phillip L. Geissler, and Peidong Yang. Self-assembly of uniform polyhedral silver nanocrystals into densest packings and exotic superlattices. *Nature Materials*, 11:131–137, 2012. doi: 10.1038/nmat3178. URL https://www.nature.com/articles/nmat3178.

[4] Da Wang, Michiel Hermes, Ramakrishna Kotni, Yaoting Wu, Nikos Tasios, Yang Liu, Bart de Nijs, Ernest B. van der Wee, Christopher B. Murray, Marjolein Dijkstra, and Alfons van Blaaderen. Interplay between spherical confinement and particle shape on the self-assembly of rounded cubes. *Nature Communications*, 9:2228, 2018. doi: 10.1038/s41467-018-04644-4. URL https://www.nature.com/articles/s41467-018-04644-4.

[5] Gurvinder Singh, Henry Chan, T. Udayabhaskararao, Elijah Gelman, Davide Peddis, Artem Baskin, Gregory Leitus, Petr Král, and Rafal Klajn. Magnetic field-induced self-assembly of iron oxide nanocubes. *Faraday Discussions*, 181:403–421, 2015. doi: 10.1039/C4FD00265B. URL https://pubs.rsc.org/en/content/articlepdf/2015/fd/c4fd00265b.

[6] Joe G. Donaldson and Sofia S. Kantorovich. Directional self-assembly of permanently magnetised

nanocubes in quasi two dimensional layers. *Nanoscale*, 7:3217–3228, 2015. doi: 10.1039/C4NR07101H. URL https://pubs.rsc.org/en/content/getauthorversionpdf/C4NR07101H.

[7] Igor Stanković, Luis Lizardi, and Carlos García. Assembly of nanocube super-structures directed by surface and magnetic interactions. *Nanoscale*, 12:19390–19403, 2020. doi: 10.1039/D0NR03485A. URL https://www.scl.rs/papers/2020_Stankovic_NanoscaleNanocubes.pdf.

[8] Erik Wetterskog, Christian Jonasson, Detlef-M. Smilgies, Vincent Schaller, Christer Johansson, and Peter Svedlindh. Colossal anisotropy of the dynamic magnetic susceptibility in low-dimensional nanocube assemblies. *ACS Nano*, 12(2):1403–1412, 2018. doi: 10.1021/acsnano.7b07745. URL https://pubs.acs.org/doi/10.1021/acsnano.7b07745.

[9] Verner Håkonsen, Gurvinder Singh, José A. De Toro, Peter S. Normile, Erik Wahlström, Jianying He, and Zhiliang Zhang. Reconfigurable mechanical anisotropy in self-assembled magnetic superstructures. *Advanced Science*, 8(8):2002683, 2021. doi: 10.1002/advs.202002683. URL https://advanced.onlinelibrary.wiley.com/doi/full/10.1002/advs.202002683.

[10] N. A. Usov and O. N. Serebryakova. Equilibrium properties of assembly of interacting superparamagnetic nanoparticles. *Scientific Reports*, 10:13677, 2020. doi: 10.1038/s41598-020-70711-w. URL https://arxiv.org/abs/2003.13355.

[11] Marian Florescu, Salvatore Torquato, and Paul J. Steinhardt. Designer disordered materials with large, complete photonic band gaps. *Proceedings of the National Academy of Sciences*, 106(49): 20658–20663, 2009. doi: 10.1073/pnas.0907744106. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC2777962/.

[12] Rose K. Cersonsky, Julia Dshemuchadse, James Antonaglia, Greg van Anders, and Sharon C. Glotzer. Pressure-tunable photonic band gaps in an entropic colloidal crystal. *Physical Review Materials*, 2: 125201, 2018. doi: 10.1103/PhysRevMaterials.2.125201. URL https://arxiv.org/abs/1803.06388.

[13] Yih Hong Lee, Chee Leng Lay, Wenxiong Shi, Hiang Kwee Lee, Yijie Yang, Shuzhou Li, and Xing Yi Ling. Creating two self-assembly micro-environments to achieve supercrystals with dual structures using polyhedral nanoparticles. *Nature Communications*, 9:2769, 2018. doi: 10.1038/s41467-018-05102-x. URL https://www.nature.com/articles/s41467-018-05102-x.

[14] Alejandro Zarzuelo Urdiales. Certified finite packings of Platonic solids: a generalized soft-to-rigid search and its computational development, 2026. URL https://doi.org/10.5281/zenodo.23306098. Zenodo preprint, version 1.0, 11 October 2026.

[15] Alejandro Zarzuelo Urdiales. Finite polyhedral fits: discrete scale jumps, orientational domains, and contact structure in 351 certified constructions, 2026. URL https://doi.org/10.5281/zenodo.23306126. Zenodo preprint, version 1.0, 11 October 2026.

[16] Alejandro Zarzuelo Urdiales. Polyhedra in polyhedra packing: 351 certified constructions, structural measurements, software, and visualizations, 2026. URL https://doi.org/10.5281/zenodo.23306130. Zenodo dataset, version 1.0, publication date 12 October 2026; original nine-family archive.

[17] konoha-ily2 and PingYou Ltd. Certified 28-cube-in-cube catalogue witness: original configuration and wall-padding refinement, 2026. URL https://github.com/alejandrozu/platonic-packing/blob/1abc63218db34cea95802aeb3b6d4a6b6fe1c61b/records/community/cubincub/cubincub_n28_pingyou_20261013.json. Original geometry and certificate by konoha-ily2; unchanged-pose wall-padding refinement by PingYou Ltd. Frozen source commit 1abc63218db34cea95802aeb3b6d4a6b6fe1c61b.

[18] Yohei Nakajima. Soft-to-rigid: Shape continuation as an alternative search for packing congruent shapes, 2026. URL https://arxiv.org/abs/2610.14463v1. Version 1, 10 October 2026.

[19] Stefan Gottschalk, Ming C. Lin, and Dinesh Manocha. OBBTree: A hierarchical structure for rapid interference detection. In *Proceedings of SIGGRAPH '96*, pages 171–180, 1996. doi: 10.1145/237170.237244. URL https://www.cs.cornell.edu/courses/cs667/2005sp/readings/gottschalk96.pdf.

[20] Roberto Moreno, Samuel Poyser, Daniel Meilak, Andrea Meo, Sarah Jenkins, Vlado K. Lazarov, Gonzalo Vallejo-Fernandez, Sara Majetich, and Richard F. L. Evans. The role of faceting and elongation on the magnetic anisotropy of magnetite Fe3O4 nanocrystals. *Scientific Reports*, 10:2722, 2020. doi: 10.1038/s41598-020-58976-7. URL https://www.nature.com/articles/s41598-020-58976-7.

[21] Andrew J. Newell, Wyn Williams, and David J. Dunlop. A generalization of the demagnetizing tensor for nonuniform magnetization. *Journal of Geophysical Research: Solid Earth*, 98(B6):9551–9555, 1993. doi: 10.1029/93JB00694. URL https://agupubs.onlinelibrary.wiley.com/doi/ab

s/10.1029/93JB00694.

[22] Dmitri Chernyshenko and Hans Fangohr. Computation of the magnetostatic interaction between linearly magnetized polyhedrons. *Journal of Magnetism and Magnetic Materials*, 412:132–137, 2016. doi: 10.1016/j.jmmm.2016.03.083. URL https://arxiv.org/abs/1601.04185.

[23] James L. G. O'Connell, William S. P. Robertson, and Benjamin S. Cazzolato. Analytic magnetic fields and semi-analytic forces and torques due to general polyhedral permanent magnets. *IEEE Transactions on Magnetics*, 56(1):7400208, 2020. doi: 10.1109/TMAG.2019.2942538. URL https://doi.org/10.1109/TMAG.2019.2942538.

[24] Dmitri Chernyshenko and Hans Fangohr. Computing the demagnetizing tensor for finite difference micromagnetic simulations via numerical integration. *Journal of Magnetism and Magnetic Materials*, 381:440–445, 2015. doi: 10.1016/j.jmmm.2015.01.013. URL https://arxiv.org/abs/1403.1978.

[25] Koji Hukushima and Koji Nemoto. Exchange Monte Carlo method and application to spin glass simulations. *Journal of the Physical Society of Japan*, 65(6):1604–1608, 1996. doi: 10.1143/JPSJ.65.1604. URL https://doi.org/10.1143/JPSJ.65.1604.

[26] Andrew Gelman and Donald B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992. doi: 10.1214/ss/1177011136. URL https://doi.org/10.1214/ss/1177011136.

[27] Aki Vehtari, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. Rank-normalization, folding, and localization: An improved $\widehat{R}$ for assessing convergence of MCMC. *Bayesian Analysis*, 16(2):667–718, 2021. doi: 10.1214/20-BA1221. URL https://arxiv.org/abs/1903.08008.

[28] James M. Flegal and Galin L. Jones. Batch means and spectral variance estimators in Markov chain Monte Carlo. *The Annals of Statistics*, 38(2):1034–1070, 2010. doi: 10.1214/09-AOS735. URL https://arxiv.org/abs/0811.1729.

[29] E. C. Fieller. Some problems in interval estimation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 16(2):175–185, 1954. doi: 10.1111/j.2517-6161.1954.tb00159.x. URL https://doi.org/10.1111/j.2517-6161.1954.tb00159.x.

[30] Jeotikanta Mohapatra, Pramanand Joshi, Hur Abbas, Markus Gusenbauer, Kaifu Bian, Ping Lu, Hongyou Fan, Thomas Schrefl, and J. Ping Liu. Superstructure magnetic anisotropy in Fe3O4 nanoparticle chains. *Nature Communications*, 16:5723, 2025. doi: 10.1038/s41467-025-60888-x. URL https://www.nature.com/articles/s41467-025-60888-x.

[31] Anwar Ahniyaz, Yasuhiro Sakamoto, and Lennart Bergström. Magnetic field-induced assembly of oriented superlattices from maghemite nanocubes. *Proceedings of the National Academy of Sciences*, 104(45):17570–17574, 2007. doi: 10.1073/pnas.0704210104. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC2077042/.

[32] National Institute of Standards and Technology. muMAG standard problem #2, 1998. URL https://www.ctcms.nist.gov/mumag/std2/spec2.html. Micromagnetic verification benchmark; exchange-length scaling specification, accessed 25 October 2026.

[33] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. doi: 10.1063/1.1699114. URL https://doi.org/10.1063/1.1699114.

[34] W. K. Hastings. Monte carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970. doi: 10.1093/biomet/57.1.97. URL https://doi.org/10.1093/biomet/57.1.97.