

Boundary sensitivity in finite-sized artificial spin ice explored via AI-assisted genetic algorithms

Received: 19 November 2025

Accepted: 11 February 2026

Cite this article as: Moon, T.J., Park, S.M., Yoon, H.G. *et al.* Boundary sensitivity in finite-sized artificial spin ice explored via AI-assisted genetic algorithms. *npj Comput Mater* (2026). <https://doi.org/10.1038/s41524-026-02016-x>

Tae Jung Moon, Seong Min Park, Han Gyu Yoon, Hee Young Kwon & Changyeon Won

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Boundary sensitivity in finite-sized artificial spin ice explored via AI-assisted genetic algorithms

Tae Jung Moon¹, Seong Min Park¹, Han Gyu Yoon¹, Hee Young Kwon^{2*}, and Changyeon Won^{1*}

¹*Department of Physics, Kyung Hee University, Seoul 02447, South Korea*

²*Center for Spintronics, Korea Institute of Science and Technology, Seoul 02792, South Korea*

*: Corresponding authors

E-mail: soky572@gmail.com, cywon@khu.ac.kr

ABSTRACT

Frustrated magnetic systems such as spin ice are key platforms for novel metamaterials. However, identifying their ground states in finite arrays is a formidable challenge, as boundary sensitivity and metastable states trap conventional optimization methods. We introduce a virtuous-cycle AI pipeline where a genetic algorithm explores the latent space of a variational autoencoder (VAE), with the best candidates progressively refining the VAE's representation. Applied to Kagome spin ice, this method reveals how the boundary magnetism is determined: boundaries break the symmetry of the $\sqrt{3} \times \sqrt{3}$ magnetic superstructure while the bulk superstructure order in the interior maintains. Furthermore, it demonstrates that high geometric confinement induces a novel quasi-ferromagnetic phase, which breaks the interior superstructure order. Our work provides a predictive framework for designing frustrated materials and demonstrates a powerful AI approach for boundary-sensitive physical systems.

Introduction

Finding the ground state in interacting many-body systems is one of the most challenging problems across various disciplines¹⁻⁵. Due to the large configuration space and rugged energy landscape, exhaustive enumeration is infeasible, and standard heuristics tend to converge to local minima rather than the global minimum. Artificial spin ice, composed of interacting nanomagnets arranged on lattices designed to realize characteristic physical properties such as geometrical frustration, residual entropy, and emergent excitations⁶⁻¹², serves as a representative platform in which searching for the ground state is particularly challenging.

Artificial Kagome spin ice is a canonical example, built on a Kagome pattern including corner-sharing triangles¹³. The nanomagnets in the system are modeled as magnetic dipoles placed on the edges of a honeycomb lattice, coupled by magnetic dipole-dipole interactions, confined to the axis along the edges, and allowed to flip only between two orientations. It is well known that each three-fold vertex on an artificial Kagome spin ice system tends to satisfy the "ice rule" (two-in one-out or one-in two-out)^{7,13} in order to minimize the local dipolar energy, leading to a macroscopically degenerate low-energy manifold¹³⁻¹⁶. In fact, since dipole-dipole interactions are long-ranged and the contributions from the magnetic dipoles beyond nearest neighbors lift this degeneracy, the true ground state is not determined solely by the ice rule. Assuming an infinite system with proper symmetries, theoretical and numerical studies predict the ground state of the artificial Kagome spin ice system, in which the magnetic dipoles and local magnetic charges are long-range ordered simultaneously¹⁶.

However, in the case of finite systems, the effects of long-range interactions depend on the shape of the system and its boundaries, which makes intuitive prediction of the ground state challenging. Moreover, due to the vast configuration space and numerous local energy minima, the difficulty of searching for the ground state increases exponentially as the system size increases.

In response to these challenges, recent machine learning approaches address the limitations of conventional optimization methods by learning the underlying correlations and constraints directly from data, thereby assisting the optimization toward physically plausible regions of configuration space^{2,17-24}. In particular, the variational autoencoder (VAE), a representative deep generative model based on artificial neural networks, has been widely utilized to capture the properties of many-body systems and to accelerate the search for low-energy configurations²⁵.

The VAE extracts essential information from high-dimensional data by compressing and reconstructing it via a compact latent-space representation^{26–28}. Thus, optimization can be performed effectively in this latent space, reducing the computational cost of searching for desired states. It is known that the genetic algorithm (GA), a conventional evolutionary optimization algorithm, performs better in latent-space than in the original data space for global minimum searches^{22,29–32}.

Building on this synergy, hybrid frameworks combining VAEs and GAs have emerged across diverse fields, such as latent-space representation for failure analysis³³, token optimization for image reconstruction³⁴, and anomaly detection in high-dimensional data³⁵. While these models demonstrate the utility of latent-space exploration, they typically decouple representation learning from evolutionary search and are often aimed at task-specific, data-driven objectives^{33–35}. In contrast, we propose a virtuous-cycle VAE–GA pipeline tailored to navigate the rugged energy landscapes of frustrated physical systems, integrating representation learning and global exploration in an adaptive closed loop. Specifically, a VAE model is first trained using an initial spin configuration dataset secured by simulated annealing based on Monte Carlo³⁶ to learn a compact latent space representation. The GA then performs selection, crossover and mutation processes directly within the latent space to propose new candidates that are energetically stabilized²². These candidates are incorporated into the dataset, and the VAE model is retrained to guide on the relevant data manifold^{37,38}. Through this iterative VAE-GA cycle, the latent representation progressively concentrates on the low-energy manifold over successive iterations^{29,37,38}. We apply this approach to artificial Kagome spin ice with open boundaries to systematically investigate the interplay between the global array geometry and local boundary terminations. Our results reveal the governing principles of ground state selection, providing a predictive framework for engineering target magnetic states in frustrated nanomagnet arrays.

Results and Discussion

VAE-GA framework for ground state search

We introduce and apply our virtuous-cycle VAE-GA optimization pipeline, a powerful and general framework for discovering low-energy states in frustrated systems. The pipeline integrates a VAE with a GA to navigate vast and complicated energy landscapes obtained by Monte Carlo simulated annealing. To demonstrate its efficacy, we apply this method to the canonical challenge of finding the ground state for geometrically frustrated Kagome spin ice under

open boundary conditions.

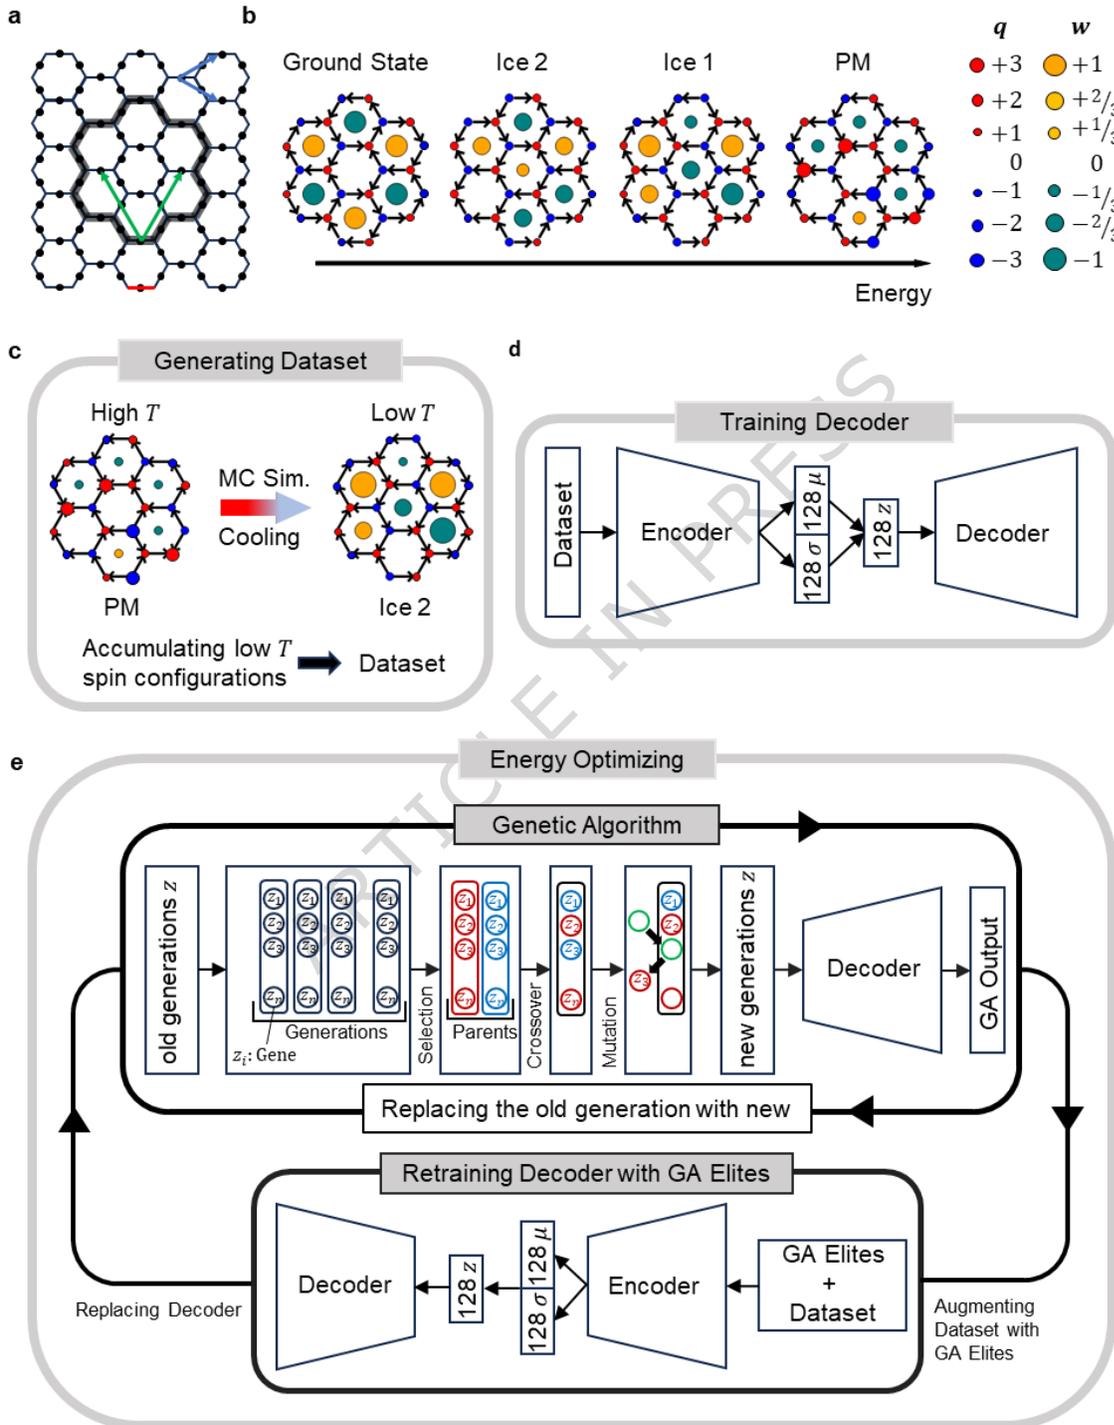


Figure 1 Schematic of the virtuous-cycle VAE-GA pipeline.

(a) The Kagome lattice, formed by placing spins at the midpoints of the edges of each hexagon in a honeycomb

structure. The blue and green arrows indicate the unit vectors of the Kagome lattice and $\sqrt{3} \times \sqrt{3}$ magnetic superstructure respectively. The size of the red bar indicating the side length of a hexagon is $\frac{1}{\sqrt{3}}$ (b) Thermodynamic phases of Kagome spin ice from the paramagnetic state to the ground-state phase, characterized by vertex charge q and loop winding number w . Positive charges are shown in red and negative charges in blue. The size of the disks at the vertices and hexagon centers is proportional to the absolute values of the vertex charge, $|q|$, and the loop winding number, $|w|$, respectively. (c) Generating procedure of the initial low-energy dataset via Monte Carlo simulated annealing. (d) Training procedure of the VAE's decoder, which learns the crucial generative mapping from the latent space z back to the physical spin configuration space. (e) The core optimization loop, where the GA discovers lower-energy states in latent space, which are then used to retrain and improve the VAE.

The physical system under investigation is the Kagome spin ice, which is constructed on a honeycomb lattice composed of corner-sharing hexagons (Fig. 1a). Spin sites (black dots) are located at the midpoint of each hexagon edge, forming the Kagome lattice. The length scale is defined by the hexagonal side length of $\frac{1}{\sqrt{3}}$ (red line), which normalizes the structural unit vectors (blue arrows) to unit length. On this lattice, the $\sqrt{3} \times \sqrt{3}$ ground state^{39,40} forms a magnetic superstructure. The highlighted region indicates the bulk interior, whose thermodynamic evolution from a high-temperature regime to the ground state is detailed in Fig. 1b. It shows the system passing through four characteristic stages, which are characterized by two key parameters. The vertex charge q is defined as the number of spins pointing inward minus those pointing outward at a vertex. The loop winding number w measures the net circulation of the six spins surrounding a hexagonal array. It is normalized to range from $+1$ (all clockwise, yellow) to -1 (all counterclockwise, teal), with fractional values indicating mixed chirality.

The progression through these phases reveals a gradual ordering process. In the paramagnetic phase, q and w are spatially not fully correlated, and all q values from $+3$ to -3 occur with similar probability. In the Ice 1 phase, the ice rule is satisfied across the lattice ($|q| = 1$), while w remains disordered due to the high degeneracy of circulation patterns. Cooling further into the Ice 2 phase produces an alternating arrangement of $q = +1$ and $q = -1$ vertices, with w forming finite chirality domains separated by domain walls. Finally, in the ground state phase, both q and w achieve perfect order, and the array realizes the $\sqrt{3} \times \sqrt{3}$ ground state^{39,40}. Here, spins around each

hexagon form closed continuous spin loops whose chirality alternates from cell to cell. The configuration satisfies the ice rule and establishes simultaneous charge and spin order, consistent with prior theoretical schematics.

Our VAE-GA optimization pipeline is designed to methodically discover the ground state through a multi-stage process, as outlined in Fig. 1c-e. (Detailed parameters for each stage of the pipeline are provided in the Methods section.)

The process begins with the generation of a large initial dataset of low-energy configurations (Fig. 1c). We use Monte Carlo simulated annealing to cool the system from a high-temperature paramagnetic state towards its ground state. However, the system's geometric frustration creates a vast energy landscape with numerous local minima, trapping the simulation in states characteristic of the Ice II regime and preventing it from consistently reaching the true ground state. This procedure ultimately yields 30,000 distinct configurations, providing the VAE with a rich statistical foundation of physically plausible spin states.

Next, this dataset is used to train the VAE (Fig. 1d). The VAE learns to compress each spin configuration into a compact, 128-dimensional latent vector z . The crucial component here is the decoder, which learns a generative mapping from this latent space back to the physical spin configuration space.

The core of our method is the iterative, closed-loop optimization cycle (Fig. 1e), where the GA and VAE work cooperatively. Instead of manipulating spins directly, the GA operates entirely within the VAE's learned latent space. Each latent vector component z_i is treated as a gene, which undergoes selection, crossover, and mutation to produce new generations of candidate solutions. The fitness of each genome is determined by the total dipolar interaction energy of its decoded spin configuration, with lower energy corresponding to higher fitness.

The system's energy is governed by the classical dipolar Hamiltonian:

$$H = D \sum_{i < j} \left[\frac{(\mathbf{S}_i \cdot \mathbf{S}_j)}{|\mathbf{r}_{ij}|^3} - \frac{3(\mathbf{S}_i \cdot \mathbf{r}_{ij})(\mathbf{S}_j \cdot \mathbf{r}_{ij})}{|\mathbf{r}_{ij}|^5} \right] = -\frac{1}{2} \sum_i \mathbf{H}_i \cdot \mathbf{S}_i \quad (1)$$

where \mathbf{S}_i and \mathbf{S}_j are the magnetic moment vectors of two spins, \mathbf{r}_{ij} is the vector connecting them, and D is the dipolar coupling constant, and \mathbf{H}_i is the magnetic dipole field acting on \mathbf{S}_i . For simplicity, all energy values reported in this work are normalized by the nearest-neighbor interaction strength, effectively setting the constant $D=1$.

The total energy is then simply the sum of these individual contributions, $E_{\text{total}} = \sum_i E_i$. A lower total energy

corresponds to a higher fitness for the GA. To avoid premature convergence to local minimum, we employ rank-based adaptive methods for crossover and mutation. This strategy uses the relative ranking of individuals, rather than their absolute energy, to maintain population diversity and navigate the rugged landscape more effectively⁴¹.

This optimization process has a two-loop structure. Within a single main cycle, the decoder's parameters remain fixed while the GA runs for multiple generations, replacing less fit genomes with new improved genomes. At the end of the cycle, the top 50 lowest-energy configurations (1% of a population of 5,000) are selected as elites. This 1% elite ratio is a deliberate choice to balance exploration and exploitation within the latent space. In frustrated magnetic systems with exceptionally rugged energy landscapes, an overly large elite pool can induce premature convergence and reduce configurational diversity, whereas an overly small pool may fail to capture the low-energy manifold efficiently. The elite configurations are appended to the training dataset, and the VAE decoder is retrained for the next cycle, enabling the generative model to progressively refine its representation of the ground-state manifold. This differs from non-iterative VAE-GA schemes³³⁻³⁵, in which the generative model is not updated using newly discovered candidates. As a result, the search progressively focuses on the true low-energy manifold and explores regions that are difficult to reach with conventional Monte Carlo simulated annealing.

Performance validation in large Kagome arrays

Our results conclusively demonstrate that the VAE-GA pipeline successfully identifies the true ground state for a large Kagome spin ice with zigzag terminations (Fig. 2), outperforming conventional Monte Carlo simulated annealing which becomes trapped in local energy minimum.

The lowest-energy configuration discovered by our VAE-GA pipeline (Fig. 2a) has a central region exhibiting the perfect $\sqrt{3} \times \sqrt{3}$ ground state order consistent with the known bulk ground state, confirming that our method reaches the ground state. Notably, spins at the boundary do not form continuous circulation patterns. Although such circulation could lower local boundary energy, it disrupts the interior $\sqrt{3} \times \sqrt{3}$ ground state and ultimately increases the total energy. In stark contrast, the lowest energy state found by extensive Monte Carlo simulated annealing is a metastable, Ice II configuration (Fig. 2b). The ice rule is satisfied with alternating vertex charges, but the hexagon loop winding does not alternate globally, instead forming finite same-chirality domains separated by domain walls. The energy difference is significant, as quantified in the energy histogram (Fig. 2c). The ground state energy identified by the VAE-GA is $\varepsilon_{GA} = \frac{E_{total}}{N} \approx -12.29$, where N is number of spins. This value lies decisively

outside and below the entire distribution of low-energy states accessible to Monte Carlo simulated annealing, proving its ability to escape the energy landscape's local minimum.

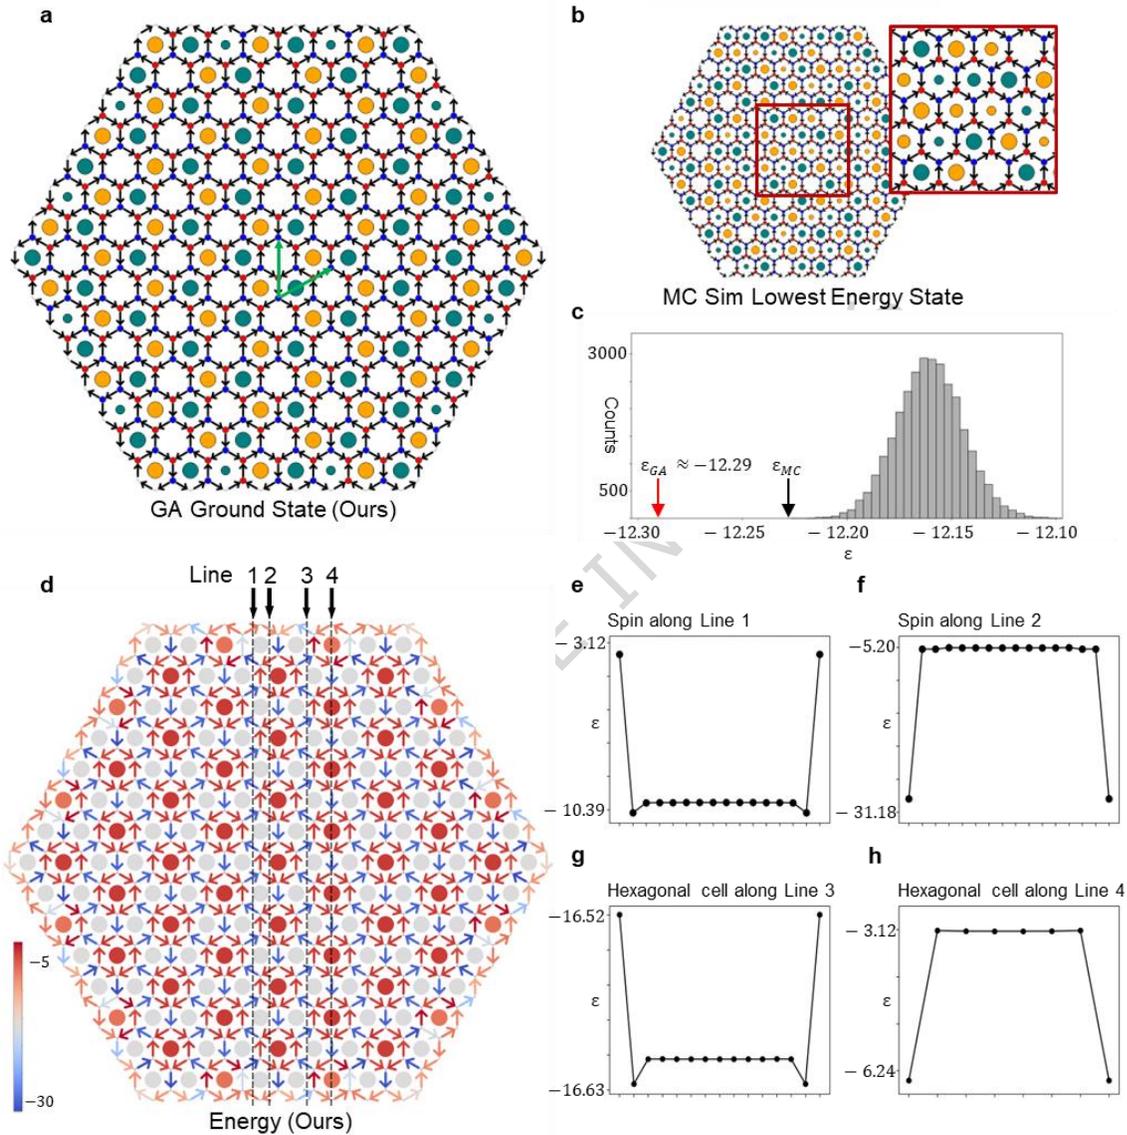


Figure 2. The ground state identified by the VAE-GA pipeline, which demonstrates superior performance over Monte Carlo simulated annealing.

(a) The ground state for a large hexagonal array, as found by our VAE-GA pipeline, exhibiting a perfect $\sqrt{3} \times \sqrt{3}$ order. (b) Lowest energy state from Monte Carlo simulated annealing, trapped in a disordered Ice II state with domain walls. (c) Energy histogram of Monte Carlo simulated annealing states, quantifying the significant energy difference

between the VAE-GA ground state (red arrow) and the Monte Carlo simulated annealing states. (d) Spatial energy map of the ground state in (a). (e–h) Energy profiles along the paths in (d), revealing that boundary perturbations decay rapidly into the bulk within a length scale of approximately two hexagonal cells

Having confidently identified the true ground state, we can use it to perform a precise analysis of its intrinsic properties. A spatial energy map of this VAE-GA derived ground state is presented (Fig. 2d), where the color of each spin represents its individual dipolar energy ε , and the color at the center of each hexagonal cell indicates the average energy of the six spins located at the edges of that cell. This map visually confirms the underlying energetics of the system, showing that the alternating $w = +1$ and $w = -1$ hexagonal cells that constitute the $\sqrt{3} \times \sqrt{3}$ ground state are the most energetically favorable. In contrast, hexagonal cells with zero or fractional winding numbers, which are prevalent in higher-energy metastable states like the one in Fig. 2b, carry a significant energy penalty. To quantify how deeply boundary effects penetrate the bulk, we examine energy profiles along the four paths indicated by dashed lines in Fig. 2d. Line 1 follows spins between a $w = +1$ and a $w = -1$ hexagonal cell. Line 2 follows spins between a $w = +1$ and a $w = 0$ hexagonal cell. Line 3 traces the centers of hexagonal cells with alternating $w = +1$ and $w = -1$ winding numbers. Line 4 traces the centers of $w = 0$ hexagonal cells.

The resulting energy profiles (Fig. 2e–h) reveal the decay of these boundary effects. For the line-1 spins (Fig. 2e), the boundary site exhibits elevated energy, which drops sharply to a minimum at the first interior site. It then rises slightly at the second interior site before stabilizing at the bulk value. For the line-2 spins (Fig. 2f), the boundary energy is lowest, increases steeply one step inward, and then reaches the bulk level without such an overshoot. The hexagonal cell-center profiles show similar trends. The line-3 hexagonal cells (Fig. 2g) have high energy at the edge that dips to a minimum and then slightly recovers before stabilizing from the third hexagonal cell inward, while the line-4 hexagonal cells (Fig. 2h) show a rapid monotonic increase toward the bulk value.

From this detailed analysis, a clear and powerful result emerges. In all cases, the influence of the open zigzag boundary decays rapidly, with local energies converging to those of an infinite system within a length scale of approximately two hexagonal cells. This rapid stabilization reflects the energetic resilience of the $\sqrt{3} \times \sqrt{3}$ magnetic order against boundary perturbations. In the Kagome lattice, long-range dipolar interactions stabilize the $\sqrt{3} \times \sqrt{3}$ superstructure,

making the energy cost of perturbing the ordered interior larger than any local energy gain from boundary rearrangements. Consequently, boundary-induced disorders have an intrinsically short correlation length. This short decay depth provides a physical justification for using sufficiently large finite arrays with open boundaries as reliable proxies for bulk ground-state properties.

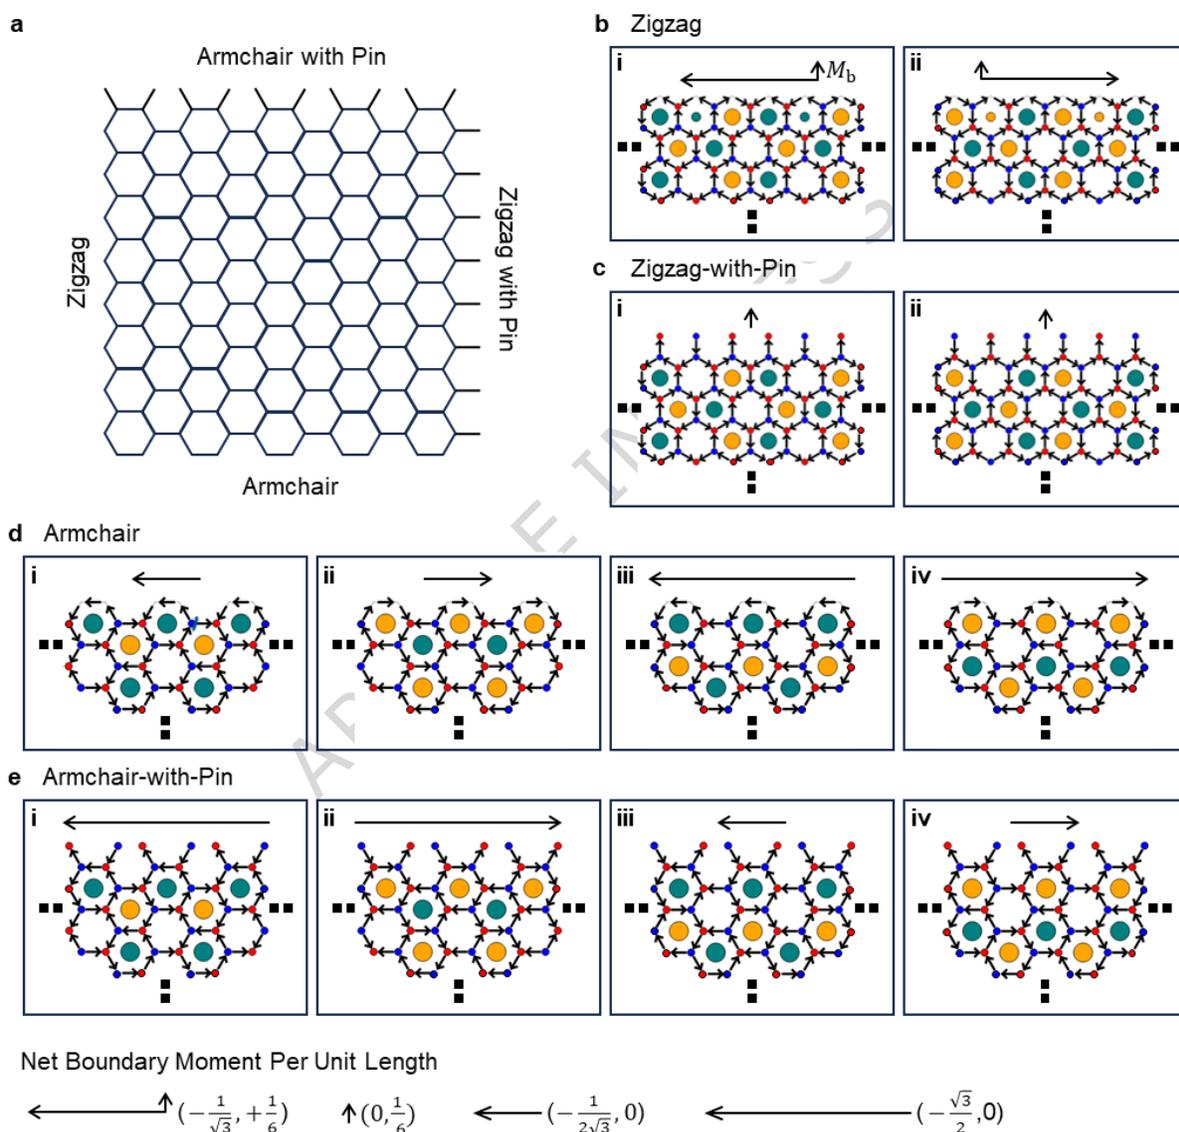


Figure 3 Ground state configurations for basic boundary terminations.

(a) A schematic illustrating the basic four boundaries (zigzag, zigzag-with-pin, armchair, and armchair-with-pin). (b-e) The ground states found by the VAE-GA for each of the four long boundary terminations, respectively. The arrows next to the boundary indicate the direction and relative magnitude of the net boundary moments. Double square dots

indicate that the visualization is truncated.

Ground state selection by specific boundary terminations

To investigate how boundaries of the Kagome lattice affect ground-state ordering, we analyzed the four most representative open-boundary geometries (Fig. 3a). The four boundary terminations shown in Fig. 3a correspond to the four basic boundary terminations, namely zigzag, zigzag-with-pin, armchair, and armchair-with-pin. We used our VAE-GA pipeline to find the ground-state configurations for each of these four boundary types, which are detailed in Fig. 3b-e. For each configuration, arrows next to the boundaries indicate the direction and relative magnitude of the net boundary magnetic moment per unit length M_b , defined as the vector sum of the outermost spins normalized by the length of the boundary.

The $\sqrt{3} \times \sqrt{3}$ magnetic superstructure in the infinite array consists of three types of unit cells characterized by different winding numbers $w = +1, -1$, and 0 . These three cells are fully exchangeable and give an overall S_3 symmetry that encompasses three Z_2 sub-symmetries of pairwise exchanges. The boundary profoundly disrupts the symmetry by either introducing fractional winding number cells or breaking Z_2 sub-symmetry, depending on the boundary types. Moreover, the boundary-induced symmetry breaking is associated with the boundary magnetization.

The zigzag termination exhibits a boundary with a repeating winding number sequence of $w = +1, -1$, and $+\frac{1}{3}$ (Fig. 3b. i). The fractional value arises because boundary vertices have a reduced number of neighbors, and this configuration stabilizes a $q = 0$ vertex that would otherwise become an energetically costly $q = 2$ vertex. Its time-reversed degenerate state simply flips the signs of all spins (Fig. 3b. ii). Due to the fractional cells, these two configurations produce a net boundary moment pointing left and right in Fig. 3b. i and ii, respectively

In the zigzag-with-pin termination, the boundary naturally induces a pinned-spin chain with an up-down-down sequence along the outermost row (Fig. 3c). This causes the fractional winding number to collapse to $w = 0$ (panel i) and minimizes boundary stray fields while maintaining the ice rule ($|q| = 1$) at all vertices. Its time-reversed counterpart is shown in panel ii. Crucially, the spin arrangement in both states causes the net boundary moment to cancel to zero. Therefore, zigzag-with-pin termination does not break the S_3 symmetry of the $\sqrt{3} \times \sqrt{3}$ magnetic

superstructure.

The armchair termination is more complex, supporting two distinct configurations of ground states (Fig. 3d). In this type of boundary, the S_3 symmetry is broken and leaves only a residual Z_2 exchange symmetry between $w = +1$ and -1 sites. The first configuration is characterized by a purely alternating winding number of $w = -1, +1, -1, +1$ (panel i) and its time-reversal (panel ii). The second configuration features a sequence with winding numbers $w = -1, 0, -1, 0$ (panel iii) and its time-reversal with $w = +1, 0, +1, 0$ (panel iv). Note that $w = \pm 1$ cells and $w = 0$ cells are not exchangeable, since the $w = 0$ cells are located half a unit cell farther from the boundary than $w = \pm 1$ cells. The realization of one configuration over the other is determined by a subtle interplay between their intrinsic energy difference and the relative alignment between the armchair cut and the first interior row of spins, a dependence we will quantify in the following section. The boundary parallel magnetic moment points left, right, left, and right in panels i to iv, respectively. The surface-induced symmetry breaking penetrates into the bulk region, thereby destroying the global symmetry of the system.

Finally, the armchair-with-pin termination (Fig. 3e) is the complementary facet created when the lattice is cut to form the armchair boundary (Fig. 3d). Since both facets originate from a single cut, they must consistently terminate the same interior bulk order, linking their ground states in a specific, non-trivial pairing. Specifically, the states in Fig. 3d(i-ii) partner with those in Fig. 3e(iii-iv), and the states in Fig. 3d(iii-iv) partner with those in Fig. 3e(i-ii).

This pairing rule has a direct consequence on the total boundary magnetization. As expected for opposing facets, each complementary pair possesses a net boundary moment of equal magnitude but opposite direction: the moment of the state in Fig. 3d(i) is equal and opposite to that of the state in Fig. 3e(iii), and the same holds for the states in Fig. 3d(iii) and Fig. 3e(i). This reveals that the magnetic configuration of a boundary is fundamentally determined by the properties of its complementary partner, as both are constrained by the need to terminate the interior bulk order.

We now investigate how the ground state configurations are combined in finite-sized arrays where multiple terminations coexist and interact. This interplay between the global array shape and the local boundary terminations ultimately selects the final ground state, as we explore under various geometries by VAE-GA pipeline (Fig. 4).

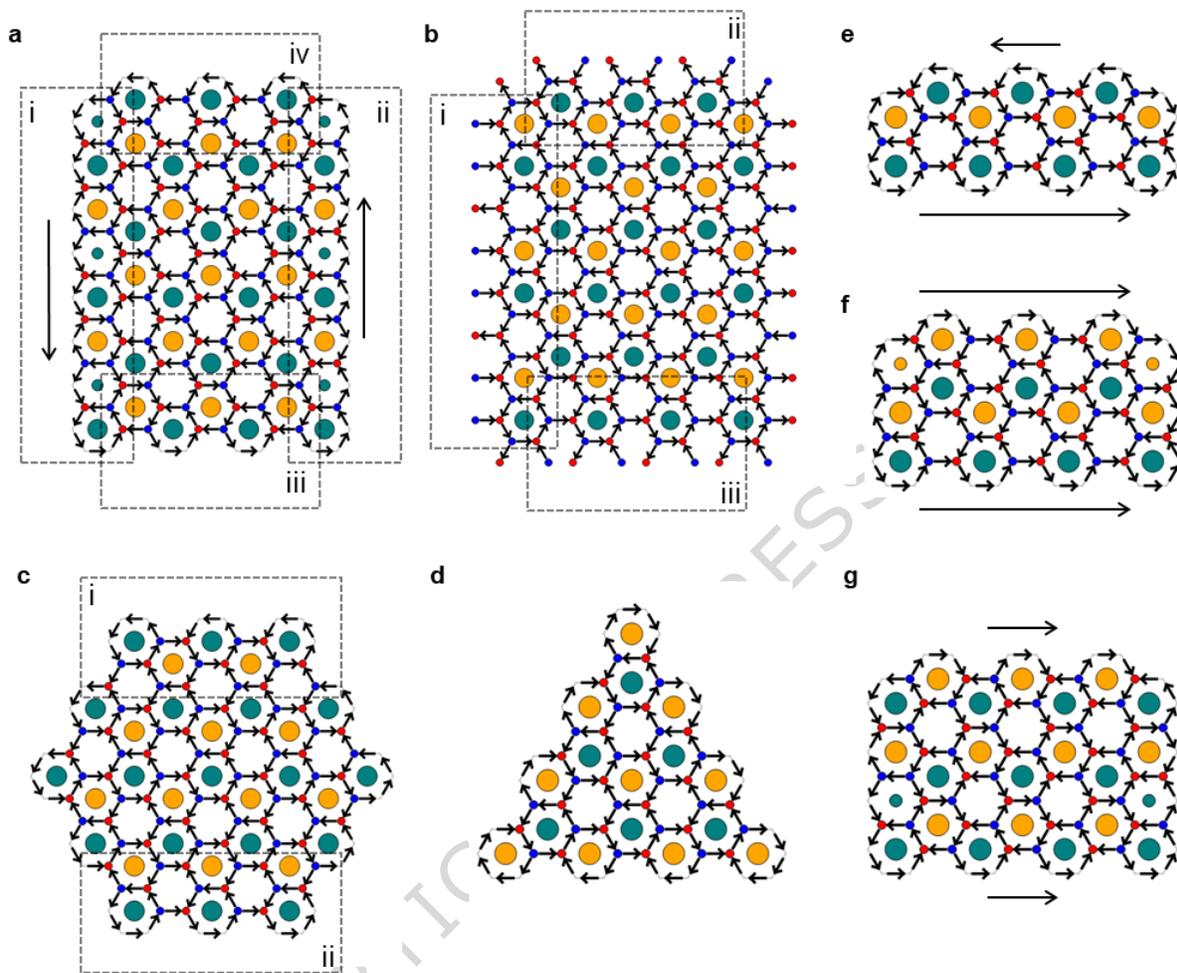


Figure 4 Geometry-Driven Ground States in Finite Kagome Spin Ice

(a) A rectangular array stabilizing different boundary states on its zigzag and armchair terminations. (b) The pinned counterpart to (a), driving a global reconfiguration of both boundary and interior states. (c) An armchair-only hexagonal array, which is forced to select two different armchair states to avoid high-energy corner defects. (d) An armchair-only triangular array, which exclusively selects the energetically favorable $w = \pm 1$ alternating state. (e-g) A series of rectangular arrays with increasing height, demonstrating deterministic switching of the boundary state. Throughout the figure, dashed boxes (labeled i-iv) highlight specific boundary terminations. The large, solid arrows in (a) and (e-g) indicate the direction of the net magnetic moment for the corresponding boundaries.

First, we analyze a rectangular array with two zigzag and two armchair terminations (Fig. 4a). A notable feature is that

the two structurally identical zigzag boundaries (panels i and ii) stabilize ground states (Fig. 3b-i) with opposite magnetic moments. Concurrently, the armchair boundaries stabilize two distinct termination types (panels iii and iv), corresponding to the $w = \pm 1$ alternating state (Fig. 3d-i) and the state containing $w = 0$ cells (Fig. 3d-iii), respectively. Despite these boundary variations, the array's interior maintains a perfect $\sqrt{3} \times \sqrt{3}$ ground state.

Next, we examine the counterpart to this array where all terminations include pins (Fig. 4b). This introduces a significant reconfiguration of both the boundary and the interior. A key driving force for this change is the reduction of high-energy boundary hexagonal cells, as the energy map (Fig. 2d) indicates that hexagonal cells with fractional or zero winding numbers are energetically costly. For the standard zigzag termination shown in Fig. 4a, simply reducing the number of high-energy $w = \frac{1}{3}$ cells is not an option, as such a configuration would violate the ice rule at the boundary. The introduction of pins, however, resolves this geometric constraint. This allows the system to find a new, stable ground state that reduces the number of high-energy boundary cells from three ($w = \frac{1}{3}$) to just two ($w = 0$) while satisfying the ice rule everywhere. This boundary optimization propagates inward, forcing a rearrangement of the interior spins. While the bulk remains in a $\sqrt{3} \times \sqrt{3}$ state, the spin arrangement of the first interior row adjacent to the armchair boundary is fundamentally altered. The specific boundary states realized are shown in the insets (panels i, ii, and iii). These states correspond to the ground states detailed in Fig. 3c-ii, Fig. 3e-i, and Fig. 3e-iii, respectively. To isolate and understand the competition between the two armchair states, we designed arrays with specific symmetries that would favor one state over the other. For an armchair-only hexagonal geometry (Fig. 4c), the system is geometrically forced to use both armchair types to avoid high-energy ice-rule violations at the corners. For instance, the top boundary (panel i) stabilizes the $w = \pm 1$ alternating state (Fig. 3d-i), while the bottom boundary (panel ii) stabilizes the state containing $w = 0$ cells (Fig. 3d-iii). In contrast, for an armchair-only triangular geometry (Fig. 4d), the symmetry allows for a uniform boundary type. Here, the system exclusively selects the $w = \pm 1$ alternating state (Fig. 3d-i), thereby revealing that this state is intrinsically lower in energy than the state containing $w = 0$ cells. Finally, we demonstrate that this intrinsic energy preference can be overridden, and the boundary state can be deterministically switched. By systematically changing the height of a rectangular array (Fig. 4e-g), we can precisely control the relative alignment between the boundary and the interior lattice. Starting from a $(\frac{7\sqrt{3}}{2} \times 2)$ array (Fig. 4e), the top boundary stabilizes the $w = \pm 1$ alternating state while the bottom stabilizes the state containing $w = 0$ cells.

Increasing the height to $(\frac{7\sqrt{3}}{2} \times 3)$ array (Fig. 4f) flips the top boundary to the less favorable state containing $w = 0$ cells, while the bottom remains unchanged. A further increase to $(\frac{7\sqrt{3}}{2} \times 4)$ array (Fig. 4g) reverts the top boundary to the $w = \pm 1$ alternating state while the bottom flips to its time-reversed variant, as visualized by the reversal of the boundary moment arrow. This demonstrates that the boundary effect is a powerful enough mechanism to force the system into a locally higher-energy boundary configuration to minimize the global energy. Across all three sizes, the interior remains robustly in the $\sqrt{3} \times \sqrt{3}$ ground state.

Taken together, these observations reveal a fundamental distinction between the two main boundary types. While both zigzag and armchair terminations robustly preserve the interior $\sqrt{3} \times \sqrt{3}$ superstructure order, the armchair termination exhibits high boundary sensitivity by selecting the energy-optimal state based on the global geometry. This implies that boundary moments can be engineered through geometric design alone, without external fields, providing concrete principles for the rational design of magnetic metamaterials with structurally dictated global textures. In contrast, zigzag terminations exhibit robust local behavior, consistent with the absence of competing low-energy boundary families.

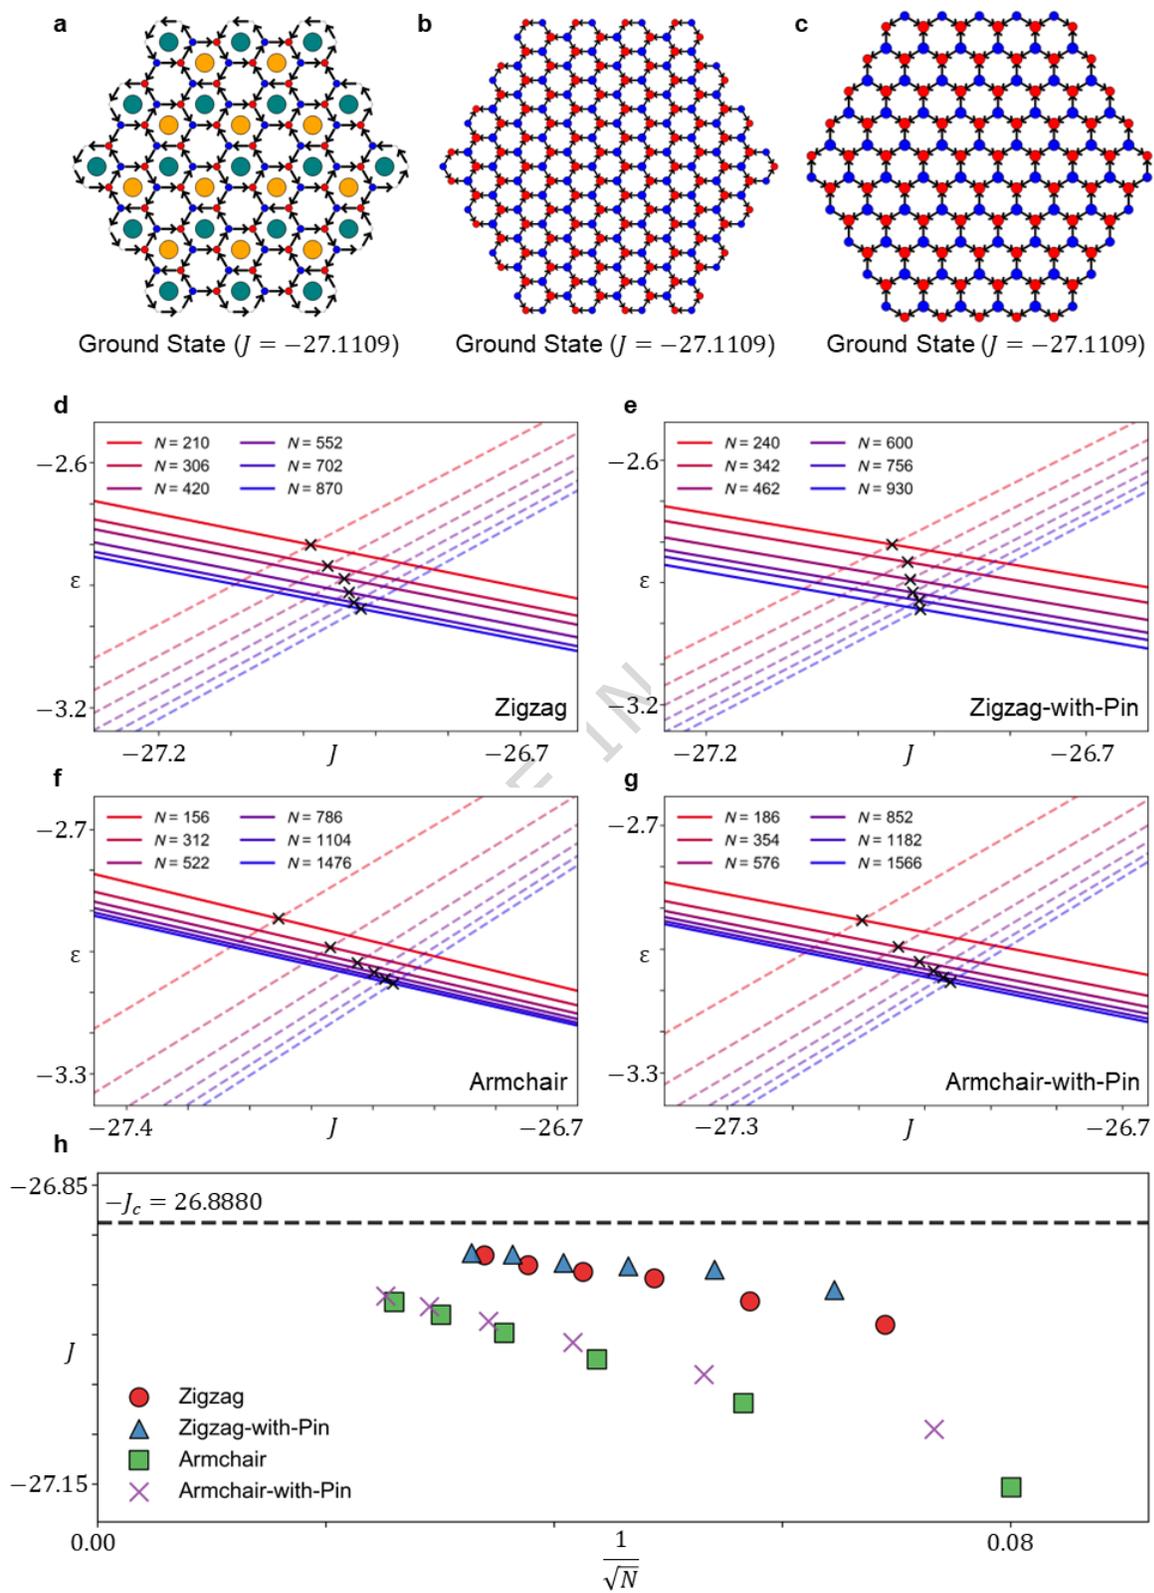


Figure 5. Influence of magnetic exchange interaction J and boundary-dependent phase transitions.

Ground state configurations identified by the VAE-GA framework under a constant exchange interaction of $J = -27.1109$. (a) An armchair lattice with $N = 156$ retains the $\sqrt{3} \times \sqrt{3}$ state, as its critical threshold $|J_c|$ is higher than the applied J . (b) A larger armchair lattice $N = 312$ under the same condition transitions to the antiferromagnetic (AFM) phase ($|q| = 3$), demonstrating that the stability of the ice phase diminishes as system size increases. (c) A zigzag lattice $N = 210$ with a size similar to (a) exhibits the AFM phase, indicating that zigzag boundaries possess a lower critical threshold compared to the armchair case of similar size. (d–g) Energy crossing diagrams for four boundary conditions: (d) Zigzag, (e) Zigzag-with-Pin, (f) Armchair, and (g) Armchair-with-Pin. The solid lines represent the energy of the $\sqrt{3} \times \sqrt{3}$ state, which includes local spin modifications induced by boundaries as discussed in Fig. 3. The dashed lines represent the energy of the AFM charge-ordered phase. The black crosses mark the critical points J_c where the phase transition occurs. (h) Convergence of the critical points toward the thermodynamic limit. The extracted J_c values are plotted against the inverse square root of the spin count, $\frac{1}{\sqrt{N}}$. The horizontal dashed line represents the theoretical limit ($J_c \approx -26.8880$) derived from infinite lattice summation, confirming that all finite-size boundaries converge to a universal bulk value.

Competition between dipolar and exchange interactions

To incorporate microscopic interactions relevant to solid-state materials^{42–46}, we extended the pure dipolar model by explicitly adding a nearest-neighbor exchange interaction term $H = -\frac{J}{2} \sum \mathbf{S}_i \cdot \mathbf{S}_j$ to Eq. (1). In the regime where $J > 0$ (ferromagnetic coupling), the exchange interaction favors spin alignments that are consistent with the existing dipolar order, thereby preserving the $\sqrt{3} \times \sqrt{3}$ ground state without inducing a phase transition. In contrast, a sufficiently strong negative J (antiferromagnetic coupling) promotes an all-in or all-out alignment at vertices ($|q| = 3$), which directly competes with the dipolar ice rule ($|q| = 1$). Therefore, we focused our investigation on the $J < 0$ regime to elucidate the phase transition driven by this competition. In this context, we define J_c as the critical exchange coupling strength at which the energies of the two competing phases become degenerate, marking the onset of the phase transition.

In this extended model, the system undergoes competition between two distinct magnetic orders. One is the $\sqrt{3} \times \sqrt{3}$

$\sqrt{3}$ dipolar state identified in this study, and the other is the AFM phase, where spins point all-in or all-out at vertices ($|q| = 3$).

Figures 5a–c illustrate the decisive impact of system size and boundary conditions on the ground state selection under a constant exchange interaction of $J = -27.1109$. For the armchair lattice with $N = 156$ (Fig. 5a), the critical point $J_c \approx -27.15$ lies deeper in the energy landscape than the applied J , allowing the system to retain the dipolar ground state. However, as the size increases to $N = 312$ under the same boundary condition (Fig. 5b), the critical point shifts to $J_c \approx -27.07$. Consequently, the applied interaction $J = -27.1109$ exceeds this threshold, triggering a transition to the AFM phase. Furthermore, the zigzag lattice $N = 210$ (Fig. 5c) already exhibits the AFM phase under the same condition. This indicates that, in the finite-size regime, zigzag boundaries possess a lower critical threshold compared to the armchair case, although they eventually converge in the thermodynamic limit.

To quantify the physical origin of this transition, we performed an energy crossing analysis for four equivalent boundary conditions (Figs. 5d–g). In these plots, the energies of the dipolar state (solid lines) and the AFM phase (dashed lines) are linear functions of J . The nearest-neighbor spin correlations in the AFM phase differ in sign from those in the dipolar ground state. Consequently, as the antiferromagnetic coupling strength $|J|$ increases (J becomes more negative), the energy of the AFM phase decreases significantly while the energy of the dipolar ground state increases, inevitably leading to a phase transition at the intersection point. We extracted precise J_c values from 24 simulation sets.

The critical strength of exchange constant J_c increases with the size of structures, furthermore the size dependence on J_c is more profound in armchair type boundaries than zigzag type boundaries. Specifically, for the Zigzag boundary, J_c shifts from -26.9900 to -26.9205 as N increases from 210 to 870. The Zigzag-with-Pin case follows a trajectory from -26.9553 to -26.9182 ($N = 240$ to 930). Meanwhile, the Armchair boundary shows a shift from -27.1533 to -26.9672 ($N = 156$ to 1476), and the Armchair-with-Pin case changes from -27.0952 to -26.9612 ($N = 186$ to 1566).

Finally, we analyzed the convergence of these critical points toward the thermodynamic limit by plotting J_c against the inverse square root of the spin count, $\frac{1}{\sqrt{N}}$ (Fig. 5h). To establish a rigorous analytical reference, we referred to the

effective nearest-neighbor coupling model, typically defined as $J_{eff} = \frac{J}{2} + \alpha D$, where the phase transition is governed by the competition between J and D . While the conventional model based solely on nearest-neighbor dipolar interactions predicts a geometric constant¹⁶ of $\alpha = 1.75$, our large-scale infinite lattice calculations reveal that long-range interactions significantly modify this coefficient depending on the magnetic order. Specifically, constructive interference increases the coefficient to $\alpha \approx 1.8775$ for the ordered ice phase, whereas screening effects reduce it to $\alpha \approx 1.4835$ for the AFM phase. By incorporating these renormalized coefficients into the energy balance equation, we derived a definitive bulk limit of $J_c \approx -26.8880$. As shown in Fig. 5h, the 24 data points scattered by boundary effects in finite sizes all converge toward this theoretical target line as the system approaches the thermodynamic limit $\frac{1}{\sqrt{N}} \rightarrow 0$. This confirms that the observed shifts in transition points are transient surface effects, providing strong evidence that our VAE-GA framework can accurately predict universal bulk properties beyond boundary constraints.

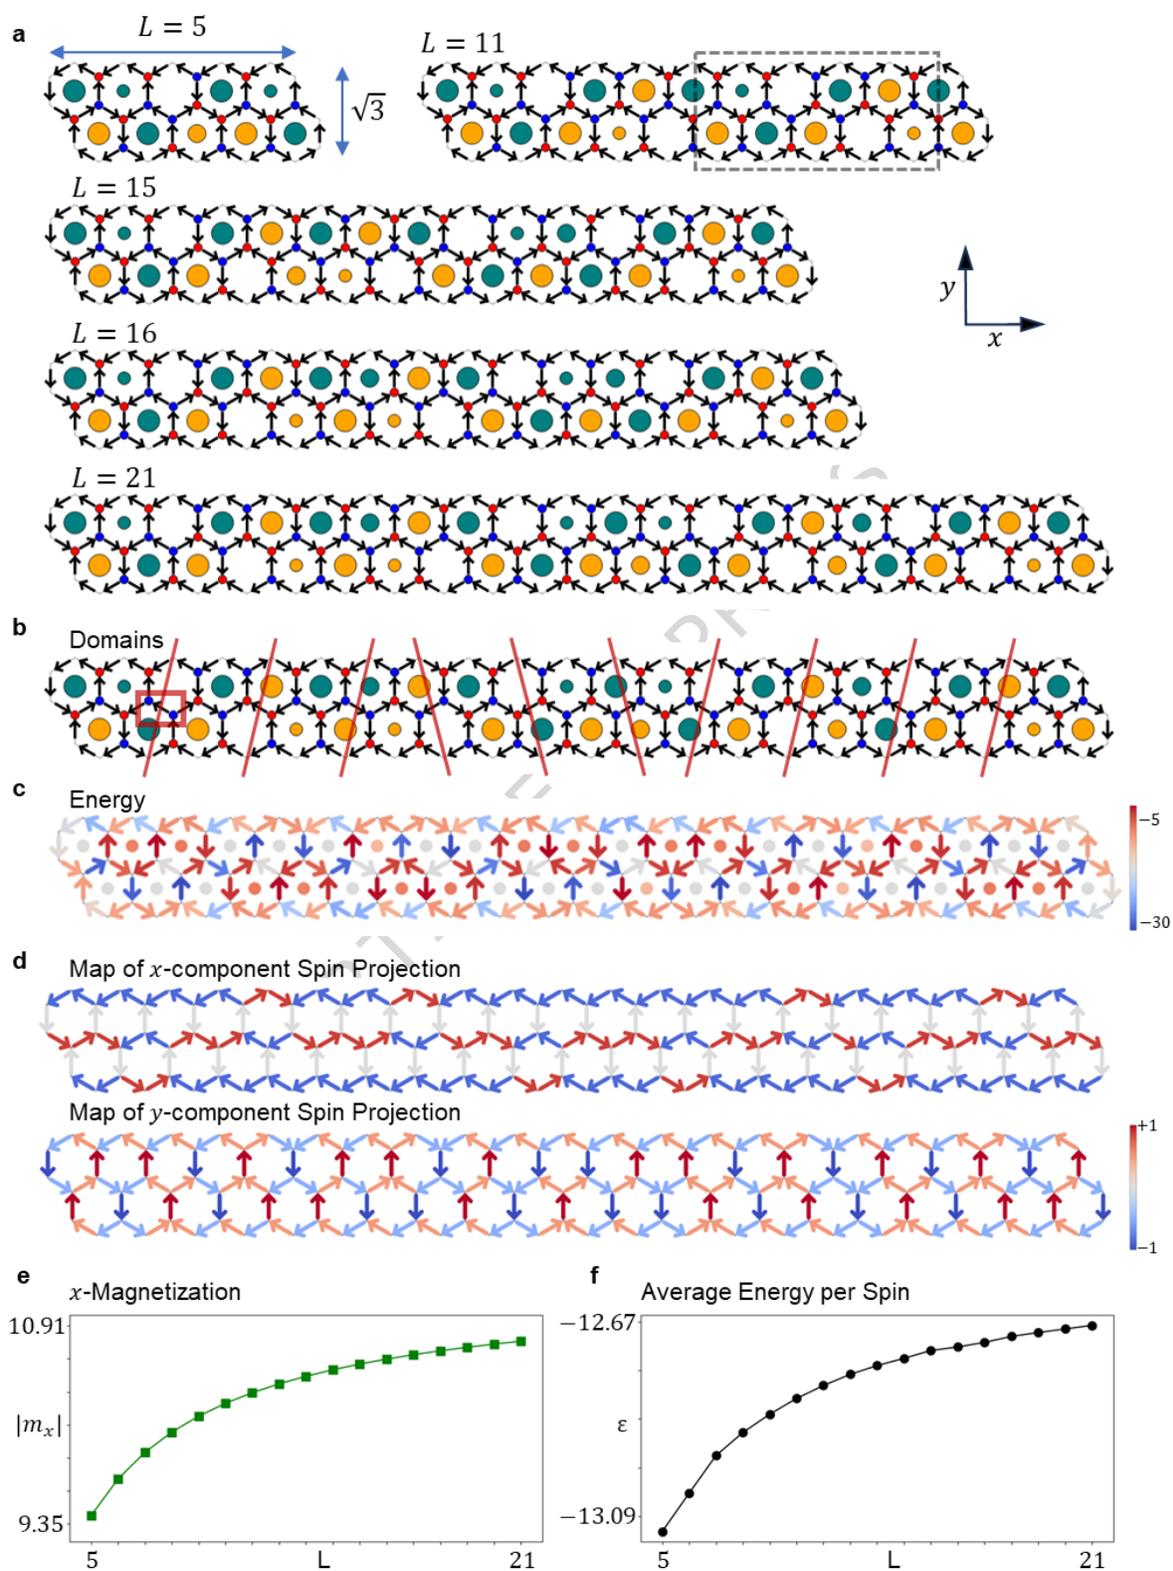


Figure 6 Emergence of a complex quasi-ferromagnetic state under extreme geometric confinement.

(a) Representative ground state spin configurations for $L \times \sqrt{3}$ arrays of varying length L with $L = 5, 11, 15, 16, 21$. The dashed box for the $L=11$ state indicates a pseudo-periodic structure that is not a true long-range order. (b) The $L=21$ ground state, illustrating how Ice II violating spins (red boxes) act as domain walls (red lines) that segment the array. (c) The spatial energy map for the $L = 21$ state (d) Spatial maps of x - and y -component spin projections for the $L = 21$ state. The x -component map (top) reveals a strong uniform moment along the zigzag boundaries, while the y -component map (bottom) shows largely compensated moments. The color scale represents the projection value from -1 (blue) to $+1$ (red). (e) x -Magnetization per spin, $|m_x|$ as a function of array length L . The saturation of $|m_x|$ to a finite value is the definitive signature of long-range order. (f) Average energy per spin as a function of array length L , showing convergence to a stable bulk value. For consistency, we exclusively show the states spontaneously magnetized along the same direction.

Quasi-ferromagnetism under geometric confinement

To probe the effects of extreme geometric confinement, we investigated highly anisotropic arrays with zigzag terminations on the top and bottom boundaries. In this quasi-one-dimensional geometry, the $\sqrt{3} \times \sqrt{3}$ ground state is destabilized. Our VAE-GA pipeline instead discovers a novel family of complex, quasi-ferromagnetic ground states, which we analyze for various array lengths (Fig. 6).

The evolution of these ground states with length, with representative examples shown in Fig. 5a, is remarkably complex. Here, we define the $(L \times \sqrt{3})$ array where L and $\sqrt{3}$ represent the maximum spin-to-spin distances along the x and y axes, respectively (Fig. 6a). While the ideal zigzag boundary rule (Fig. 3b) holds for short lengths ($L = 3, 4$), it breaks down completely for $L \geq 5$. The system enters a disordered Ice I phase but never achieves the alternating charge order of the Ice II phase. This is due to the emergence of spins that are surrounded by same-signed magnetic charges, explicitly violating the Ice II rule. These violating spins always appear in even numbers.

Over the full range of L studied (see Supplementary Fig. S1), we found no simple global ordering principle. This complexity is exemplified by comparing states of different lengths. A striking example of this complexity occurs at $L = 11$, where the system forms a pseudo-periodic pattern with a repeating unit of 5 hexagonal cells, suggestive of a superlattice (dashed box). However, this apparent order is not a true long-range phenomenon but rather an emergent

property specific to this length. This period-5 structure is completely absent at longer lengths, such as $L = 16$ and 21 . For instance, the $L = 15$ state exhibits mirror symmetry in its winding number pattern, but this symmetry is completely broken in the $L = 16$ state. The $L = 21$ state, the longest array studied, appears highly disordered, with its structure defined by internal domain boundaries, further demonstrating the absence of a simple repeating structural unit.

The origin of this complex ordering lies in a competition between bulk and boundary energies. The ground state for $L = 21$, (Fig. 6b) reveals that the Ice II violating spins are not randomly distributed but instead act as domain walls (red lines) that segment the array into smaller, ordered domains. The spatial energy map (Fig. 6c) reveals the energetic consequence of this domain structure.

This cost is offset by an energetic gain at the zigzag boundary, as revealed by the magnetization maps (Fig. 6d). The x -component map shows a strong tendency for spins along the zigzag boundary to align, creating a uniform boundary moment (here, in the $-x$ direction). This is a direct violation of the alternating spin rule for the ideal zigzag termination (Fig. 3b). The y -component map, in contrast, shows a more disordered arrangement with largely compensated moments, consistent with near-zero net y -magnetization. The key insight is that for this specific geometry, the energetic gain from forming this uniform moment along the zigzag boundary outweighs the energetic penalty of creating domain walls in the bulk. This behavior is in stark contrast to that of the large hexagonal array (Fig. 2a), where preserving the bulk order was paramount.

This emergent long-range order is quantitatively captured by magnetization per spin, m . This dimensionless quantity is defined as the vector average of the individual spin orientations, $m = \frac{1}{N} \sum S_i$. We focus on its x -component, $m_x = \frac{1}{N} \sum S_{x,i}$. Its physical magnitude is absorbed into the overall energy normalization of the Hamiltonian (Eq. 1). The $|m_x|$ is plotted as a function of array length L (Fig. 6e). The $|m_x|$ grows and saturates towards a finite value, the definitive signature of long-range ferromagnetic-like order. For one of the two degenerate ground states, the number of spins aligned with the majority x -direction is $4L$, while the number of minority-aligned spins is $2L+1$. The linear growth of this difference $2L - 1$ ensures that the total moment, M_x , also grows linearly with L , while the per-spin moment, m_x , correctly converges to a non-zero bulk value.

The average energy per spin (Fig. 6f) generally converges, confirming the system settles into a well-defined thermodynamic state. However, its convergence is non-monotonic, reflecting the complex structural rearrangements.

For instance, the energy step between $L = 14$ and 15 is notably larger than the step between $L = 13$ and 14. Finding these states represents a formidable computational challenge, requiring up to 200 optimization cycles for $L=21$, which underscores the power of our virtuous-cycle pipeline.

While our discussion has primarily focused on artificial spin ice, experimental systems such as $\text{Dy}_2\text{Ti}_2\text{O}_7$ and $\text{Ho}_2\text{Ti}_2\text{O}_7$ provide realistic spin ice platforms⁴²⁻⁴⁴, and related frustrated systems include Kagome-lattice-based metal-organic frameworks. The experimental determination of a ground state is an extremely challenging problem because the complex energy landscape resulting from long-range dipolar interactions and dynamic constraints on excitations lead to extremely slow relaxation even in the absence of disorder^{42,45}. Our framework can be extended to investigate the roles of boundaries in these systems by incorporating interaction parameters derived from first-principles calculations, such as density functional theory. In these realistic systems, boundaries do not merely truncate links as assumed in artificial spin ice; rather, symmetry breaking can modify the local magnetic structure near the edges, leading to exchange couplings that deviate from bulk values. By integrating these refined interaction parameters, our VAE-GA pipeline can provide a practical route for ground-state searches in finite geometries, as the GA fitness evaluation can effectively accommodate spatially varying interactions.

In summary, our comprehensive investigation reveals that the ground state of finite Kagome spin ice is not solely determined by its bulk properties but is instead governed by a complex interplay between global geometry, local boundary terminations, and energetic competition. We have shown that while some terminations like zigzag exhibit robust, predictable behavior, others like armchair are highly sensitive and configurable, offering a powerful tool to control the system's magnetic texture via registry and shape. Furthermore; under extreme geometric confinement, we discovered that this boundary-bulk competition can completely destabilize the canonical $\sqrt{3} \times \sqrt{3}$ state, leading to the emergence of novel, quasi-ferromagnetic phases with intricate domain structures. These findings provide a clear roadmap for the rational design of frustrated metamaterials. This advance was enabled by our virtuous-cycle VAE-GA pipeline, which iteratively retrains a generative model on elite configurations discovered during the search. This self-improving mechanism navigates rugged energy landscapes that trap conventional methods and represents a powerful, generalizable tool for complex many-body problems.

Method

System Construction and Initial Dataset Generation

We first generated a honeycomb lattice by placing hexagonal centers in a grid. The side length of each hexagon was set to $\frac{1}{\sqrt{3}}$, defining the fundamental length scale of the system. The Kagome lattice was then formed by defining a spin at the midpoint of each edge of this honeycomb structure. Various finite array geometries, such as the hexagons, rectangles, and triangles explored in the main text, were created by starting with a large grid and selectively removing specific hexagonal cells to "carve out" the desired shape and boundary terminations. For terminations requiring pins (e.g., "zigzag with pin"), additional spin sites were geometrically added at the boundary, extending outward from specific vertices of the boundary hexagonal cells.

The initial seed dataset for the VAE-GA pipeline was generated using Monte Carlo simulated annealing. For each geometry, we generated a training set of 30,000 low energy spin configurations. The simulation started from a high-temperature paramagnetic state ($T_{\text{start}} = 3.0$) and cooled over 100,000 iterations to a low-temperature regime ($T_{\text{end}} = 0.01$), with 95% of the iterations dedicated to the annealing schedule. The energy of each spin configuration was calculated using the classical dipolar Hamiltonian defined in Eq. (1), with the interaction range set to 1000.0, a value larger than any system size to effectively include all pairwise interactions.

The Virtuous-Cycle VAE-GA Pipeline

Our core methodology is a closed-loop, virtuous-cycle optimization pipeline that integrates a VAE for representation learning with a Genetic Algorithm (GA) for global exploration.

The VAE architecture consists of a symmetric encoder and decoder; each built as a fully connected network. The encoder consisted of three hidden layers (512, 384, and 256 neurons, respectively), and the decoder had a symmetric structure (256, 384, and 512 neurons). All hidden layers utilized the LeakyReLU activation function and were followed by Batch Normalization. The VAE compresses each N-dimensional spin configuration into a 128-dimensional latent vector z . The VAE was trained using the Adam optimizer with a learning rate of 1×10^{-3} . The loss function included a mean squared error (MSE) reconstruction term and a Kullback–Leibler (KL) divergence term, with the KL-term weight set to 0.005.

The GA operates entirely in the 128-dimensional latent space learned by the VAE. The fitness of each individual (a latent vector z) is defined as the negative of the total dipolar energy of its decoded spin configuration. We used a

population size of 5,000 individuals. The rugged energy landscape of spin ice contains numerous metastable states with subtle energy differences, where a purely fitness-based approach can prematurely converge in local minimum. Therefore, we employed rank-based adaptive methods for both crossover and mutation. This strategy uses the relative ranking of individuals, rather than their absolute fitness, to create a more stable selection pressure that maintains population diversity and effectively navigates the complex landscape. The selection of parents for the next generation was performed using stochastic remainder selection. For the periodic deep search cycles employed for the most challenging geometries, the mutation rate was linearly annealed from an initial value of 1.0 down to 0.01 over the first half of the generations, then held constant at 0.01 for the remainder. For standard cycles, it was kept constant at 1.0.

The full optimization process is an iterative loop. First, the VAE is trained on the current dataset for up to 100 epochs with a batch size of 500 and an early stopping condition (monitor='loss', patience=10). The trained decoder is then used as the fitness function for the GA, which runs for a set number of generations to explore the latent space. For most geometries, a standard search of 2,000 generations was sufficient, while for the highly complex $L \times \sqrt{3}$ arrays, we employed periodic deep search cycles of 20,000 generations for every 20 main cycles. At the end of the GA run, the 50 individuals with the highest fitness are identified. These elite configurations are used to augment the training dataset for the subsequent cycle. The VAE is then retrained on this larger dataset, which is now enriched with the newly discovered low-energy configurations, completing the cycle. This process was repeated for up to 200 cycles for the most challenging geometries, allowing the VAE's latent representation to progressively focus on the true low-energy manifold of the system.

Code availability

The Python code used to implement the method in this study is available on GitHub (<https://github.com/NanomagLab/Kagome-AI-Optimizer>)

Funding

This research was supported by the National Research Foundation (NRF) of Korea funded by the Korean Government (NRF-2023R1A2C1006050) and (NRF-2021R1C1C2093113).

Author Contributions

T.J.M. developed the algorithms and performed the experiments. The main results were discussed and interpreted with contributions from S.M.P., H.G.Y., H.Y.K., and C.W. The research was supervised jointly by H.Y.K. and C.W. All authors contributed to the final version of the manuscript.

Competing Interests

The authors declare no competing interests.

Reference

1. Barahona, F. On the computational complexity of ising spin glass models. *J. Phys. A Math. Gen.* **15**, (1982).
2. Fan, C. *et al.* Searching for spin glass ground states through deep reinforcement learning. *Nat. Commun.* **14**, (2023).
3. De Simone, C. *et al.* Exact ground states of Ising spin glasses: New experimental results with a branch-and-cut algorithm. *J. Stat. Phys.* **80**, (1995).
4. Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B. & Troyansky, L. Determining computational complexity from characteristic 'phase transitions'. *Nature* **400**, (1999).
5. Hogg, T., Huberman, B. A. & Williams, C. P. Phase transitions and the search problem. *Artificial Intelligence* vol. 81 Preprint at [https://doi.org/10.1016/0004-3702\(95\)00044-5](https://doi.org/10.1016/0004-3702(95)00044-5) (1996).
6. Heyderman, L. J. & Stamps, R. L. Artificial ferroic systems: Novel functionality from structure, interactions and dynamics. *Journal of Physics Condensed Matter* **25**, (2013).
7. Wang, R. F. *et al.* Artificial 'spin ice' in a geometrically frustrated lattice of nanoscale ferromagnetic islands. *Nature* **439**, (2006).
8. Skjærvø, S. H., Marrows, C. H., Stamps, R. L. & Heyderman, L. J. Advances in artificial spin ice. *Nature Reviews Physics* vol. 2 Preprint at <https://doi.org/10.1038/s42254-019-0118-3> (2020).
9. Nisoli, C., Moessner, R. & Schiffer, P. Colloquium: Artificial spin ice: Designing and imaging magnetic frustration. *Rev. Mod. Phys.* **85**, (2013).

10. Rougemaille, N. & Canals, B. Cooperative magnetic phenomena in artificial spin systems: spin liquids, Coulomb phase and fragmentation of magnetism – a colloquium. *European Physical Journal B* **92**, (2019).
11. Farhan, A. *et al.* Emergent magnetic monopole dynamics in macroscopically degenerate artificial spin ice. *Sci. Adv.* **5**, (2019).
12. Mengotti, E. *et al.* Real-space observation of emergent magnetic monopoles and associated Dirac strings in artificial kagome spin ice. *Nat. Phys.* **7**, (2011).
13. Qi, Y., Brintlinger, T. & Cumings, J. Direct observation of the ice rule in an artificial kagome spin ice. *Phys. Rev. B Condens. Matter Mater. Phys.* **77**, (2008).
14. Gilbert, I. *et al.* Emergent ice rule and magnetic charge screening from vertex frustration in artificial spin ice. *Nat. Phys.* **10**, (2014).
15. Canals, B. *et al.* Fragmentation of magnetism in artificial kagome dipolar spin ice. *Nat. Commun.* **7**, (2016).
16. Chern, G. W., Mellado, P. & Tchernyshyov, O. Two-stage ordering of spins in dipolar spin ice on the kagome lattice. *Phys. Rev. Lett.* **106**, (2011).
17. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, (2017).
18. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science (1979)*. **355**, (2017).
19. Wang, H. *et al.* Scientific discovery in the age of artificial intelligence. *Nature* vol. 620 Preprint at <https://doi.org/10.1038/s41586-023-06221-2> (2023).
20. Wetzels, S. J. Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders. *Phys. Rev. E* **96**, (2017).
21. Carleo, G., Nomura, Y. & Imada, M. Constructing exact representations of quantum many-body systems with deep neural networks. *Nat. Commun.* **9**, (2018).
22. Park, S. M. *et al.* Optimization of physical quantities in the autoencoder latent space. *Sci. Rep.* **12**, (2022).
23. Li, J. *et al.* An Auto-Encoder with Genetic Algorithm for High Dimensional Data: Towards Accurate and Interpretable Outlier Detection. *Algorithms* **15**, (2022).

24. Kwon, H. Y. *et al.* Searching for the ground state of complex spin-ice systems using deep learning techniques. *Sci. Rep.* **12**, (2022).
25. Kwon, H. Y. *et al.* Magnetic State Generation using Hamiltonian Guided Variational Autoencoder with Spin Structure Stabilization. *Advanced Science* **8**, (2021).
26. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. in *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings* (2014). doi:10.61603/ceas.v2i1.33.
27. Kingma, D. P. & Welling, M. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning* vol. 12 Preprint at <https://doi.org/10.1561/22000000056> (2019).
28. Doersch, C. Variational Autoencoders Tutorial. *ArXiv* (2016).
29. Bojanowski, P., Joulin, A., Paz, D. L. & Szlam, A. Optimizing the latent space of generative networks. in *35th International Conference on Machine Learning, ICML 2018* vol. 2 (2018).
30. Bentley, P. J., Lim, S. L., Gaier, A. & Tran, L. COIL: Constrained Optimization in Learned Latent Space: Learning Representations for Valid Solutions. in *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (2022).
31. Bentley, P. J., Lim, S. L., Gaier, A. & Tran, L. Evolving Through the Looking Glass: Learning Improved Search Spaces with Variational Autoencoders. in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol. 13398 LNCS (2022).
32. Volz, V. *et al.* Evolving Mario levels in the latent space of a deep convolutional generative adversarial network. in *GECCO 2018 - Proceedings of the 2018 Genetic and Evolutionary Computation Conference* (2018). doi:10.1145/3205455.3205517.
33. Rammal, A., Ezukwoke, K., Hoayek, A. & Batton-Hubert, M. Unsupervised approach for an optimal representation of the latent space of a failure analysis dataset. *Journal of Supercomputing* **80**, (2024).
34. Jiang, J., Kim, D., Kim, B. & Shin, Y. G. GA-VAE: Enhancing Local Feature Representation in VQ-VAE Through Genetic Algorithm-Based Token Optimization. *IEEE Access* **13**, (2025).
35. Wan, J., Chen, Y. & Gao, C. The VAE-FastGA anomaly detection model based on subspace and weakly correlated ultra-high-dimensional data. *Journal of Supercomputing* **81**, (2025).

36. Kincaid, R. K. & Ninh, A. Simulated Annealing. in *Springer Optimization and Its Applications* vol. 204 (2023).
37. Tripp, A., Daxberger, E. & Hernández-Lobato, J. M. Sample-efficient optimization in the latent space of deep generative models via weighted retraining. in *Advances in Neural Information Processing Systems* vols 2020-December (2020).
38. Lookman, T., Balachandran, P. V., Xue, D. & Yuan, R. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials* vol. 5 Preprint at <https://doi.org/10.1038/s41524-019-0153-8> (2019).
39. Gartside, J. C. *et al.* Realization of ground state in artificial kagome spin ice via topological defect-driven magnetic writing. *Nat. Nanotechnol.* **13**, (2018).
40. Zhao, K. *et al.* Realization of the kagome spin ice state in a frustrated intermetallic compound. *Science (1979)*. **367**, (2020).
41. Basak, A. A Rank based Adaptive Mutation in Genetic Algorithm. *Int. J. Comput. Appl.* **175**, (2020).
42. Bermúdez, M. M., Borzi, R. A., Tennant, D. A. & Grigera, S. A. Low-temperature state of the spin-ice material Dy₂Ti₂O₇: Reverse Monte Carlo investigation. *Phys. Rev. B* **112**, (2025).
43. Yavors'kii, T., Fennell, T., Gingras, M. J. P. & Bramwell, S. T. Dy₂Ti₂O₇ spin ice: A test case for emergent clusters in a frustrated magnet. *Phys. Rev. Lett.* **101**, (2008).
44. Borzi, R. A. *et al.* Intermediate magnetization state and competing orders in Dy₂ Ti₂ O₇ and Ho₂ Ti₂ O₇. *Nat. Commun.* **7**, (2016).
45. Samarakoon, A. M. *et al.* Machine-learning-assisted insight into spin ice Dy₂Ti₂O₇. *Nat. Commun.* **11**, (2020).
46. Henelius, P. *et al.* Refrustration and competing orders in the prototypical Dy₂Ti₂ O₇ spin ice material. *Phys. Rev. B* **93**, (2016).