

Deep generative learning of magnetic frustration in artificial spin ice from magnetic force microscopy images

Received: 14 October 2025

Accepted: 29 April 2026

Cite this article as: Neogi, A., Mishra, S., Iyer, P.P. *et al.* Deep generative learning of magnetic frustration in artificial spin ice from magnetic force microscopy images. *npj Comput Mater* (2026). <https://doi.org/10.1038/s41524-026-02124-8>

Arnab Neogi, Suryakant Mishra, Prasad P. Iyer, Tzu-Ming Lu, Ezra Bussmann, Sergei Tretiak, Andrew C. Jones & Jian-Xin Zhu

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.

Deep Generative Learning of Magnetic Frustration in Artificial Spin Ice from Magnetic Force Microscopy Images

Arnab Neogi^{1,2*†}, Suryakant Mishra^{2†}, Prasad P. Iyer³,
Tzu-Ming Lu³, Ezra Bussmann³, Sergei Tretiak^{1,2*},
Andrew C. Jones^{2*}, Jian-Xin Zhu^{1,2*}

¹Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States.

²Center for Integrated Nanotechnologies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States.

³Center for Integrated Nanotechnologies, Sandia National Laboratory, Albuquerque, New Mexico 87185, United States.

*Corresponding author(s). E-mail(s): aneogi2@lanl.gov; serg@lanl.gov; acj@lanl.gov; jxzhu@lanl.gov;

Contributing authors: mishras@lanl.gov;

†*These authors contributed equally to this work.

Abstract

Increasingly large datasets of microscopic images with nanoscale resolution facilitate the development of machine learning methods to identify and analyze subtle physical phenomena embedded within the images. In this work, microscopic images of honeycomb lattice spin-ice samples serve as datasets from which we automate the calculation of net magnetic moments and directional orientations of spin-ice configurations. In the first stage of our workflow, machine learning models are trained to accurately predict magnetic moments and directions within spin-ice structures. Variational Autoencoders (VAEs), an emergent unsupervised deep learning technique, are employed to generate high-quality synthetic magnetic force microscopy (MFM) images and extract latent feature representations, thereby reducing experimental and segmentation errors. The second stage of proposed methodology enables precise identification and prediction of frustrated vertices and nanomagnetic segments, effectively correlating structural and functional aspects of microscopic images. This facilitates the design of optimized

spin-ice configurations with controlled frustration patterns, enabling potential on-demand synthesis.

Keywords: Variational AutoEncoder, Generative Machine Learning, Artificial Spin-Ice, Magnetic Force Microscope

1 Introduction

Artificial spin ice (ASI) systems are engineered arrays of nanoscale magnetic elements designed to mimic the frustrated spin configurations of natural spin-ice materials [1]. These patterned nanomagnet lattices provide a versatile platform for exploring fundamental magnetic phenomena that were previously observed only in bulk spin systems or theoretical models [2],[3],[4]. In particular, ASI exhibits signatures of magnetic frustration, characterized by a large degeneracy of low-energy configurations and the emergence of excitations resembling magnetic monopoles, all governed by tunable geometric constraints [2],[3]. Studying these effects in ASI has broadened our understanding of frustration and emergent magnetism under controllable settings, bridging the gap between theoretical models and real materials. Beyond fundamental physics, ASI has promising applications as a platform for new magnetic and spintronic technologies. By tailoring the geometry and coupling of its nanomagnets, an ASI array can function as a reprogrammable magnonic crystal with tunable spin-wave spectra for microwave signal processing and quantum information devices [1],[5],[6]. ASI-based networks have also been proposed for unconventional computing architectures. For example, specific magnetically frustrated nanomagnet arrangements can be used for logic gate implementations or neuromorphic computing [7],[8]. Moreover, the ability to manipulate ASI's emergent magnetic monopole defects using external magnetic fields, electric currents, or thermal gradients [9],[10],[11] paves the way for in-situ control of magnetic charge transport, enabling novel active metamaterials and next-generation functional devices.

Magnetic Force Microscopy is a powerful technique for exploring and characterizing complex magnetization patterns in ASI systems. MFM is a variant of atomic force microscopy that uses a magnetized probe tip to scan above the sample surface and sense the stray magnetic fields produced by individual nanomagnets. This method provides high-resolution, non-destructive imaging of local magnetic configurations in an ASI lattice [12]. The fabrication method of ASI sample used for magnetic phase measurement is explained in section S1. Unlike bulk magnetometry methods that measure only a sample's aggregate response, MFM directly visualizes the magnetization of individual bars and the state of each vertex, enabling real-space analysis of spin interactions at the nanoscale. The technique's strong sensitivity to nanoscale stray fields makes it particularly effective for capturing magnetization reversal processes, thermally activated excitations, and field-induced phase transitions in ASI systems [13]. Figure 1 illustrates an MFM measurement on an ASI sample. In Fig. 1a, schematic of a variable external B field MFM setup shows a magnetized tip scanning over the patterned nanomagnet array to detect local magnetic forces. Fig. 1b and 1c exemplify

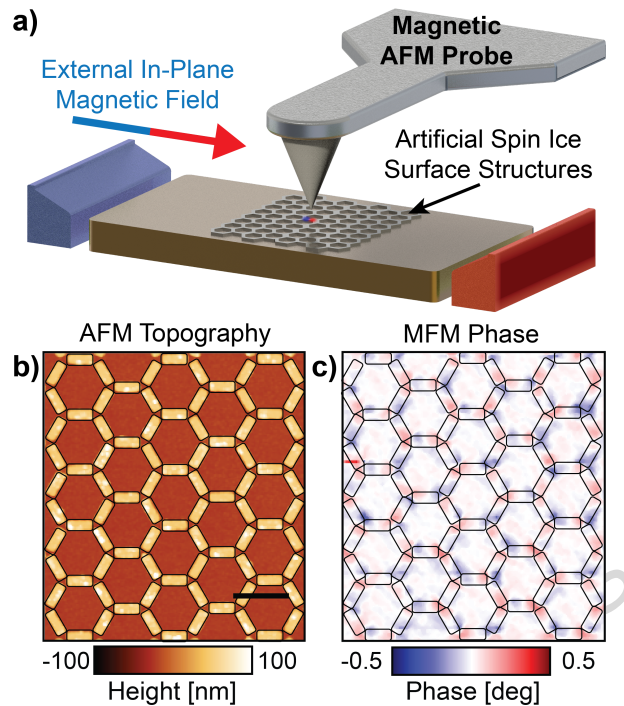


Fig. 1 Magnetic AFM setup and representative AFM topography and MFM phase images of the artificial spin ice sample. (a) A schematic representation of magnetic-AFM illustrating the experimental setup for ASI sample measurement, along with its corresponding (b) morphology with the scale bar of $1\mu\text{m}$. and (c) magnetic phase image. The red and blue colors in the phase image represent the North and South poles of the magnetic bars, respectively. Outlines of the ASI elements determined from the AFM topography are depicted by the black lines in both the AFM topography and MFM images (b,c) as a guide to the eye.

the two complementary outputs of the measurement: the sample's surface morphology (topographic image) and the corresponding magnetic phase image, respectively. In the MFM phase image, each nanobar's stray field produces a bipolar contrast with red and blue regions indicating the north and south magnetic poles at opposite ends of the bar. This contrast explicitly reveals the magnetization direction of every element. This direct imaging of nanoscale dipole orientations and polarities enables immediate identification of the magnetic state of each nanomagnet in the ASI lattice. Continued improvements in MFM instrumentation have further enhanced its ability to probe ASI. Modern MFM probes with higher spatial resolution and lower magnetic moment coatings minimize tip-induced perturbations while increasing image contrast, now resolving magnetic features as small as 10 nm [14]. Additionally, the integration of variable in-plane magnetic field sample holders into MFM systems enables the real-time observation of ASI magnetization dynamics on time-scales longer than that of the MFM scanned frame-rate which is typically in the order of several minutes. This time resolution is sufficient to capture slow magnetic field-driven phase transitions in-situ while performing MFM imaging. These advances have accelerated ASI research

by allowing in-depth exploration of various topological or frustration-induced configurations and by opening the door to spin-logic experiments under realistic conditions [7]. As MFM continues to evolve, it remains a critical tool for understanding and harnessing the unique properties of ASI systems for future technological applications. While MFM provides rich visual data on ASI, analyzing the resulting images to quantify spin configurations can be challenging and labor-intensive. The increasing volume of high-resolution microscopy data has motivated the use of machine learning (ML) techniques in condensed matter physics [15][16] to automate image analysis and feature extraction. In the context of ASI, however, traditional image-processing methods often rely on heuristic thresholds or manual identification of nanomagnets [17], which can introduce errors due to noise, instrumental artifacts, and user bias. This limitation creates a need for robust data-driven approaches capable of reliably extracting meaningful physical information, such as each nanomagnet’s moment orientation and the energy or “frustration” state of each vertex from MFM images. Such automated analysis is particularly important given the highly complex, correlated magnetization patterns that ASI exhibits [18], which can be difficult to interpret with conventional techniques.

In this work, we address these challenges by integrating MFM imaging with an unsupervised deep learning approach for quantitative analysis of spin-ice configurations. We employ a VAE to learn latent feature representations from MFM phase images of a honeycomb ASI lattice. This data-driven approach enables automated identification of each nanomagnet’s magnetic moment direction and the detection of frustrated which are high-energy vertex states directly from the images. The VAE model effectively captures the high-dimensional correlations in the ASI’s magnetic phase patterns while reducing experimental noise and segmentation errors. Furthermore, the learned latent space can be used to generate synthetic MFM images that replicate experimental features, offering deeper insight into how frustration manifests and can be controlled in these networks. By capturing subtle signatures of emergent frustration that might be missed by manual analysis, our framework provides a powerful tool for discovering and designing optimized spin-ice states with tailored properties, addressing the highly tunable interactions and reprogrammable magnetism of ASI systems [19]. Ultimately, by combining machine learning with advanced microscopy, this approach provides a scalable and precise platform for characterizing ASI systems, leading to on-demand “frustration engineering” in magnetic metamaterials [20] and advancing spin-based information processing technologies.

Our overall workflow for image analysis shown in Fig. 2 comprises two distinct components. The first component of our framework processes raw MFM data to extract each nanomagnet’s magnetic moment. We begin by segmenting the ASI sample’s images to isolate individual nanomagnets as distinct regions. In practice, the high-contrast AFM morphology image of the sample, such as shown in Fig. 3(a), is first converted to grayscale and thresholded to obtain an initial mask of all nanomagnets. This mask is then refined using a Euclidean distance transform and a watershed algorithm, which sharpen boundaries and separate any touching or overlapping nanomagnets, as described in Methods Section 4.2. The outcome is a segmentation of the morphology into uniquely labeled nanomagnet regions referred in Fig. 3(b).

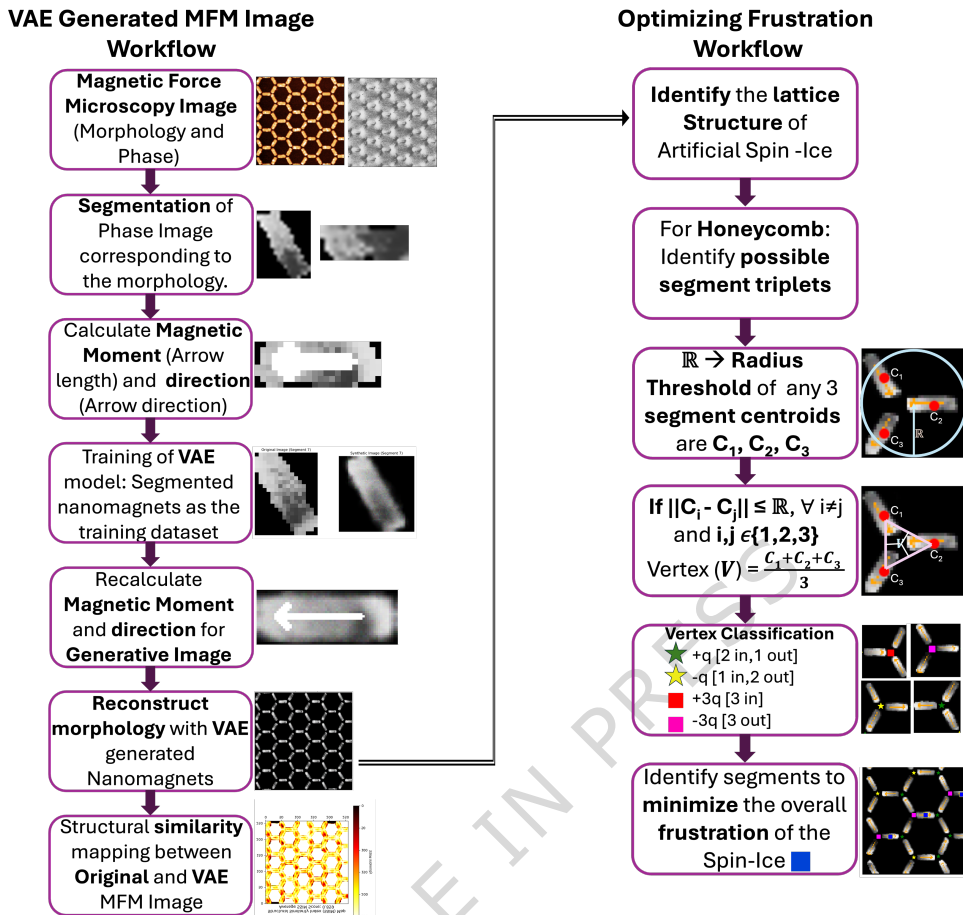


Fig. 2 Workflow for VAE generated MFM reconstruction and frustration optimization in the artificial spin-ice lattice. The scheme of the overall workflow comprising two distinct components. First, we apply a segmentation algorithm to partition the MFM image into distinct magnetic domains. From the original MFM segments, we train a VAE model to generate synthetic MFM segments. In this representation, the direction of each arrow indicates the nano-magnet dipole direction, while the arrow length corresponds to the normalized magnetic moment of the dipole. The segment position and orientation are retrieved from the original MFM image and reconstructed. The second stage of our workflow focuses on classifying the vertices of a honeycomb lattice spin-ice sample into high energy or highly frustrated ($\pm 3q$) and low energy or less frustrated ($\pm q$) vertices. Furthermore, our algorithm not only identifies the high energy monopoles of the spin-ice lattice but also determines which segments, if toggled, would minimize the overall frustration.

This automated segmentation ensures each bar's magnetic signal is cleanly isolated with accurate boundaries, providing the foundation for all subsequent analysis. Next, we map these segmented regions onto the corresponding MFM phase image to extract the magnetic signal of each nanomagnet. The MFM phase channel in Fig. 3(c), captures the bipolar contrast of each nanobar's stray field, with a bright lobe




Fig. 3 Morphology segmentation and mapping onto the MFM phase image to extract individual nanomagnet segments. (a) Morphology of the ASI system scanned using MFM (b) Uniquely colored segmentation of the morphology capturing the shape and position of each segments. (c) Gray-scale MFM phase image. (d) Uniquely colored morphological segmentation image overlaid on the gray-scale MFM image. (e) Examples of segments from the MFM image segmented based on the overlay morphological image representing nano-magnets contributing to the overall layout of the artificial spin ice sample.

at one end and a dark lobe at the other end indicating the north and south magnetic poles, respectively. By overlaying the colored segmentation mask on the grayscale MFM phase image as in Fig. 3(d), we align each nanomagnet's physical outline with

its magnetic contrast pattern. Each segmented region now cleanly corresponds to a single nanomagnet’s phase signal. From this overlay, individual MFM segments are extracted for analysis, as illustrated in Fig. 3(e), which shows examples of 64x64 pixel cutouts of one per nanomagnet. These isolated segment images serve as inputs to our deep learning pipeline. We record the original position and label of each segment via connected-component analysis so that the lattice can be reconstructed after processing. This preparation step yields a set of well-defined nanomagnet images and ensures that the magnetic signal of each nanobar is considered independently during subsequent analysis.

For each segmented nanomagnet, we then compute its magnetic moment vector, comprising the in-plane orientation and a relative magnitude, directly from the MFM phase contrast. We first estimate the long-axis orientation of a nanomagnet by fitting an ellipse to its segmented shape; the ellipse’s major axis provides an initial guess for the nanomagnet’s dipole alignment. However, establishing which end of the nanomagnet corresponds to the north pole, i.e. the direction of the dipole moment, requires more than simply comparing the average brightness of the two halves of the MFM segment. To robustly determine the polarity, our algorithm goes beyond a mean brightness test and employs a radial brightness gradient analysis as discussed in Methods Section 4.2.

In this radial analysis, each nanomagnet segment is virtually split along its major axis, the presumed dipole axis, and we scan the MFM phase intensity outward from the center toward both ends of the nanomagnet. All segmented nanomagnets from the example in Fig. 3(e) are reassembled in their original lattice positions. One particular nanomagnet, highlighted and enlarged in the right panel of Fig. 4(a), exhibits an anomalous bright spot due to a local imaging artifact, which can confound naive brightness-based methods. Our algorithm addresses this by examining the gradient of the phase signal along radial lines: a proper dipole pattern shows a steadily increasing brightness gradient toward the true north end, the bright pole and a decreasing gradient toward the south end, the dark pole. In the highlighted segment of Fig. 4(a), the measured brightness gradient along one half of the nanomagnet initially decreases which is expected for the south end but then experiences a sudden increase beyond a certain point. This abrupt reversal, marked by a jump above a set threshold, signals the presence of an artifact or “adversarial” feature in the image. The algorithm automatically flags the region beyond this point as unreliable and excludes that portion of the segment from the moment calculation. By filtering out such spurious contrast deviations, we ensure that only the genuine bipolar contrast of the nanomagnet is used to determine its magnetic moment.

The remaining valid contrast profile after artifact removal allows us to precisely determine which end of the nanomagnet is the north pole. In the example segment of Fig. 4(a), once the aberrant region is ignored, the brightness consistently increases toward one end, correctly identifying that end as the magnet’s north pole. Fig. 4(b) validates the result: the arrow overlaid on the zoomed segment points toward the true north end (bright lobe), neglecting the misleading artifact. Through this radial gradient method, our framework obtains a robust determination of each nanomagnet’s dipole orientation despite the presence of minor image artifacts or noise. We assign an arrow to each segment pointing from the south to north pole, and scale the arrow’s

length in proportion to the contrast difference between the two halves of the segment. The arrow direction and length thus represent the inferred magnetic moment orientation and a normalized strength for each nanomagnet. At this stage, we have a set of segmented nanomagnets, each labeled with an arrow indicating its magnetization direction and relative magnitude.

Using this information, we next train a Variational Autoencoder to generate refined, synthetic MFM images of the nanomagnets. The VAE is an unsupervised generative model that learns a latent feature representation of individual nanomagnet phase profiles as described in Methods Section 4.2. In our implementation, each 64x64 pixel nanomagnet segment is fed into an encoder network that compresses it into a 60 dimensional latent vector, and then a decoder network reconstructs the image from this latent code. By training on the ensemble of experimental segments, the VAE learns to reconstruct nanomagnets with realistic bipolar patterns while effectively denoising the data and correcting spurious irregularities. Once trained, the decoder can produce a clean synthetic version of each nanomagnet segment given its original image as input. We pass all segmented nanomagnet images through the VAE, obtaining a set of synthetic MFM segments that closely resemble the experimental ones but with improved consistency and reduced noise.

Each synthetic segment is then placed back into its original spatial location within the lattice using the recorded coordinates from the initial segmentation. We apply the same radial brightness-gradient procedure to each synthetic segment to determine its dipole orientation, ensuring consistency with the method used on the raw data. The result is a reconstructed MFM image of the entire ASI sample, where every nanomagnet is now represented by a VAE-generated, denoised phase image and is annotated with a confident magnetic moment vector (arrow direction and length). This VAE-driven reconstruction serves two important purposes: (a) it reduces the influence of noise or ambiguous contrast in the raw images, and (b) it provides a more reliable basis for identifying each nanomagnet's moment. In essence, the left branch of our workflow of Fig. 2, generates an accurate "magnetic moment map" for the lattice, derived through a combination of precise image segmentation and generative modeling.

With a reliable map of each nanomagnet's moment, the second stage of our pipeline focuses on analyzing and optimizing magnetic frustration in the lattice shown in Fig. 2, right column. Using the set of arrow-labeled nanomagnets from the first stage, we systematically identify all vertex interaction sites in the honeycomb ASI geometry described in Methods Section 4.2. Each vertex is the junction where three neighboring nanomagnets meet. To find these, we compute the centroids of all segmented nanomagnets and search for triplets of centroids that fall within a predefined distance of one another (on the order of the lattice spacing). Every such triplet is registered as a vertex, with the vertex location defined as the average of the three segment centroids. Once the vertices are identified, we determine the magnetic "charge" or energy state of each vertex by examining the orientations of its three constituent moments. In a trivalent spin-ice lattice, a vertex is frustrated or high energy if all three nanomagnet moments point either inward toward the vertex or outward away from it, a "3-in" or "3-out" configuration. This corresponds to a net magnetic charge of $\pm 3q$ at the vertex, essentially a monopole-like excitation. Conversely, a satisfied or low energy




Fig. 4 Reconstruction of segmented nanomagnets and identification of artifacts during magnetic moment determination. (a) The segmented nano-magnets from Figure 3(e) are reconstructed with position of each segment retrieved from the labeled segment data of the morphology of spin-ice sample. The right side is a zoomed in nano-magnet with artifact, classified as an adversary attack on the expected dipole image. An expected dipole nano-magnet image should have a bright(North) and a dark end (South). The segment is subdivided into two halves, each half is radially scanned calculating brightness gradient. The gradient decreases along the South end of the nano-magnet while increases along the North end. Reaching the artifact radial area, the decreasing brightness gradient experiences sudden increase ($>$ threshold contrast) and classifies the rest of the segment as an adversary. The erroneous part is not included in calculating the magnetic moment calculation. (b) Validating the calculation process, the arrow of the zoomed in segment points towards the North neglecting the adversary.

vertex has two moments pointing in and one out, or two out and one in, a “2-in/1-out” or “2-out/1-in” configuration, resulting a net charge of $\pm q$. We classify each identified vertex by comparing the directions of its three arrows, thus mapping it to either a high energy monopole ($\pm 3q$) state or a low energy dipole ($\pm q$) state. In this way, the framework automatically produces a frustration map of the entire lattice: all high energy vertices are flagged as sites of emergent monopoles, whereas $\pm q$ vertices indicate locally satisfied spin-ice ordering. Visualization of this classification on two experimental dataset is shown in Figure 5(a and b).

Beyond detecting frustration, our approach also optimizes the spin configuration by suggesting targeted moment reversals to alleviate high-energy states. In the final stage of the workflow, the algorithm identifies which specific nanomagnet flips would convert frustrated vertices into lower-energy configurations. For each nanomagnet that participates in one or more $\pm 3q$ (monopole) vertices, we virtually “toggle” its magnetic moment i.e., flip its arrow direction and then recompute the vertex charges in its neighborhood. If a particular flip eliminates one or more 3-in/3-out vertices thus reducing the total count of frustrated vertices without introducing new frustration elsewhere, that flip is marked as beneficial and is provisionally adopted. We iterate this

procedure, testing different candidate flips in sequence while keeping any previously accepted flips fixed. This greedy optimization strategy prioritizes flips that relieve multiple frustration points simultaneously and avoids undoing earlier improvements. After systematically exploring the candidates, the end result is an adjusted ASI configuration with a minimized number of high-energy vertices. In Figure 2 (right column), this process is depicted by highlighting the nanomagnets selected for flipping and showing the consequent disappearance of monopole defects in the frustration map. The entire frustration optimization routine thus provides a data-driven means not only to diagnose where frustration is present, but also to prescribe how the sample’s magnetization should be reconfigured to achieve a lower-energy state. By combining the VAE informed moment identification with a n i t e r a t i v e frustration minimization strategy, our framework offers a blueprint for “frustration engineering” by actively tuning the magnetic state of the ASI lattice to control and reduce or even deliberately increase magnetic frustration in a systematic way. This two-part machine learning approach enables precise, high-throughput analysis of spin-ice microstates and provides guidance for experimentally realizing more stable or selectively frustrated artificial spin ice configurations.

2 Results

2.1 Variational Autoencoder Reconstruction

To this end we have shown that the Variational Autoencoder model can accurately identify the magnetic moments of individual nanomagnet segments in the artificial spin ice lattice. By harnessing its ability to extract high-dimensional latent features and reconstruct the Magnetic Force Microscopy images, the VAE enables precise determination of dipole orientations. This, in turn, allows for a systematic frustration analysis of the ASI system, classifying high-energy and low-energy vertex states based on their magnetic configurations. The VAE model is trained on segmented nanomagnet structures extracted from MFM images, learning a compressed latent space representation of these structures. Once trained, the VAE reconstructs synthetic MFM images, preserving essential magnetic features while mitigating noise and experimental artifacts.

Fig. 5(a) displays MFM images recorded under an external magnetic field of 90 mT, applied from right to left. The applied field enforces a preferential alignment of nanomagnet moments, influencing the initial spin-ice configuration. The second row in Fig. 5(b) shows the relaxed ASI lattice, 15 mins after applying magnetic field of -25mT, where negative sign corresponds to the opposite direction of initial magnetic field condition, for relaxation period. During this time, spontaneous moment flips occur, allowing the system to transition toward a lower-energy state. In both cases, the second column presents the segmented nanomagnets obtained from the morphology image, where each segment is labeled, and arrows indicate their magnetic dipole moment directions. The third column in Fig. 5 presents the VAE-generated synthetic MFM images, where individual nanomagnets are reconstructed and assigned dipole orientations based on the trained model. The latent space representation learned by the VAE allows for smooth reconstructions that are robust to noise and minor artifacts,




Fig. 5 Progression from experimental MFM images to VAE assisted moment maps, vertex classification, and frustration minimization by selected flips. The first row (a) presents MFM images recorded under an external magnetic field of 90 mT, applied from right to left. The second row (b) corresponds to the relaxed state of the ASI lattice after relaxation of 15 mins, allowing spontaneous flipping of nanomagnet moments due to thermal activation. The second column displays the reconstructed MFM images, pertaining to the corresponding morphology image, where individual nanomagnets are segmented and labeled, with arrows indicating the dipole moment direction. The third column shows the synthetic nanomagnet reconstructions generated training VAE model. The dipole moment directions are calculated following the method outlined in Section 4.2. The highlighted segments in the original MFM images compared with the synthetic segments, reveals discrepancies in the top row, where miscalculations occur in moment direction, while the dipoles in the relaxed bottom row remain inconclusive. The vertices of the honeycomb lattice are classified into high-energy monopoles and low-energy dipoles: Frustrated monopoles are further categorized into $+3q$ (red squares) and $-3q$ (magenta squares). Lower-energy dipoles are classified as $+q$ (green stars) and $-q$ (yellow stars). The nanomagnet segments contributing to high-energy monopoles are identified, and the frustration of the ASI lattice is iteratively evaluated. Finally, nanomagnets to be toggled for frustration minimization are detected to optimize the ASI configuration in (c) using greedy optimization algorithm described in Section 2.2.

serving as a method to determine magnetic moment directions with higher accuracy. Fig. 5(a) emphasizes discrepancies between the original and VAE-generated images.

The model identifies cases where moment direction assignments are incorrect in the experimental MFM images when using direct image segmentation, highlighting the advantage of learned representations. After relaxation of the ASI lattice, dipole orientations remain ambiguous in certain regions. The VAE enables a systematic approach to analyzing these fluctuations by reconstructing the most probable moment configurations given the observed data. Thus the VAE generated ASI system accomplishes two key purposes: first, it addresses the uncertain magnetic configurations present in the original MFM images, and second, it identifies inconclusive dipole moments within the microscopic data. The algorithm identifies inconclusive dipole moments of the microscopic images. As a consequential advantage, the synthetic MFM images aids in a more precise classification of lattice vertices.

2.2 Greedy Optimization of Frustration of Spin Ice Lattice

The process of frustration minimization involves identifying and selectively toggling the magnetization direction of nanomagnets in the ASI lattice to achieve a lower-energy state while maintaining the overall structure of the spin ice. This is a two stage methodology consisting of determining magnetic moments of the segments to be toggled (flipped from 0 to 1 or vice versa) and evaluating the impact of toggling on the overall frustration state.

Algorithm:

The objective of the algorithm is that given (i) a set of vertices with their associated nanomagnet segments and (ii) an experimentally inferred dipole “Direction” for each segment, we seek a minimal set of segment toggles i.e., magnetic moment direction flips, that reduces the total number of frustrated vertices as determined by the vertex classification rule described in Section 4.2.

Inputs

- Vertices: Table containing each vertex ID and its associated segments which is three segments per vertex for honeycomb lattice, and vertex coordinates (x,y) .
- Magnetic Moments: Table containing each segment ID and its inferred dipole Direction and other metadata.
- Synthetic Segment Positions: Table containing VAE generated segment centroids (x,y) .
- Segment to Vertex Mapping: Dictionary mapping each segment to the one or two vertices that share it.
- Classify detailed Frustration(v , segments): Function that returns the frustration state of vertex v using the current segment directions and the geometric rule including the “horizontal” segment identification from segment orientation.

Outputs

- Updated Magnetic moments table with a boolean flag Toggle indicating whether a segment direction was flipped.
- Toggle Tracker: Audit table listing each accepted toggle listing vertex ID, segment ID, role, original and new binary directions.

Step 0: Preprocessing.

1. Load the vertex table, the segment dipole-direction table, and the segment centroid position table.
2. Normalize segment identifiers by extracting numerical segment IDs from filenames where necessary.
3. Construct the following mappings:
 - `segment_directions` : Segment \rightarrow Direction, storing the current dipole direction of each segment.
 - `segment_positions` : Segment $\rightarrow (x, y)$, storing the centroid coordinates of each segment.
 - `segment_to_vertex_mapping` : Segment \rightarrow [Vertex IDs], obtained by iterating over the associated segments of each vertex.
4. Initialize a boolean flag `Toggle = False` for all segments.

Step 1: Identification of frustrated vertices.

1. Load the initial vertex frustration analysis table, or compute it by evaluating the vertex classification function for all vertices.
2. Define the set of initially frustrated vertices as

$$\mathcal{F}_0 = \{v \mid \text{Frustration_Status}(v) = \text{"Frustrated"}\}. \quad (1)$$

3. Let the initial global frustration count be

$$N_0 = |\mathcal{F}_0|. \quad (2)$$

Step 2: Greedy frustration minimization (local search).

For each vertex $v \in \mathcal{F}_0$, perform the following steps:

1. Determine candidate segments for toggling.

- Extract the three associated segments of vertex v :

$$S_v = \{s_1, s_2, s_3\}. \quad (3)$$

- Identify the horizontal segment s_h by fitting an ellipse to each segment mask and selecting the segment whose orientation angle lies within a fixed tolerance of 0° or 90° .
- Label the remaining two segments as angled segments s_{a1} and s_{a2} .

2. Attempt single-segment toggles and evaluate global improvement. Consider toggling the segments in the fixed order

$$[s_h, s_{a1}, s_{a2}]. \quad (4)$$

For each candidate segment s (with $\text{role} \in \{\text{Horizontal}, \text{Angled}_1, \text{Angled}_2\}$):

- (a) Store the current dipole direction

$$d = \text{Direction}(s). \quad (5)$$

- (b) Apply a tentative toggle

$$d' = 1 - d, \quad (6)$$

and update both the segment-direction table and the global direction mapping.

- (c) **Recompute the global frustration count** by evaluating the frustration classifier for all vertices:

$$N' = \sum_{u \in \mathcal{V}} \mathbf{1} \{ \text{classify_detailed_frustration}(u) \text{ contains "Frustrated"} \}. \quad (7)$$

- (d) **Accept or reject the toggle.**

- If $N' < N_0$, accept the toggle:
 - Set **Toggle** = True for segment s .
 - Append a record to the toggle-tracking table containing the vertex ID, segment ID, segment role, original direction d , and new direction d' .
 - Update the global frustration count $N_0 \leftarrow N'$.
 - Stop further testing for vertex v and proceed to the next vertex.
- Otherwise, reject the toggle and revert the dipole direction $d' \rightarrow d$ in both the table and the global mapping.

The procedure terminates once all vertices in the initially frustrated set \mathcal{F}_0 have been processed.

2.3 Selecting Candidate Segments for Toggling

The selection process prioritizes the segment that has the highest influence on multiple frustrated vertices. Each segment is classified as either horizontal or angled based on its orientation in the lattice. Horizontal segments (aligned with 0° or 90°) are treated as reference points for toggling decisions. Segments associated with the highest number of frustrated vertices are given priority for toggling. If a single segment contributes to frustration in two or more vertices, toggling it may result in a greater reduction in overall frustration. The spatial arrangement of segments is another criteria of the selection process where the each segment is analyzed to ensure that toggling one segment does not introduce new frustration states elsewhere in the lattice. For each frustrated vertex, the selected segment's magnetic moment is toggled, based on the binary classification of magnetic moments as explained in Methods section 4.2 and the frustration classification of all affected vertices is re-evaluated. If toggling the segment results in a net decrease in frustration across the lattice, the change is accepted; otherwise, the segment's original state is restored.

In the toggling process, the first step is the extraction of original magnetic moment direction of the segment and temporarily toggling the segment's direction. Then we recalculate the frustration classification for all affected vertices and compare the new frustration count with the previous count. If the frustration count decreases, the new magnetic moment direction is retained; otherwise, the segment is reverted to

its original orientation. All toggled segments are stored in a tracking table to ensure that subsequent toggles do not interfere with previously optimized regions.

2.4 Evaluating the Effectiveness of Toggling

The effectiveness of the toggling strategy is measured by comparing the frustration count before and after the optimization process. The final optimized ASI configuration is saved, along with a record of which segments were toggled. The final frustration map is visualized by classifying vertices into their respective frustration categories. The segments to be toggled are marked with blue squares in Fig. 5(c) that reduces the overall frustration, generating a low energy state ASI system. By iteratively adjusting the nanomagnet orientations, the VAE-generated ASI lattice is refined into a more stable configuration, minimizing the overall frustration in the system. This allows precise control over frustration states in artificial spin ice systems, providing a template for generating on-demand spin ice sample with the ability to reduce energy or if required, introduce high energy vertices at specific locations. The frustration minimization procedure implemented here is intentionally designed as a greedy, image-driven optimization algorithm that operates directly on experimentally inferred magnetic configurations rather than on simulated or thermally equilibrated spin states. At each step, only single-segment toggles are considered, and a toggle is accepted exclusively when it produces a global reduction in the number of frustrated vertices, ensuring that improvements are assessed at the lattice level rather than locally at an individual vertex. This conservative strategy preserves interpretability and allows each accepted toggle to be directly associated with a specific reduction in frustration. While the greedy nature of the algorithm does not guarantee convergence to a globally optimal minimum-frustration state, it provides a deterministic, transparent, and experimentally actionable framework well suited for post-characterization analysis of MFM images as illustrated in Fig. 5(c).

3 Discussion

The integration of VAEs with MFM imaging has enabled a novel and robust approach for analyzing frustration in artificial spin ice systems. Traditional methodologies for determining magnetic moment orientations and frustration states often face limitations due to experimental noise, segmentation inaccuracies, and uncertainties in moment determination. The VAE-based generative model addresses these challenges by providing a data-driven framework capable of reconstructing high-fidelity synthetic MFM images, thereby allowing for a more precise identification of frustration vertices and their associated nanomagnetic configurations.

The primary advantage of employing VAEs lies in their ability to extract latent feature representations from experimental MFM data, effectively mitigating segmentation errors and compensating for inconsistencies in the direct image-based calculations. By training on segmented MFM images, the VAE learns to encode essential spatial and magnetic features, facilitating improved prediction of moment orientations

even in cases where the original data is ambiguous or partially corrupted. The reconstructed synthetic images retain critical structural and magnetic information while reducing noise-induced distortions, allowing for a systematic and automated approach to frustration classification. The capability of VAEs to generalize from learned features corrects magnetic moment miscalculations present in raw experimental images, as observed in the comparison between original and synthetic reconstructions. This approach is particularly useful in identifying high-energy monopole configurations ($\pm 3q$) and lower-energy dipole states ($\pm q$), enabling deeper understanding of frustration dynamics within the ASI lattice. Furthermore, the automated classification of frustration states assists in the optimization of ASI configurations through targeted toggling of nanomagnet orientations. By iteratively evaluating the impact of flipping specific nanomagnets, the framework efficiently identifies an optimal set of toggled segments that minimizes overall lattice frustration. This process provides a blueprint for designing spin-ice configurations with tailored frustration properties, offering a pathway toward controlled magnetic metamaterials with tunable energy states.

In this work, we have demonstrated that deep learning-based generative models can significantly improve the accuracy and robustness of frustration classification in ASI systems. The ability to synthesize high-quality MFM images, correct for segmentation errors, and systematically classify magnetic moment orientations highlights the power of machine learning in advancing experimental condensed matter physics. The automated frustration analysis and optimization framework presented here lays the groundwork for next-generation artificial spin ice designs with precise control over magnetic configurations, contributing to the broader field of engineered magnetic materials. The findings of this study have broad implications for the design and control of ASI-based systems in various applications, including reconfigurable magnonics, spintronics, and quantum information processing. The ability to systematically predict and minimize frustration through deep learning approaches enables the development of programmable ASI structures with precise energy landscapes.

The present approach also has limitations that should be noted explicitly. Magnetic Force Microscopy is a scanning probe technique and its time resolution is limited by the scan time, typically several minutes, so it images steady state conditions and can only track slow field driven evolution rather than fast non equilibrium dynamics. The VAE does not generate new spatial information beyond that contained in the experimental data, so the synthetic outputs should not be interpreted as true super resolution, even if they appear smoother due to noise and artifact suppression. Because direct ground truth for individual nanomagnet moments is not experimentally accessible in these systems, the VAE reconstructions and inferred dipole assignments are validated through expert guided visual inspection, internal consistency checks, quantitative image fidelity metrics such as MSE and SSIM as shown in Fig. S2, and comparison between direct segmentation based and VAE assisted assignments, so the method provides probabilistic refinement rather than definitive ground truth recovery. A direct quantitative benchmark against linear dimensionality reduction approaches such as PCA is not included here and is left for future work. Our future work will also explore expanding the VAE framework to incorporate multi-scale analysis, allowing for frustration predictions in larger and more complex ASI lattices. Additionally, integrating reinforcement

learning techniques with the existing toggling strategy could further enhance the efficiency of frustration minimization by optimizing flipping sequences dynamically. Experimental validation of the proposed synthetic reconstructions through direct magnetic imaging and comparison with theoretical models may provide further insights into the reliability and accuracy of the machine learning approach.

4 Methods

4.1 Magnetic Force Microscopy Measurements

Nanoscale topography and MFM phase shift images were collected using an Atomic Force Microscope system (DriveAFM, Nanosurf AG). Measurements were performed using ‘lift mode’ to record phase contrast variations associated with the magnetic phase of the ASI sample surface. First, a measurement of the topography of the sample surface was performed using tapping-mode feedback of the AFM probe via a line trace and retrace scan. Next, to reduce the effect of magnetic phase artifacts associated with the sample topography, the probe is then separated by a controlled distance from the surface (~ 70 nm), and a subsequent, elevated, trace, and retrace line scan is performed while recording the phase shift of the oscillating AFM probe. The external, in-plane magnetic field was produced using a variable magnetic field sample holder (VMFSH, NanoSurf) with a base that contains the permanent magnets and an integrated calibrated Hall sensor. Magnetic atomic force microscopy probes were used (Co/Cr coated Bruker MESP-LM-V2) with a spring constant of 2.2 N/m, and resonance frequency of ~ 65 kHz. These probes were magnetized in an out-of-plane direction using a permanent magnet prior to measurements.

All MFM images analyzed in this work were acquired with a well-defined in-plane magnetic field applied parallel to the image plane and oriented from right to left with respect to the displayed image frame. This orientation is used consistently throughout the manuscript. For the dataset shown in Fig. 5(a), the ASI sample was imaged under an applied in-plane field of 90 mT directed from right to left. For the relaxed configuration shown in Fig. 5(b), the applied field was subsequently reversed to -25 mT, corresponding to the opposite field polarity along the same right-to-left axis and the system was allowed to relax for 15 minutes prior to imaging. To avoid ambiguity, the image orientation is fixed across all figures and the applied field direction is explicitly indicated where relevant.

4.2 Automating Net Magnetization Calculation

To automate the calculation and analysis of net magnetization in ASI systems, we initiate an image processing pipeline that extracts and uniquely labels individual nanomagnetic segments from the MFM morphology image. The accurate segmentation forms the foundation for computing net magnetic moments in subsequent stages of the workflow.

Segmentation of Nanomagnets

The input morphology image (Fig. 3(a)) is first converted to grayscale, followed by the application of a binary thresholding mask to distinguish magnetic regions from the

background. This step ensures that only relevant segments are considered. Euclidean distance transform refines the segmentation by enhancing the separation between foreground structures and suppressing noise, thereby improving segmentation accuracy. It converts the binary mask into a grayscale map, where brighter intensities represent regions farther from the boundary keeping only high-intensity regions, ensuring robust object detection.

To identify individual nanomagnetic segments, Connected Components Analysis (CCA) from the OpenCV2 computer vision library [21] is employed. This step assigns unique labels to distinct magnetic structures. However, in cases where segments are closely packed, CCA alone may not effectively separate overlapping structures. To address this, the Watershed Segmentation Algorithm [22] is applied to refine boundaries further. The watershed algorithm uses gradient information to distinguish touching objects by identifying high-gradient regions as boundaries, ensuring each nanomagnet is distinctly labeled and effectively separated.

Following segmentation, contours are extracted for each labeled segment using OpenCV's "findContours()" function which identifies the outer boundary of each detected region by tracing along intensity changes in the binary mask, where each segmented nanomagnet has a unique label. OpenCV's "cv2.findContours()" is applied to this mask with external retrieval mode ("cv2.RETR_EXTERNAL") to detect only the outermost contour, ignoring any nested inner contours. The simplified chain approximation reduces redundant points and store only the essential boundary points of each segment.

Only segments with an area exceeding a predefined threshold are retained to filter out small, incomplete, or boundary-clipped structures. These selected segments are then assigned unique colors using an HSV colormap, enhancing visual differentiation. The center coordinates of each segment are determined using the mean position of the contour points, and each segment is numerically labeled for reference and tracking. Finally, the segmentation map is compared with the original morphology image to validate the accuracy of the detected structures.

The individual segments are reassembled for full segmentation visualization, which ensures that the extracted segments are correctly placed within the original morphology image. The individual segmented nanomagnets, stored as separate image files, are sorted and iteratively placed back into the original layout. To correctly reposition each segment, the corresponding label position is retrieved from the segmentation mask mapping segments to their original positions.

The MFM image with magnetic contrast data, is segmented similarly using the reconstructed segmentation map, such that each nanomagnetic region is extracted with accurate boundaries. This method aligns the spatial regions in the MFM image with the segmented morphology image, allowing further analysis of magnetic properties at the nanoscale and we validate the segmentation process by comparing the extracted segment count with the number of unique labels. This step ensures that the MFM image is segmented accurately, preserving the spatial and morphological characteristics of each nanomagnet for further magnetic analysis.

Magnetic Moment Calculation

Next we calculate the magnetic moment direction and strength of each nanomagnet from the segmented MFM image which involves segment extraction, brightness-based dipole determination, and arrow visualization to indicate the moment direction. The segment image is converted to grayscale, and binary thresholding is applied to create a mask that separates the magnetic structure from the background. Thresholding isolates the nanomagnet segment by setting pixel intensities above a defined threshold to white (255) and others to black (0). The contours of the binary mask are extracted to identify the outer boundary of the nanomagnet. The bounding box of the segment is calculated using contour fitting, providing the position (x, y) , width, and height and the diagonal length of the bounding box is computed to determine the appropriate arrow size. The orientation of each segment is estimated by fitting an ellipse to the extracted contour. The major axis of the ellipse represents the longest dimension of the nanomagnet, which is used as the magnetization direction. The orientation angle of the ellipse is extracted and rotated by 90° to align with the expected moment direction.

To establish the direction of the magnetic moment, our algorithm analyzes brightness variations across the nanomagnet. The segment is split into two halves along the major axis of the ellipse and brightness values from both halves are extracted, and sudden intensity changes are identified using gradient analysis. The radial gradient of brightness are calculated starting from the separation line of towards the edge of the contour. If there is a sudden intensity change along the major axis of the ellipse, we identify it as the artifact of the dipole as shown in the zoomed in segment image in Fig. 4, considering that each segment is expected to have a brightness gradient from bright to dark, inferring bright side as the head of the dipole. Filtered brightness values are computed to reduce noise effects and capture a more reliable contrast between the two sides. The side with higher brightness is inferred to be the head of the dipole, as MFM contrast reflects magnetization variations. If $I(x, y)$ be the grayscale intensity at pixel (x, y) . The brightness of each half is computed as:

$$B_{side_1} = \frac{1}{N_1} \sum_{(x,y) \in side_1} I(x, y) \quad (8)$$

$$B_{side_2} = \frac{1}{N_2} \sum_{(x,y) \in side_2} I(x, y) \quad (9)$$

where: B_{side_1} and B_{side_2} are the mean brightness values of the two halves, N_1 and N_2 are the number of pixels in each half.

The brightness difference between the two halves is normalized to scale the arrow length, so that it varies proportionally with the MFM contrast. The arrow length is constrained within a defined range to maintain visual clarity and the direction of the arrow is determined based on the brighter half, with the arrowhead pointing towards it.

The contrast strength (C_s), which defines magnetization strength, is given by:

$$C_s = |B_{side_1} - B_{side_2}| \quad (10)$$

Since pixel intensities in grayscale images range from 0 to 255, the contrast is normalized: $C_s^{norm} = \frac{C_s}{255}$ so that C_s^{norm} ranges from [0,1]. With normalizing, the contrast

strength is scaled relative to the maximum brightness range. The arrow length is then linearly scaled between the predefined minimum and maximum arrow lengths such that:

$$L_{arrow} = L_{min} + C_s^{norm} \cdot (L_{max} - L_{min}) \quad (11)$$

The L_{max} and L_{min} are chosen to be 90% and 55% respectively of the diagonal length of the identified contour (L_{diag}) where $L_{diag} = \sqrt{w^2 + h^2}$. The final arrow direction is determined by identifying which half of the segment is brighter:

$$\theta_{arrow} = \begin{cases} \theta_{major} + 90, & \text{if } B_{side_1} > B_{side_2} \\ \theta_{major} + 270, & \text{otherwise} \end{cases} \quad (12)$$

The θ_{major} determines the major axis of the fitted ellipse and if $Side_1$ is brighter the arrow points along the major axis, otherwise, the arrow is rotated by 180° to indicate reversal.

This magnetic moment and strength is iterated over all the segmented nanomagnet from the MFM image and then reassembled following the morphology image. The net magnetic moment, M_{net} , of the ASI sample is calculated as the sum of arrow lengths, where the arrow length represents the relative strength of the magnetic moment for each segment.

$$M_{net} = \sum_{i=1}^N L_{arrow,i} \quad (13)$$

here L_{arrow} is the arrow length corresponding to the magnetic moment strength of the i -th segment and N is the total number of segments.

To generalize the scaling of the net magnetic moment from arbitrary units (a.u.) to physical units (mT) use the unit used in experiment, we define a scaling transformation that maps the computed image-based net magnetization to the experimentally measured range. From the experimental section, we define the real magnetic moment range: $Real_{Range} = M_{max} - M_{min}$. From image processing algorithm, we define the arbitrary range: $Arbitrary_{Range} = L_{max} - L_{min}$. Thus the scaling factor is computed as:

$$S = \frac{Real_{Range}}{Arbitrary_{Range}} \quad (14)$$

$$M_{real} = \left(\frac{M_{max} - M_{min}}{L_{max} - L_{min}} \right) \cdot M_{image} \quad (15)$$

$$M_{real} = S \cdot M_{image}, \quad (16)$$

where M_{real} is the real net magnetic moment and M_{image} is the net magnetic moment computed from image processing (in arbitrary units, a.u.).

Variational Auto Encoder Architecture

We implemented VAE model to learn and generate synthetic MFM images from an input dataset of segmented nanomagnets. The workflow follows the schematic architecture provided in the Fig.6, representing a structured representation of encoder, latent space, and decoder operations.

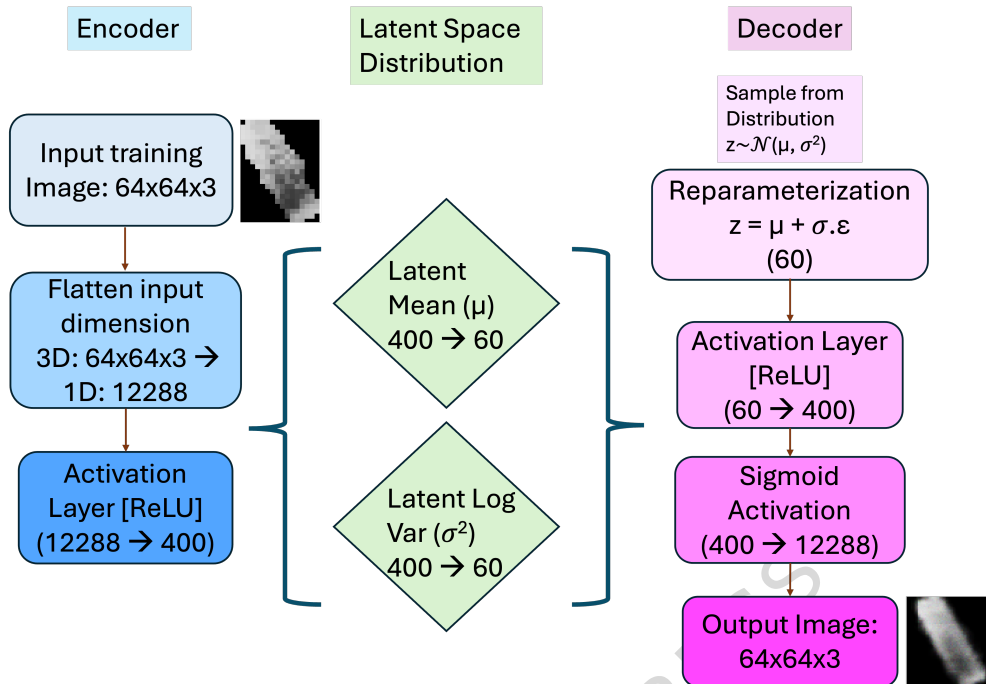


Fig. 6 Variational autoencoder architecture used to generate synthetic MFM segments from the training images. Synthetic MFM segments are generated using VAE architecture. It consists of the following steps: The training image of dimension ($H=64 \times W=64 \times \text{RGB-Channel}=3$) is flattened into a vector of size 12288 for processing. The encoder part of the architecture consists of a fully connected layer with input size of 12288, flattened image and ReLU activation function outputs a compressed representation of 400 while retaining the relevant information. The encoder then splits into two separate layers to model the latent space. The mean (μ) of the latent Gaussian distribution represents the central tendency of the latent space for each given input. It determines the "position" of the latent vector in the latent space of dimension 60 and the Latent log layer produces the log variance ($\log\sigma^2$) of the latent Gaussian distribution. The latent vector "z" is sampled from the Gaussian distribution defined by μ and σ : $z = \mu + \sigma \cdot \epsilon$, where $\epsilon \sim N(0, 1)$, is the random noise sampled from standard normal distribution. The first fully connected dense decoder layer begins to expand the compressed latent vector back into a higher-dimensional representation with ReLU activation function capturing the non-linear transformations. The Sigmoid activation of the 2nd fully connected dense layer of the decoder ensures the output is normalized matching the input image's pixel intensity range. BCE loss function measures the reconstruction loss and KL divergence regularizes the latent space by encouraging the learned distribution $q(z|x)$ to be close to the prior distribution $p(z) \sim N(0, 1)$.

The original dataset is initialized using a CSV file containing segment metadata and each image is retrieved using an index-based file naming scheme. The segment image is converted from BGR to RGB format for compatibility with the PyTorch [23] module tensors. The grayscale thresholding is applied to extract relevant nanomagnetic features, normalized by scaling pixel values [0,1] and reshaped to (C, H, W) tensor format. The dataset is split into training (80%) and validation (20%) subsets to prevent the overfitting and custom collate function filters out corrupted or missing images. The VAE follows a structured encoder-bottleneck-decoder approach:

Encoder: This section of the model compresses the input image into a lower-dimensional latent representation. The input image is flattened from (64,64,3) to a 1D vector of size 12,288. A fully connected layer reduces the feature dimension to 400 and the activation function ReLU introduces the non-linearity. Two separate fully connected (FC) layers compute the latent mean vector (μ), representing the central position of each data point in latent space, and the log variance vector ($\log \sigma^2$), indicating how much the data spreads around this position. To address the non-differentiability of directly sampling from the distribution defined by μ and σ^2 , which would prevent gradient backpropagation and model training, the model employs the reparameterization trick, sampling from a standard Gaussian distribution and then scaling and shifting this sample using the learned μ and σ . The randomness is introduced only from ϵ , which is independent of the model parameters. The backpropagation gradients flow through μ and σ , allowing the network to learn both the mean and variance effectively:

$$z = \mu + \sigma \cdot \epsilon, \quad (17)$$

where ϵ follows a standard normal distribution $N(0,1)$. This step establishes stochasticity in image generation, allowing continuous latent space sampling. This probabilistic space allows the model to learn meaningful and continuous representations of the input (e.g., shape, color, or orientation of an object). Unlike standard autoencoders, which encode input into a deterministic latent vector, VAEs treat the latent space as a probability distribution which assures that the latent space is smooth and continuous, allowing for interpolation and generation of new data points. The latent dimension, chosen as 60, describes the size of each latent vector z , the compressed representation of each sample. The latent space features are learned representations that encode important attributes of the input data. Each dimension in z captures different aspects of the input, such as: global features like shape, orientation, size, or brightness of an image and local features like texture, edges, or specific patterns. In an image dataset, some latent dimensions might encode properties like image brightness in z_1 , orientation of objects in z_2 or texture details in z_3 . However, the exact meaning of each dimension is not explicitly defined but is learned implicitly during training.

Decoder: The latent representation, z (dim = 60) is passed through a fully connected ReLU layer to map it back to 400 dimensions and another fully connected layer expands it back to 12,288 dimensions. A sigmoid activation function ensures that the pixel values are bounded between 0 and 1, so that the reconstructed output can be reshaped back into the original image dimensions (64,64,3). The loss function of the VAE model consists of two key components:

(a) Reconstruction Loss (Binary Cross-Entropy - BCE): The reconstruction loss measures how accurately the decoder can reconstruct the original input images from the latent representation. Specifically, the binary cross-entropy (BCE) loss quantifies the difference between the input image and the reconstructed image on a pixel-by-pixel basis. BCE loss is defined as:

$$\mathcal{L}_{\text{recon}} = \sum [x_i \log(\hat{x}_i) + (1 - x_i) \log(1 - \hat{x}_i)] \quad (18)$$

where, x_i is the original input image's pixel value at position i , normalized between 0 and 1, and \hat{x}_i is the reconstructed image's pixel value at the corresponding position i , also normalized between 0 and 1. This equation computes the pixel-wise reconstruction loss, summing over all pixels in the image. Minimizing this loss term encourages the decoder to generate output images closely matching the input images.

(b) KL Divergence (KLD) Loss: The KL Divergence loss acts as a regularization term, enforcing the latent space distribution to approximate a standard Gaussian distribution ($\mathcal{N}(0, 1)$). By doing so, it prevents the model from overfitting and encourages the latent space to have desirable properties, such as continuity and smoothness. The KL divergence loss is computed as:

$$\mathcal{L}_{\text{KL}} = -\frac{1}{2} \sum_j (1 + \log \sigma_j^2 - \mu_j^2 - \sigma_j^2) \quad (19)$$

where μ_j is the mean of the latent Gaussian distribution for the latent dimension j and σ_j^2 is the variance of the latent Gaussian distribution for the latent dimension j . The $\log \sigma_j^2$ is computed directly by the model, and this equation ensures that the learned latent distributions do not deviate significantly from a normal distribution, thus regularizing the learned latent representations.

The total loss function for the VAE combines both reconstruction and regularization losses as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}} + \beta \cdot \mathcal{L}_{\text{KL}} \quad (20)$$

where the hyperparameter β controls the relative importance or balance between reconstruction accuracy and latent space regularization. Adjusting β allows for fine-tuning of the latent representation and reconstruction quality according to specific application needs.

The VAE model is trained using PyTorch's Adam optimizer [24], with learning rate of 0.005, batch size of 64 for 100 epochs. During each epoch the images are loaded and flattened. The VAE forward pass computes latent distribution (μ, σ), sampled latent vector z , the reconstructed image. The loss function is then calculated and backpropagated while the optimizer updates model weights to improve reconstruction quality. The VAE training converges rapidly, and extending training beyond ~ 40 epochs does not materially change the reconstructed dipole orientations or vertex frustration classification. The training and validation loss evolution is shown in Fig. S1. After training the test MFM images are passed through the VAE to generate synthetic reconstructed images. The pipeline for synthetic image reconstruction can be summarized as follows: original test images are encoded into the latent space, and then the decoder generates synthetic reconstructions. The reconstructed images are saved and compared with the original samples. For validating the synthetic image quality, we compare original with generated images visually as well as evaluate pixel-wise mean squared error (MSE), SSIM and absolute pixel difference between original and synthetic samples as described in the S3, to verify that the latent space retains key magnetic features.

Frustration Analysis

The generated MFM image of honeycomb ASI is then processed for classifying frustrations and also identifying nanomagnets whose directions are needed to be toggled to minimize frustrations. The frustration classification has 3 stages:

- (a) Identification of segment centroid
- (b) Honeycomb vertices identification
- (c) Classification of the vertices

Identification of Segment Centroid: For each synthetic MFM segment, converted to grayscale, the binary threshold is applied to remove the background and isolating the nanomagnetic structure. A bounding box is computed around the segment such that the boundary matches with the original segment. The connected component mask extracts contours representing the original segment with an area above a predefined threshold of 20 pixels and the rest are discarded to avoid artifacts or partial segments. Once segment bounding boxes are established, each synthetic segment is resized to match the bounding box of its original counterpart and overlaid onto the reassembled image at the correct location as explained in the Section 4.2. The centroid of each segment is calculated using:

$$x_{center} = x + \frac{w}{2}, y_{center} = y + \frac{h}{2} \quad (21)$$

where (x,y) are the top-left coordinates of the bounding box, and (w,h) are its width and height. Each centroid position is stored in a list alongside its segment label for reference as Segment ID, the unique identifier for each nanomagnetic segment, x-coordinate which is the horizontal position of the centroid and y-coordinate stores vertical position of the centroid.

Vertex Identification of Honeycomb lattice Spin Ice: In an artificial spin-ice honeycomb lattice, a vertex is defined as the meeting point of three adjacent nanomagnets [18]. Each vertex corresponds to a trivalent interaction, where the three magnetic dipoles contribute to the local frustration state. So from the calculated centroid we can deduce that each vertex must be formed by three neighboring segment centroids. To identify vertices, a maximum distance threshold (radius = 30 pixels) is applied which defines the maximum allowable distance between two segment centroids to be considered part of the same vertex. To identify all possible vertices, our algorithm considers every unique triplet of segments from the dataset which stores the segment ID, and their corresponding centroid (x,y). Using combinatorial selection, we iterate over all segment triplets, $Triplet = \{(x_1, y_1), (x_2, y_2), (x_3, y_3)\}$, and each possible set of three segments is extracted and their centroid positions are retrieved.

For each triplet of segments, the pairwise Euclidean distances are computed to check for spatial connectivity:

$$d_{12} = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2} \quad (22)$$

$$d_{13} = \sqrt{(x_3 - x_1)^2 + (y_3 - y_1)^2} \quad (23)$$

$$d_{23} = \sqrt{(x_3 - x_2)^2 + (y_3 - y_2)^2} \quad (24)$$

If all three distances are within the predefined radius threshold (r), the segments are considered to form a valid vertex; $d_{12}, d_{13}, d_{23} \leq r$.

To calculate the vertex centroid, the vertex location is computed as the centroid of the triangle formed by the three segments:

$$x_{vertex} = \frac{x_1 + x_2 + x_3}{3}, y_{vertex} = \frac{y_1 + y_2 + y_3}{3} \quad (25)$$

Thus each vertex is assigned a unique central coordinate.

The energy state of each vertex is determined by its three associated segments. Therefore, we classify the vertices to identify the high-energy points within the honeycomb lattice.

Frustration Classification: In a nanoscale system, frustration occurs when competing magnetic interactions prevent a system from reaching a globally minimized energy state [25]. Each vertex in the honeycomb lattice consists of three interacting nanomagnets, forming a three-state system where the dipole orientations define its frustration state.

The frustration states in the VAE-generated ASI system are classified as high energy and low energy vertices. The high energy vertices are sub-classified into $+3q$ (3-in), direction of three arrows from the associated segments convergent towards the vertex and $-3q$ (3-out), where the three arrows are divergent from the vertex. The low energy vertices are also sub-classified into $+q$ (2-in, 1-out) and $-q$ (2-out, 1 in). Thus the analysis is based on the magnetic moments of nanomagnet segments and their interactions at lattice vertices.

From above sections we identified each segment and vertex positions of the VAE generated ASI system and the magnetic moment orientations are extracted from the synthetic magnetic moment dataset the method of which are explained in section 4.2. The vertex segment associations are retrieved to map each segment to its corresponding vertex and each segment's magnetic moment is then mapped to represent the magnetic moment set (M_s) for a given vertex consisting of three segment IDs (S_1, S_2, S_3): $M_s = \{S_1, S_2, S_3\}$. For each segment, we extract its binary direction (0,1) from the dataset and assign magnetic moments pointing towards right as 0 and left as 1.

A vertex is classified based on the relationship between the horizontal segment (H), the two angled segments (A_1, A_2) and the corresponding vertex position. If magnetic moment of the two angled segments match while the horizontal segment opposes i.e., $A_1 = A_2 \neq H$, the system is unable to achieve a minimum energy state, leading to frustration and if the magnetic moment of all three segments are equal i.e., $A_1 = A_2 = H$, the system achieves a stable low-energy configuration. Now we further classify the frustrated vertex and we consider the position of vertex with respect to the horizontal segment as the reference point. To identify the horizontal segments, we extracted the contours of the segment and apply an ellipse fit to estimate the orientation. We calculate the segment's angle (θ) from its bounding contour, S_i for the i -th segment,

as given by:

$$\theta = \text{fitEllipse}(S_i) \quad (26)$$

The angle of the horizontal segment (θ_h) associated with each vertex is then identified with an angle closest to 0° or 90° , as

$$\theta_h = \min(|\theta - 0^\circ|, |\theta - 90^\circ|) \quad (27)$$

The frustrated vertex class might have two conditions: (a) the detected horizontal segment is shared between two frustrated vertices and (b) the frustrated vertices are disconnected. For the first case, if a vertex lies to the right of the horizontal segment and the magnetic moment associated with the horizontal segment is 1, the right vertex is classified as "+3q" and assigned red squares as shown in Fig. 5 and the left connected vertex is classified as "-3q", denoted by magenta squares.

If a vertex lies to the left to the horizontal segment and the magnetic moment of the horizontal segment (M_H) is 0, then the left vertex is classified as "+3q" and the corresponding right vertex is "-3q". This classification is based on energy asymmetry in the ASI system, where relative positions impact frustration strength. For each vertex, the algorithm compares its x-coordinate to the x-coordinate of the horizontal segment.

If $x_{vertex} > x_{segment}$ and $M_H = 1$: +3q and the other connected vertex is -3q

If $x_{vertex} < x_{segment}$ and $M_H = 0$: +3q and the other connected vertex is -3q.

For the second case where the horizontal segment is not shared by two frustrated vertices, we remove the logic where the other connected vertex is the opposite frustration. So the logic can be simplified as:

If $x_{vertex} > x_{segment}$ and $M_H = 1$: +3q and

if $x_{vertex} < x_{segment}$ and $M_H = 0$: +3q.

The low energy state class of vertices, where the interactions between the three segments forming a vertex are minimally strained, are also sub classified with reference to the relative position and magnetic moment of the horizontal segment.

If the vertex is to the left of the horizontal segment ($x_{vertex} < x_{segment}$),

For $M_H = 1$, the vertex is +q denoted by green star and

for $M_H = 0$, the vertex is -q denoted by yellow star.

If the vertex is to the right of the horizontal segment ($x_{vertex} > x_{segment}$),

For $M_H = 0$, the vertex is +q denoted by green star in Fig. 5 and

for $M_H = 1$, the vertex is -q denoted by yellow star.

In some cases, low energy states may still exist even when $A_1 \neq A_2$, meaning that magnetic moment of the two angled segments do not point in the same direction. In such cases,

If $x_{vertex} < x_{segment}$ and $M_H = 1$, it is classified as +q (Green star) and $M_H = 0$ it is classified as -q (Yellow star).

If $x_{vertex} > x_{segment}$ and $M_H = 0$, it is classified as +q (Green star) and $M_H = 1$ it is classified as -q (Yellow star).

This detailed classification method helps in distinguishing stable (low-energy) regions in the spin-ice system, which is crucial for analyzing magnetic order and energy distribution.

5 Declarations

5.1 Data availability

The representative datasets are currently under institutional review by the Export Control and Classification teams. The authors will make the representative datasets available through a public repository following completion of this review.

5.2 Code availability

The code associated with this work is currently under institutional review by the Export Control and Classification teams. The authors will make the code available through a public repository following completion of this review and the required approval process by NNSA patent counsel, after which a final notice will be issued to release the code.

5.3 Acknowledgements

This work at Los Alamos was carried out under the auspices of the U.S. Department of Energy (DOE) National Nuclear Security Administration (NNSA) under Contract No. 89233218CNA000001. It was supported by Center for Integrated Nanotechnologies (CINT), a DOE BES user facility, in partnership with the LANL Institutional Computing Program for computational resources. The fabrication of ASI samples was performed, in part, at CINT, Sandia National Laboratories, a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. DOE's National Nuclear Security Administration under contract DE-NA-0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

5.4 Author contribution

A.N. developed all the codes for image processing, Auto-Encoder model development, frustration prediction method development and prepared the manuscript. S.M. performed the MFM measurements, generated microscopic image data, all related experiments and prepared the manuscript. P.P.I., T.M.L. and E.B. prepared the Artificial Spin Ice sample, S.T., A.C.J. and J.X.Z. supervised the project in conception and discussion and prepared the manuscript.

5.5 Conflict of interest/Competing interests

The authors declare no competing interests.

ARTICLE IN PRESS

References

- [1] Skjærvø, S.H., Marrows, C.H., Stamps, R.L., Heyderman, L.J.: Advances in artificial spin ice. *Nature Reviews Physics* **2**(1), 13–28 (2020)
- [2] Gilbert, I., Chern, G.-W., Zhang, S., O'Brien, L., Fore, B., Nisoli, C., Schiffer, P.: Emergent ice rule and magnetic charge screening from vertex frustration in artificial spin ice. *Nature Physics* **10**(9), 670–675 (2014)
- [3] Lendinez, S., Jungfleisch, M.: Magnetization dynamics in artificial spin ice. *Journal of Physics: Condensed Matter* **32**(1), 013001 (2019)
- [4] Zhang, S., Gilbert, I., Nisoli, C., Chern, G.-W., Erickson, M.J., O'Brien, L., Leighton, C., Lammert, P.E., Crespi, V.H., Schiffer, P.: Crystallites of magnetic charges in artificial spin ice. *Nature* **500**(7464), 553–557 (2013)
- [5] Chumak, A.V., Kabos, P., Wu, M., Abert, C., Adelmann, C., Adeyeye, A., Åkerman, J., Aliev, F.G., Anane, A., Awad, A., *et al.*: Advances in magnetics roadmap on spin-wave computing. *IEEE Transactions on Magnetics* **58**(6), 1–72 (2022)
- [6] Krawczyk, M., Grundler, D.: Review and prospects of magnonic crystals and devices with reprogrammable band structure. *Journal of physics: Condensed matter* **26**(12), 123202 (2014)
- [7] Sievers, S., Braun, K.-F., Eberbeck, D., Gustafsson, S., Olsson, E., Schumacher, H.W., Siegner, U.: Quantitative measurement of the magnetic moment of individual magnetic nanoparticles by magnetic force microscopy. *small* **8**(17), 2675–2679 (2012)
- [8] Hu, W., Zhang, Z., Liao, Y., Li, Q., Shi, Y., Zhang, H., Zhang, X., Niu, C., Wu, Y., Yu, W., *et al.*: Distinguishing artificial spin ice states using magnetoresistance effect for neuromorphic computing. *Nature Communications* **14**(1), 2562 (2023)
- [9] Castelnovo, C., Moessner, R., Sondhi, S.L.: Magnetic monopoles in spin ice. *Nature* **451**(7174), 42–45 (2008)
- [10] Morris, D.J.P., Tennant, D.A., Grigera, S.A., Klemke, B., Castelnovo, C., Moessner, R., Czternasty, C., Meissner, M., Rule, K., Hoffmann, J.-U., *et al.*: Dirac strings and magnetic monopoles in the spin ice $\text{Dy}_2\text{Ti}_2\text{O}_7$. *Science* **326**(5951), 411–414 (2009)
- [11] Nisoli, C., Li, J., Ke, X., Garand, D., Schiffer, P., Crespi, V.H.: Effective temperature in an interacting vertex system: Theory and experiment on artificial spin ice. *Physical Review Letters* **105**(4), 047205 (2010)
- [12] Ladak, S., Read, D., Perkins, G., Cohen, L., Branford, W.: Direct observation of

- magnetic monopole defects in an artificial spin-ice system. *Nature Physics* **6**(5), 359–363 (2010)
- [13] Gartside, J.C., Arroo, D.M., Burn, D.M., Bemmer, V.L., Moskalenko, A., Cohen, L.F., Branford, W.R.: Realization of ground state in artificial kagome spin ice via topological defect-driven magnetic writing. *Nature nanotechnology* **13**(1), 53–58 (2018)
- [14] Kazakova, O., Puttock, R., Barton, C., Corte-León, H., Jaafar, M., Neu, V., Asenjo, A.: Frontiers of magnetic force microscopy. *Journal of applied Physics* **125**(6) (2019)
- [15] Kalinin, S.V., Sumpster, B.G., Archibald, R.K.: Big-deep-smart data in imaging for guiding materials design. *Nature materials* **14**(10), 973–980 (2015)
- [16] Ziatdinov, M., Dyck, O., Maksov, A., Li, X., Sang, X., Xiao, K., Unocic, R.R., Vasudevan, R., Jesse, S., Kalinin, S.V.: Deep learning of atomically resolved scanning transmission electron microscopy images: chemical identification and tracking local transformations. *ACS nano* **11**(12), 12742–12752 (2017)
- [17] Wang, .R., Nisoli, C., Freitas, R., Li, J., McConville, W., Cooley, B., Lund, M., Samarth, N., Leighton, C., Crespi, V., *et al.*: Artificial ‘spin ice’ in a geometrically frustrated lattice of nanoscale ferromagnetic islands. *Nature* **439**(7074), 303–306 (2006)
- [18] Nisoli, C., Moessner, R., Schiffer, P.: Colloquium: Artificial spin ice: Designing and imaging magnetic frustration. *Reviews of Modern Physics* **85**(4), 1473–1490 (2013)
- [19] Nisoli, C., Wang, R., Li, J., McConville, W.F., Lammert, .f.P.E., Schiffer, P., Crespi, V.H.: Ground state lost but degeneracy found: the effective thermodynamics of artificial spin ice. *Physical review letters* **98**(21), 217203 (2007)
- [20] Kempinger, S., Huang, Y.-S., Lammert, P., Vogel, M., Hoffmann, A., Crespi, V.H., Schiffer, P., Samarth, N.: Field-tunable interactions and frustration in underlayer-mediated artificial spin ice. *Physical review letters* **127**(11), 117203 (2021)
- [21] Bradski, G.: The OpenCV Library. *Dr. Dobb’s Journal of Software Tools* (2000)
- [22] Beucher, S.: Use of watersheds in contour detection. In: *Proc. Int. Workshop on Image Processing*, Sept. 1979, pp. 17–21 (1979)
- [23] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., *et al.*: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)

- [24] Zhang, Z.: Improved adam optimizer for deep neural networks. In: 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS), pp. 1–2 (2018). Ieee
- [25] Drisko, J., Marsh, T., Cumings, J.: Topological frustration of artificial spin ice. Nature communications **8**(1), 14009 (2017)

ARTICLE IN PRESS

Figure legends

Fig. 1 "Magnetic AFM setup and representative AFM topography and MFM phase images of the artificial spin ice sample."

Fig. 2 "Workflow for VAE generated MFM reconstruction and frustration optimization in the artificial spin ice lattice."

Fig. 3 "Morphology segmentation and mapping onto the MFM phase image to extract individual nanomagnet segments."

Fig. 4 "Reconstruction of segmented nanomagnets and identification of artifacts during magnetic moment determination."

Fig. 5 "Progression from experimental MFM images to VAE assisted moment maps, vertex classification, and frustration minimization by selected flips."

Fig. 6 "Variational autoencoder architecture used to generate synthetic MFM segments from the training images."

ARTICLE IN PRESS