Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantized composition penalties alter state selection in phase-field Ising models

Zhiyi Liu
Sichuan Qiming Daren Technology Co., Ltd.
Anping Liu
Sichuan Qiming Daren Technology Co., Ltd.
Zicong Zhu
zhuzc0907@outlook.com

Sichuan Qiming Daren Technology Co., Ltd.

Article

Keywords: phase-field models, Ising encoding, composition constraints, coefficient quantization; 20 quadratic unconstrained binary optimization, energy fidelity.

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11603792/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Quantized composition penalties alter state selection in phase-field Ising models

Zhiyi Liu[1,2], Anping Liu[1], Zicong Zhu[1,*]

[1]Tianli Qiming AI Research Institute, Sichuan Qiming Daren Technology Co., Ltd., Chengdu, China

[2]International Research Center for Complexity Sciences, Hangzhou International Innovation Institute, Beihang University, Hangzhou, China

[*]Correspondence: zhuzc0907@outlook.com

Abstract

Phase-field models seek low-energy spatial structures at fixed composition, but finite-precision Ising formulations can change which structures are selected. We investigate how rounding alters energy ordering among composition-preserving states and why protecting the composition penalty may not improve physical accuracy. Using a one-dimensional model with four allowed local field values and an exact binary encoding before rounding, we combine analytic energy decomposition with exhaustive enumeration of all configurations and tied minima. Rounding the composition penalty creates artificial preferences among feasible states; jointly rounding physical and penalty coefficients can reverse physical-energy ordering even when residual penalty energies are identical. In the calibrated comparison, exact penalty preservation reduces reversals but consumes coefficient range, making physical energy differences harder to resolve and increasing ties with inferior states. Averaged over 96 held-out instances and six prespecified precisions, the separately calibrated penalty-preserving policy lowers the fraction of encoded minima satisfying composition and the physical-energy tolerance by 2.14 percentage points relative to conventional rounding. A separate continuum comparison identifies a physical minimum excluded by the allowed field values. These results connect coefficient errors to selected structures and explain how increased inferior ties can outweigh reduced ordering reversals.

Keywords: phase-field models; Ising encoding; composition constraints; coefficient quantization; quadratic unconstrained binary optimization; energy fidelity.

Introduction

Phase-field models describe spatially varying local composition and assign free energy to both the local composition and its variation across space [1]. Equilibrium structures minimize this energy subject to composition constraints. A programmable Ising machine instead minimizes an objective built from binary variables and their pair interactions. Connecting these objectives requires spatial discretization, a finite set of allowed local values (the local alphabet), binary encoding, constraint penalties and coefficient quantization, which rounds coefficients to representable values. These steps can change the energy ordering of candidate structures, so an exact minimum of the encoded objective need not minimize the original physical energy.

Mapping phase-field equilibrium problems to Ising formulations is already established [2]. Aoki and colleagues used factorization-machine corrections to reduce discrepancies between a quadratic unconstrained binary optimization (QUBO) representation and phase-field free energy [3]; related materials optimization with factorization-machine QUBOs also predates the present work [4]. In microstructure calculations with elastic interactions, precision and interaction truncation can affect the computed equilibria [5]. These studies motivate examining the changes introduced after an otherwise exact finite-state encoding.

The general penalty dilemma is also known. Too little penalty allows constraint violation, whereas a large penalty can obscure the original objective after normalization [6, 7]. Finite-precision QUBO compression studies conditions for preserving minimizing solutions while reducing dynamic range [8, 9]. Linear penalties, variable reduction and constraint-dependent coefficient conditioning provide additional alternatives in appropriate settings [10–12]. Zeng and colleagues also describe quantization-induced minima with different quality under the original objective [12]. Building on these observations, we ask which feasible field configurations coefficient rounding favors and how changes in reversals and ties combine to determine physical fidelity.

We study this question in an exactly encoded four-level phase field. In exact arithmetic, the squared composition penalty is zero for every feasible structure. Rounding its couplings introduces an artificial preference within that set, which can be expressed directly through the populations of the four physical field levels. Integer alignment preserves the penalty on an exact coefficient lattice while consuming part of the range available to resolve physical energy differences.

The analysis establishes two connected results. First, an exact decomposition links rounding to physical-state selection: the complete objective contains both physical-coefficient error and a joint-rounding interaction, and a validation example identifies the on-site and neighboring couplings that reverse physical ordering. Second, enumerating all configurations tied for the lowest encoded energy (the complete ground set) explains why protecting the penalty can lower mean target-energy fidelity even when strict reversals become less frequent. A strict reversal leaves no physically acceptable state among the lowest-energy feasible states. Under the separately calibrated policies, additional ties with inferior states outweigh the reduction in reversal loss. An independent Fourier–Galerkin reference complements these results by identifying a physical minimum excluded by the finite alphabet.

Results

An exact restricted encoding isolates later compilation error

We use a standard phase-field free energy with a symmetric double-well local potential and a gradient cost [1, 13], together with a prescribed spatial bias. All lengths and energies are dimensionless. The one-dimensional domain has length $L = 8$ and periodic boundaries, so its ends are connected. We record the local composition variable ϕ_i at n equally spaced grid points, with n even and spacing $h = L/n$. For a binary-mixture interpretation, the local fractions of components A and B are $(1 + \phi_i)/2$ and $(1 - \phi_i)/2$: $\phi_i = 1$ and -1 represent the two pure components. The

discrete free energy is

$$F_h(\phi) = h\sum_{i=1}^{n}\left[\frac{(\phi_i^2-1)^2}{4} + v_i\phi_i\right] + \frac{\kappa}{2h}\sum_{i=1}^{n}(\phi_{i+1}-\phi_i)^2, \quad \sum_i \phi_i = 0, \tag{1}$$

where $\phi_{n+1} = \phi_1$. The double-well term favors local values near ± 1; the gradient term penalizes differences between neighbors, with strength $\kappa > 0$. The bias v_i gives a position-dependent preference: positive v_i favors negative ϕ_i. The constraint $\sum_i \phi_i = 0$ fixes equal total amounts of A and B. Thus F_h is the total energy of a configuration, not the energy at one point. Nodal refers to evaluation from grid-point values; Methods explains the factors h and $1/h$.

We restrict each ϕ_i to $\{-1, -1/3, 1/3, 1\}$. Each value is represented exactly by two binary variables $b_{i0}, b_{i1} \in \{0, 1\}$, or by two Ising spins $s_{i\alpha} = 2b_{i\alpha} - 1 \in \{-1, 1\}$ with $\alpha \in \{0, 1\}$:

$$\phi_i = -1 + \frac{2}{3}b_{i0} + \frac{4}{3}b_{i1} = \frac{s_{i0} + 2s_{i1}}{3}, \quad C = \frac{1}{2}\sum_i (s_{i0} + 2s_{i1}). \tag{2}$$

The two spins at each site encode its local composition. The integer-valued composition error C measures departure from the prescribed zero mean: $\sum_i \phi_i = 2C/3$, so the constraint is satisfied exactly when $C = 0$.

For this restricted alphabet, the quartic local term reduces exactly to a quadratic spin interaction,

$$\frac{(\phi_i^2-1)^2}{4} = \frac{8}{81}(1 - s_{i0}s_{i1}). \tag{3}$$

Together with the quadratic gradient term and linear bias, this gives an exact QUBO or Ising representation without additional auxiliary spins. The within-site pair coefficient is $4\kappa/(9h) - 8h/81$, neighboring weighted-spin coefficients are $-\kappa w_\alpha w_\beta/(9h)$ with $w = (1, 2)$, and local fields are $hv_i w_\alpha/3$. Full constants and derivations are provided in the Supplementary Information. This exact representation targets the four-level nodal objective in Eq. (1); other alphabets require their corresponding encodings [14–16].

We enforce composition by adding a penalty, $H_P = F_h + PC^2$, where $P > 0$ controls the cost of violating the constraint. The penalty is zero for every feasible configuration, meaning one with $C = 0$. The penalty adds the constant $5Pn/4$ and pair couplings $(P/2)w_a w_b$ between all distinct encoding spins. Thus a local physical ring becomes a densely coupled encoded problem, even though no auxiliary spins are needed. In the tested sizes, the representation uses 12 or 16 spins and up to 66 or 120 pair coefficients.

A zero-on-feasible penalty becomes a spurious ordering term

Compilation here converts exact Ising coefficients to values supported by a finite coefficient grid. Let θ collect the coefficients multiplying individual spins (fields) and spin pairs (couplings). Conventional shared-range compilation sets $R = \max_a |\theta_a|$, uses the symmetric integer budget $q = 2^{B-1} - 1$, and replaces

$$\delta = \frac{R}{q}, \qquad \theta_a \longmapsto \widetilde{\theta}_a = \delta\,\mathrm{rint}\left(\frac{\theta_a}{\delta}\right), \tag{4}$$

where δ is the spacing between representable coefficients and rint rounds to the nearest integer, choosing the even integer at a midpoint. The signed bit budget B specifies $2q+1$ coefficient levels, including zero; for example, four bits use integer levels $-7, \ldots, 7$, multiplied by δ.

For the pure penalty PC^2 with $P > 0$, $R = 2P$. Rounding acts on the expanded pair coefficients, so it can disrupt the cancellations that make the exact penalty zero when $C = 0$. For example, at four bits with $P = 1$, the step is $\delta = 2/7$, and the coefficients $1/2, 1, 2$ become $4/7, 8/7, 2$. In general, the normalized pair coefficients have three values, $1/4$, $1/2$ and 1. For $B \geq 3$, q is congruent to 3 modulo 4 and these values round to $1/4 + 1/(4q)$, $1/2 + 1/(2q)$ and 1. Define sums over the two encoding-spin groups,

$$X = \sum_i s_{i0}, \qquad Y = \sum_i s_{i1}. \tag{5}$$

Within the feasible sector, the set of configurations satisfying $C = 0$, $X = -2Y$. Substitution into the rounded pair energy gives

$$\widetilde{H}_{\text{penalty}}\big|_{C=0} = \text{constant} - \frac{P}{q} Y^2. \tag{6}$$

The term favors large $|Y|$ even though all these states have exactly zero composition error. This statistic has a direct field interpretation. If N_a counts sites at field level a, composition feasibility gives

$$Y = 2(N_{-1} - N_{+1}) = -\frac{9}{4} \sum_i \phi_i^3. \tag{7}$$

Thus the artifact favors an imbalance between the two saturated levels, compensated by the intermediate levels so that the total composition remains zero. For eight sites, two sites at -1 and six at $+1/3$ have $Y = 4$; reversing all signs gives $Y = -4$. The 56 selected states are these two occupancy patterns with every placement of the two saturated sites. The penalty alone does not distinguish their spatial arrangements; interfacial preference comes from the gradient term in Eq. (1).

For n divisible by four, feasible states include $Y = 0$ and $|Y| = n/2$, so the spurious feasible energy spread is $Pn^2/(4q)$. Exhaustive enumeration at $n = 8$ confirms the expression (Fig. 1). There are 8,092 exactly feasible states, all having the same zero unrounded penalty. At each tested precision from 4 to 12 bits, shared-range rounding selects only 56 global minima, all feasible and all with $Y^2 = 16$. In contrast, the mean Y^2 under uniform counting of all feasible states is approximately 1.6609. This example demonstrates why perfect composition feasibility does not imply preserved structure ranking.

The full objective contains a joint-rounding interaction

Having isolated the penalty's preference, we now include the physical coefficients to identify the terms that determine state selection. Let $Q_\delta(u) = \delta\,\text{rint}(u/\delta)$ for a fixed coefficient step δ. Separately rounding the three pure-penalty couplings gives the feasible-sector coefficient

$$\beta(\delta, P) = 2Q_\delta(P/2) - 2Q_\delta(P) + \tfrac{1}{2} Q_\delta(2P). \tag{8}$$

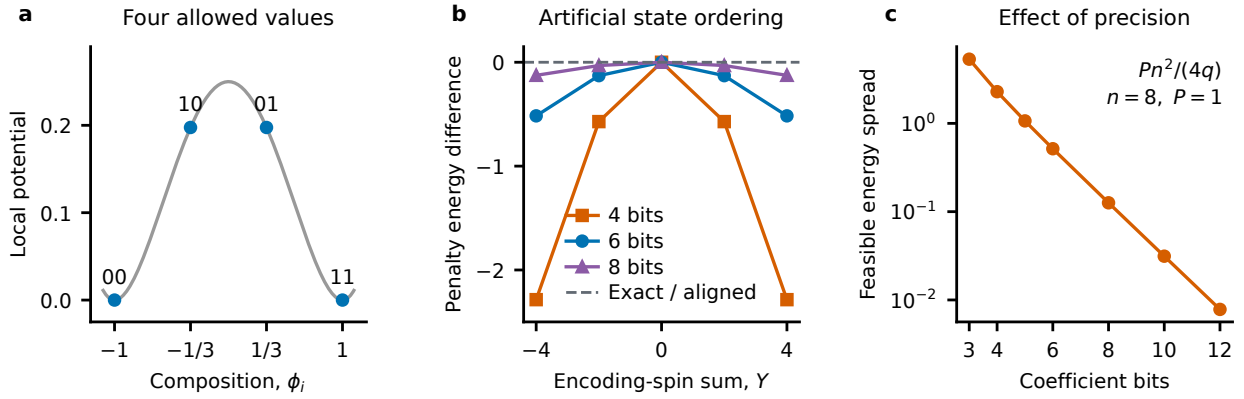


Figure 1: Quantizing a composition penalty creates artificial ordering among feasible states. **a**, Four allowed composition values and binary labels (b_{i0}, b_{i1}) on the local double-well potential. **b**, Pure-penalty energy relative to $Y = 0$ for $n = 8$ and $P = 1$. Conventional rounding gives $-PY^2/q$ at 4, 6 and 8 bits; exact and aligned penalties remain constant. **c**, Feasible-sector energy spread $Pn^2/(4q)$ versus precision, including a three-bit stress case outside the primary benchmark. Here P is the penalty weight, $q = 2^{B-1} - 1$ is the integer coefficient budget at B bits, and $Y = \sum_i s_{i1}$ sums the second encoding-spin group. All quantities are dimensionless. The pure-penalty calculation agrees with exhaustive enumeration.

The penalty contribution is then a constant plus βY^2. In exact arithmetic, the pure-penalty grid and nearest-even rule give $\beta = -P/q$, whereas a grid that exactly represents all penalty couplings gives $\beta = 0$. Equation (8) also applies when the full objective determines a different step.

Rounding a sum need not give the same result as rounding its two parts separately. We separate physical rounding error from the additional change caused by combining physical and penalty coefficients before rounding. Write the nonconstant physical and penalty coefficient vectors as $\boldsymbol{a}$ and $\boldsymbol{b}$, the actual compiled coefficients as $\widetilde{\boldsymbol{\theta}}$, and let $\boldsymbol{\psi}(s)$ collect individual spins s_a and pair products $s_a s_b$. Define

$$e_F(s) = \boldsymbol{\psi}(s)^\top [Q_\delta(\boldsymbol{a}) - \boldsymbol{a}], \tag{9}$$

$$I(s) = \boldsymbol{\psi}(s)^\top [\widetilde{\boldsymbol{\theta}} - Q_\delta(\boldsymbol{a}) - Q_\delta(\boldsymbol{b})]. \tag{10}$$

Here e_F is physical rounding error and I is the residual between actual compilation and the separately rounded blocks. For conventional compilation, $\widetilde{\boldsymbol{\theta}} = Q_\delta(\boldsymbol{a} + \boldsymbol{b})$, so I is the joint-rounding interaction. The same fixed step is used throughout; neither block is renormalized separately. For any two feasible states s, t, write $\Delta g = g(s) - g(t)$. Their compiled energy difference satisfies

$$\Delta\widetilde{H} = \Delta F_h + \beta\Delta(Y^2) + \Delta e_F + \Delta I. \tag{11}$$

All constant offsets cancel. The four terms distinguish the original physical-energy difference, the residual penalty preference, physical rounding and joint rounding. A positive ΔF_h but negative $\Delta\widetilde{H}$ means that compilation assigns a lower encoded energy to the physically worse state. For the separately compiled integer-aligned construction below, $\beta = 0$ and $I = 0$, leaving only physical rounding error.

Reanalysis of all 5,824 saved validation and test configurations recovered every archived coeffi-

cient vector and global ground set. A fixed rule selected illustrative states from the validation set only. In the first strict-reversal case ($n = 8$, $\kappa = 0.5$, bias amplitude $A = 0.02$, four bits), the selected compiled minimum is physically worse than the exact physical optimum by 0.169023. The two states have the same Y^2, so the penalty term contributes zero to their difference. Physical rounding contributes 0.000289, while the joint-rounding interaction contributes -0.507937, giving a compiled difference of -0.338624 (Fig. 2). The compiled objective therefore favors the inferior state through an interaction that the pure-penalty formula alone cannot explain.

Joint rounding changes two coupling classes in this example. At $P = 4/27$ and $\delta = 8/189$, it subtracts one grid unit from each within-site coupling and each neighboring pair of first-group spins relative to separately rounding the physical and penalty blocks. All remaining coefficients have zero interaction, so

$$I(s) = -\delta \left(\sum_i s_{i0} s_{i1} + \sum_i s_{i0} s_{i+1,0} \right). \tag{12}$$

The two sums increase by 4 and 8 in the inferior state, respectively. Thus $\Delta I = -12\delta = -0.507937$: the changed couplings favor that state strongly enough to overcome its physical energy disadvantage. Coefficient-level verification is given in the Supplementary Information. The within-site term favors saturated field values $\phi_i = \pm 1$. The neighboring term favors alignment of the first encoding spins; these spins vary nonmonotonically with ϕ_i, so this preference can accompany a larger physical gradient energy.

The state profiles also explain the physical differences. In the reversal example, the selected field lowers its local potential energy but increases gradient and bias energy, giving a net physical loss. In the tie examples, the states have identical local and gradient energies; their spatial placements differ in the bias field. Conventional rounding erases that distinction for one inferior state, and the coarser aligned grid admits an additional inferior tie. All preserve composition. Local-potential, gradient and bias energies, the complete state vectors and each term in Eq. (11) are reported in the Supplementary Information. The terms form an exact algebraic energy balance and may cancel.

Constraint and physical ordering require separate margins

Let F_* denote the exact minimum of F_h over $C = 0$. An infeasible state can have lower physical energy than the feasible optimum. The critical exact-arithmetic penalty compares this energy advantage with its squared composition error:

$$P_{\mathrm{crit}} = \max_{C(s) \neq 0} \frac{F_* - F_h(s)}{C(s)^2}. \tag{13}$$

For $P > P_{\mathrm{crit}}$, every exact global minimizer is feasible. At equality, an infeasible state can tie the feasible optimum. We calculate this exhaustive diagnostic for the small systems used here and store its nonnegative part; compiler calibration uses the development sets described below.

Finite precision adds a distinct requirement. Suppose two configurations differ in k of $2n$ spins. Exactly k field features and $k(2n - k)$ pair features may change sign. If the rounding step is $\delta = R/q$, their energy-difference error is bounded by $\delta k(2n - k + 1)$ and hence by

$$B_{\mathrm{ord}} = \delta n(n + 1). \tag{14}$$

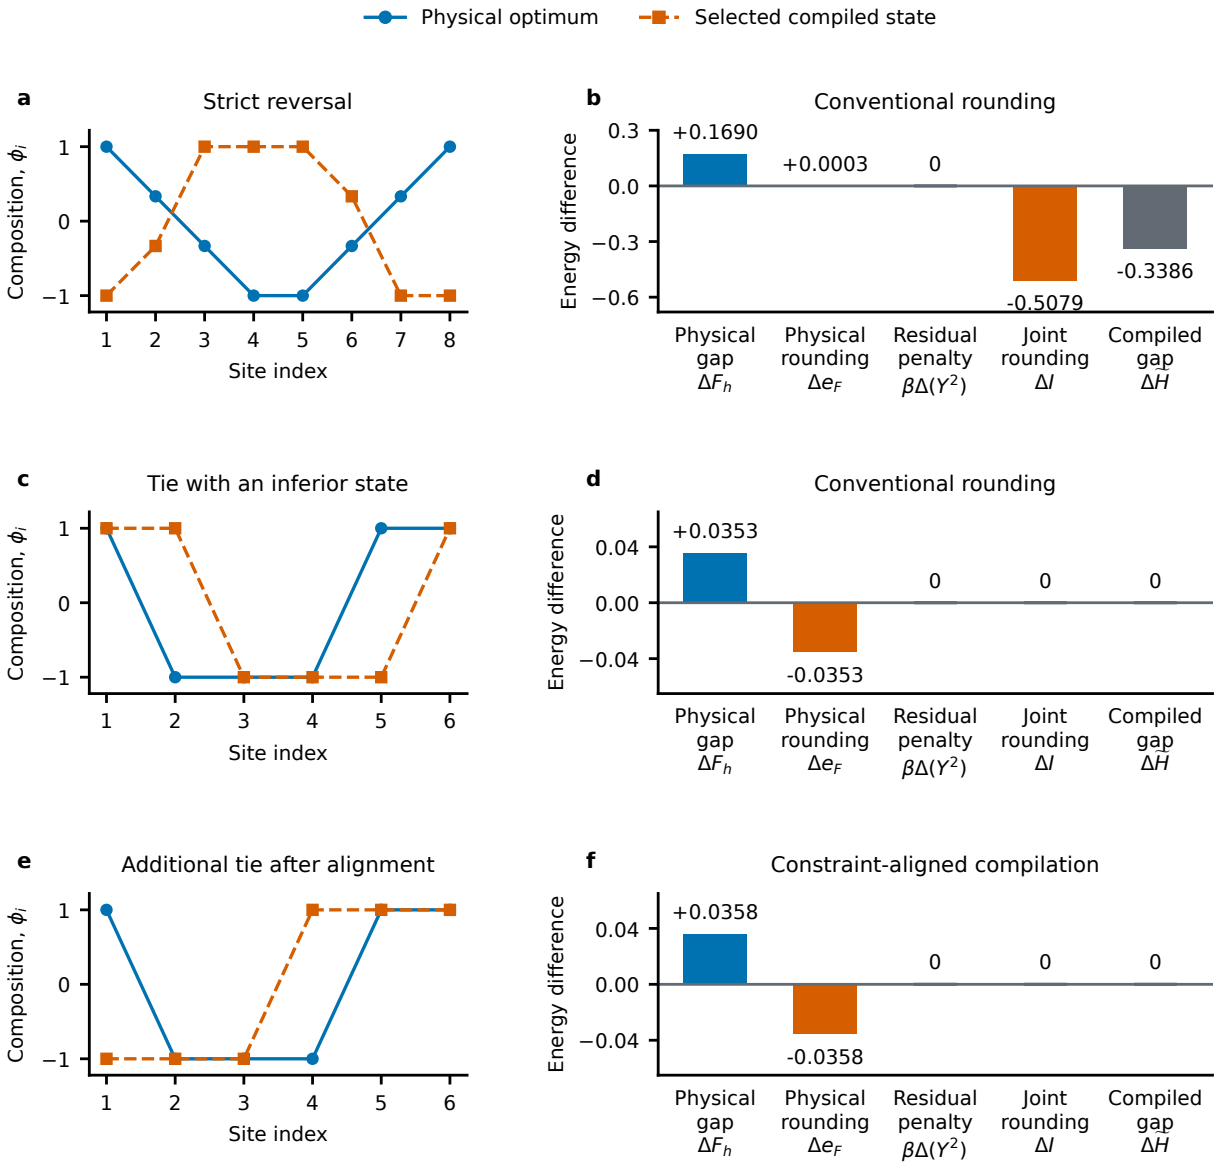


Figure 2: Physical states and energy changes in fixed validation examples. a,b, Strict reversal under conventional rounding. **c,d**, A tie with a physically inferior state under conventional rounding. **e,f**, An additional inferior tie after alignment on the same instance and allocation as c,d. The examples follow a fixed validation-selection rule and use four bits with penalty allocation $\ell = 1$, which sets the penalty weight through Eq. (15). Profiles compare a physical optimum with a selected inferior compiled minimum; all satisfy zero composition. Dimensionless energy differences are selected state minus optimum: physical gap ΔF_h, physical rounding Δe_F, residual penalty $\beta\Delta\langle Y^2\rangle$, joint-rounding interaction ΔI, and their compiled sum $\Delta\widetilde{H}$. Panels b,d,f have individually labelled vertical scales. Full states and energy components are reported in the Supplementary Information.

Define $\Delta_c(P) = \min_{C\neq 0}(F_h + PC^2 - F_*)$ and the smallest positive feasible physical gap $\Delta_f = \min_{C=0,F_h>F_*}(F_h - F_*)$. The first gap separates the feasible optimum from constraint-violating states; the second separates it from physically worse feasible states. The sufficient conditions $\Delta_c > B_{\mathrm{ord}}$ and $\Delta_f > B_{\mathrm{ord}}$ ensure that every quantized global minimizer is feasible and physically optimal. Degenerate physical ground states are retained; they are not assigned an artificial positive gap.

These bounds instantiate established penalty and perturbation principles [6–8]. Their role is to distinguish constraint enforcement from physical ordering. A growing penalty can improve the first margin while increasing R and degrading the second. A conservative explicit penalty interval is derived in the Supplementary Information, but can be empty even for successful compilations. In the complete held-out configuration grid, 306 of 2,496 conventional compilations and 249 of 2,496 aligned compilations satisfied the sufficient exact-optimum certificate. Independent exhaustive checking confirmed every certified configuration and resolved cases outside the certificate.

Integer alignment preserves the penalty block at a range cost

Integer alignment reserves exact integer coefficients for the composition penalty. Let M be the maximum magnitude of an unpenalized physical coefficient. The integer allocation $\ell \geq 1$ reserves at most 4ℓ of the q positive coefficient units for the penalty, leaving $q - 4\ell$ for the physical coefficients. With $4\ell < q$, set

$$\delta = \frac{M}{q - 4\ell}, \qquad P = 2\ell\delta, \tag{15}$$

round each physical coefficient θ_F to an integer $k_F = \mathrm{rint}(\theta_F/\delta)$, and add integer pair coefficients $\ell w_a w_b$. All stored coefficients remain within $[-q, q]$. On the original dimensionless energy scale, the resulting Hamiltonian is

$$\widetilde{H}_{\mathrm{aligned}} = F_{h,\mathrm{rounded}} + 2\ell\delta C^2 + \mathrm{constant}. \tag{16}$$

The penalty is exactly constant on $C = 0$ before any subsequent coefficient error. Reserving more range for the penalty increases δ, leaving a coarser grid for physical coefficients. Feasible states with different physical energies can then receive the same compiled energy, admitting inferior states into the minimum-energy set (Fig. 2). A feasible compiled global minimizer has physical energy excess $F_h - F_*$ at most $\delta n(n+1)$, but feasibility itself needs a separate argument.

For each instance, precision and allocation, we compare aligned compilation with conventional rounding at the same effective P in Eq. (15). Conventional normalization may yield a finer step $R/q \leq \delta$ because physical and penalty coefficients can partially cancel. Thus equal bit budgets and penalties do not imply equal physical resolution. The calibrated primary policies can choose different allocations; those policy-level comparisons are distinguished from same-allocation mechanism comparisons.

At a common coefficient step, the aligned construction agrees with direct rounding of the full objective away from nearest-even half ties. Stabilizing half-grid arithmetic removed all matched-grid discrepancies in this benchmark and left every locked-policy ground set unchanged. The comparison therefore concerns coefficient-grid and penalty allocation; implementation details are given in the Supplementary Information.

Held-out fidelity does not improve on average

Each instance is a phase-field problem specified by its grid size, gradient coefficient and spatial bias. The benchmark contains 16 training instances, 16 validation instances and 96 test instances. Training screens penalty allocations and validation selects the final allocation; these sets calibrate the coefficient representation. Each development set contains one instance for each combination of $n \in \{6, 8\}$, $\kappa \in \{0.1, 0.5, 1, 1.8\}$ and bias amplitude $A \in \{0.02, 0.08\}$. The test set contains six independent pairs of angles shifting the spatial bias waves in each of these 16 parameter combinations (Methods). Coefficient budgets are 4, 5, 6, 8, 10 and 12 bits. Allocation candidates come from six prespecified fractions, with duplicate integer allocations removed at low precision. Both methods are calibrated separately at each precision using development data only, then held fixed for testing.

For the complete set G_Q of minimum-energy encoded states, we define composition feasibility and target-energy fidelity as

$$f_C = \frac{|\{s \in G_Q : C(s) = 0\}|}{|G_Q|}, \tag{17}$$

$$f_E = \frac{|\{s \in G_Q : C(s) = 0,\ F_h(s) \le F_* + \varepsilon L\}|}{|G_Q|}, \quad \varepsilon = 10^{-3}. \tag{18}$$

Here $|\cdot|$ counts states. Both fractions give equal weight to every encoded global minimum, retaining all ties: one acceptable state among two tied minima gives $f_E = 1/2$. Target-energy fidelity counts states meeting the prespecified physical tolerance; mean physical regret, reported separately, averages $(F_h - F_*)/L$ over feasible minima. These measures describe the proportion of acceptable states and energy-error size among feasible minima, respectively.

Averaging equally over the six precisions within each instance, conventional rounding achieved mean $f_E = 0.796393$, compared with 0.774972 for aligned compilation. The paired aligned-minus-conventional difference was -0.021420 (95% stratified instance-bootstrap confidence interval (CI) $[-0.032742, -0.010147]$; two-sided random-sign $p = 0.01046$, Holm-adjusted $p = 0.02092$ across two primary outcomes). Aligned compilation therefore has 2.14-percentage-point lower target-energy fidelity in the declared precision mixture.

The precision dependence is informative (Table 1; Fig. 3). Aligned fidelity is lower at 4 and 5 bits, higher at 6 and 8 bits, and equal to conventional fidelity at 10 and 12 bits, where both reach one under the declared tolerance. At eight bits, the separately selected allocations are $\ell = 20$ for conventional rounding and $\ell = 10$ for alignment. At either common allocation, 10 or 20, conventional fidelity is higher (Supplementary Information, Table S11). The eight-bit advantage therefore belongs to the calibrated policy comparison. The six precisions are repeated conditions on the same instances.

Mean composition feasibility was 1.000000 for conventional rounding and 0.999074 for aligned compilation. The paired difference was -0.000926 (95% CI $[-0.002315, 0]$; Holm-adjusted $p = 0.49858$). Thus composition feasibility remains high while target-energy fidelity varies substantially. Mean ground-set multiplicity also increases under alignment. At five bits, alignment has lower fidelity but smaller mean feasible regret, illustrating the distinction between the two measures; Supplementary Table S12 reports the energy-error values at every precision.

To explain this difference, we classified minima within the feasible sector independently of global

feasibility. A configuration is *preserved* when every feasible-sector minimum lies within the stated physical tolerance; a *tie expansion* contains both acceptable and inferior states; a *strict reversal* has no acceptable state among its feasible-sector minima. On the 576 locked test conditions per method, aligned compilation reduced strict reversals from 48 to 10 but increased tie expansions from 110 to 164 (Table 2). Weighting the inferior states by their fraction of each complete ground set connects these events to the primary result: alignment reduces reversal loss by 6.60 percentage points, but adds 8.65 points of tie-expansion loss and 0.09 points of infeasibility loss. The net increase in loss is 2.14 points, accounting for the average fidelity decrease. This decomposition retains the original equal weighting of instances and precisions. Across all 48 conventional strict-reversal conditions, the 260 feasible-sector minima and their physical optima have $Y^2 = 0$. The residual penalty contribution to these comparisons is therefore zero: their ordering changes require the physical-rounding and interaction terms, rather than the pure-penalty preference alone.

Table 1: Target-energy fidelity under the locked policies by signed coefficient precision. Values are means across 96 independent instances with 95% stratified instance-bootstrap confidence intervals (CIs). Each method selects its allocation on development data. At eight bits, conventional/aligned allocations are $\ell = 20/10$; they match at the other precisions. Target-energy fidelity f_E is the fraction of each complete encoded ground set satisfying composition and physical-energy tolerance $10^{-3}L$.

Bits	Conventional rounding	Constraint-aligned compilation
	Mean f_E [95% CI]	
4	0.3838 [0.3591, 0.4100]	0.2525 [0.2306, 0.2752]
5	0.6997 [0.6632, 0.7326]	0.6139 [0.5743, 0.6536]
6	0.7418 [0.7088, 0.7726]	0.8060 [0.7661, 0.8429]
8	0.9531 [0.9323, 0.9722]	0.9774 [0.9601, 0.9931]
10	1.0000 [1.0000, 1.0000]	1.0000 [1.0000, 1.0000]
12	1.0000 [1.0000, 1.0000]	1.0000 [1.0000, 1.0000]

Allocation, tolerance and residual errors limit the mechanism

The complete allocation grid shows that preserving the penalty block does not eliminate the need to balance constraint range against physical resolution (Fig. 4). Increasing the requested penalty allocation can improve one regime while degrading another. At low bit budgets, different requested fractions collapse to the same integer allocation, producing repeated heatmap cells. At fixed instance, precision and allocation, both methods use the same effective penalty, whereas their rounding steps can differ. The primary comparison uses allocations locked on the development data.

We reconstructed relative energy error on the entire feasible sector, subtracting its mean offset before calculating root-mean-square error (RMSE). A global additive offset cannot change ranking, whereas even a small relative error can matter near a physical energy gap. Figure 4c relates this error to target-energy fidelity across the locked configurations. An exploratory tolerance analysis retained the same ground sets and selected parameters. The lower pooled aligned fidelity persists at zero additional tolerance and at $\varepsilon = 0.01$, while the fraction of acceptable states changes with the threshold (Supplementary Table S6).

Residual coefficient errors were evaluated at locked 6- and 8-bit settings using four common Gaussian draws per instance and standard deviations (SDs) 0.05 and 0.25 in each compiler's coefficient-

Table 2: Ordering events and contributions to target-energy fidelity loss. Counts refer to 576 instance–precision conditions per method, from 96 independent instances. Exact and near-optimal classifications use tolerances 10^{-10} and $0.001L$, respectively. A strict reversal excludes acceptable states from the feasible-sector minima; a tie expansion retains both acceptable and inferior states. Loss contributions weight each condition equally and sum to $100(1-\overline{f}_E)$. The final column is aligned minus conventional, calculated before rounding; it reports condition-count differences in the upper sections and percentage-point differences in the loss section.

Event or loss component	Conventional	Aligned	Difference
Exact optimum: condition counts			
Preserved	353	332	-21
Tie expansion	163	221	+58
Strict reversal	60	23	-37
Near optimum: condition counts			
Preserved	418	402	-16
Tie expansion	110	164	+54
Strict reversal	48	10	-38
Global feasibility: condition counts			
All global minima feasible	576	574	-2
Mixed feasible/infeasible minima	0	2	+2
No global minimum feasible	0	0	0
Near-optimal fidelity loss: percentage points			
Inferior feasible states: strict reversal	8.33	1.74	-6.60
Inferior feasible states: tie expansion	12.03	20.67	+8.65
Infeasible states	0.00	0.09	+0.09
Total loss	**20.36**	**22.50**	**+2.14**

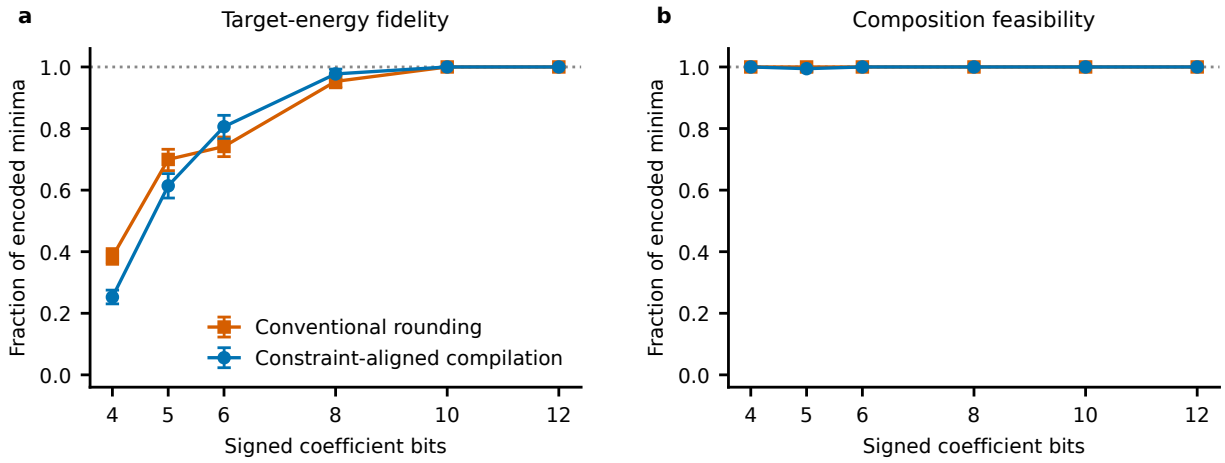


Figure 3: Feasibility and physical fidelity separate on held-out instances. Fractions of each complete encoded ground set satisfying **a**, exact composition and physical-energy tolerance $10^{-3}L$ (interval length L), or **b**, composition alone. Markers are means across 96 independent test instances; bars are 95% stratified instance-bootstrap intervals from 10,000 resamples within 16 cells. Both methods are separately calibrated on development data and locked before testing. At eight bits, conventional/aligned allocations are $\ell = 20/10$; they match at the other precisions. The six precisions are repeated conditions on the same instances. Dotted lines indicate a counting fraction of one.

grid units (Fig. 5). These errors perturb every field and pair slot, including zero entries, and can destroy the exact penalty-block property after compilation. The corresponding fidelity curves are shown in Fig. 5, with fidelity and feasibility values in Supplementary Table S5. A small perturbation can resolve existing ties and thereby increase the counting fraction. The synthetic errors have equal standard deviations in grid units; their physical-energy scales differ when the compiler steps differ. As with analog-error studies more generally [17], ideal optimization of the perturbed coefficients need not solve the intended problem.

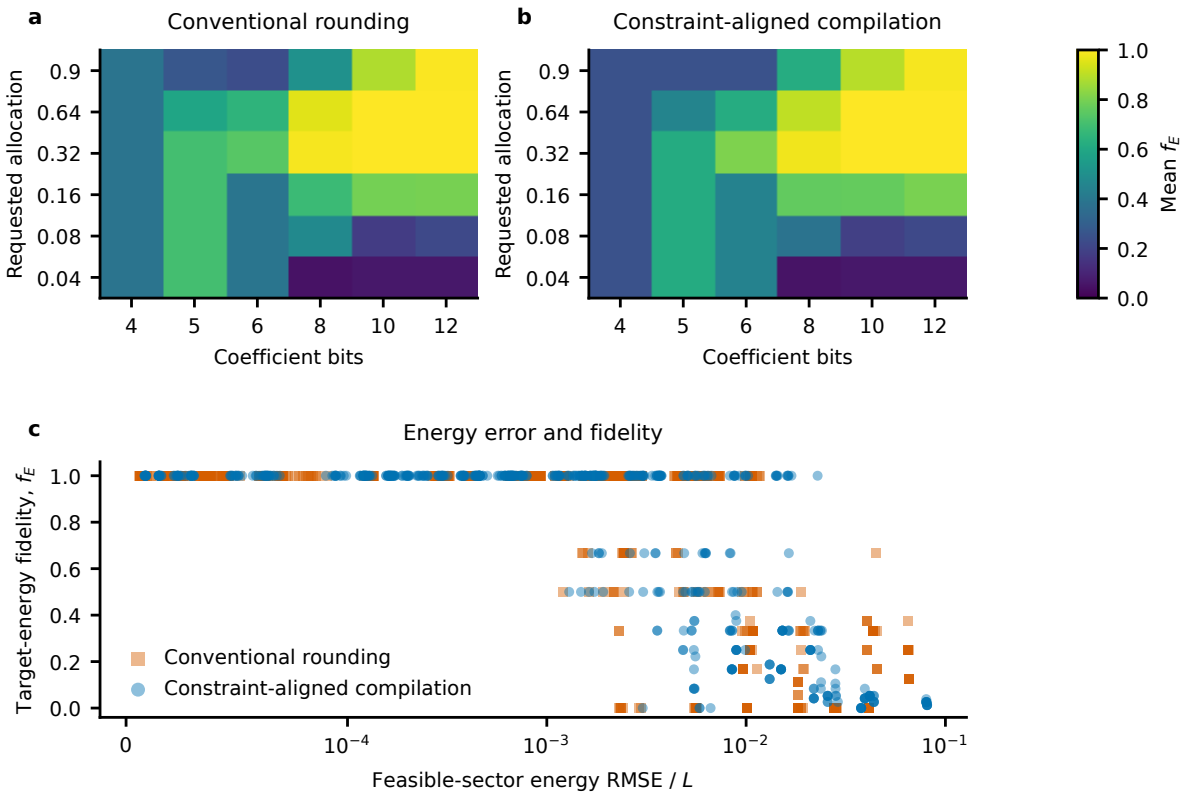


Figure 4: Penalty preservation and coefficient-range allocation. a,b, Mean target-energy fidelity f_E across 96 test instances over the requested-allocation grid. Fidelity is the fraction of each complete ground set satisfying composition and physical-energy tolerance $10^{-3}L$. Same-allocation method pairs use equal effective penalties but may have different rounding steps. Requested fractions can map to the same integer allocation at low precision, giving repeated cells. Both heatmaps use the same colour scale. **c**, Feasible-sector energy root-mean-square error (RMSE), after subtracting its mean offset and dividing by interval length L, versus f_E. Each method contributes 576 locked cases: 96 instances at six precisions. The horizontal scale is symmetric-logarithmic, with a linear region below 10^{-4}.

Continuum references expose a separate representation limit

We separate two questions: which states are excluded by the four-level restriction, and how the energy evaluation changes for the same grid-point values. For a field varying continuously in

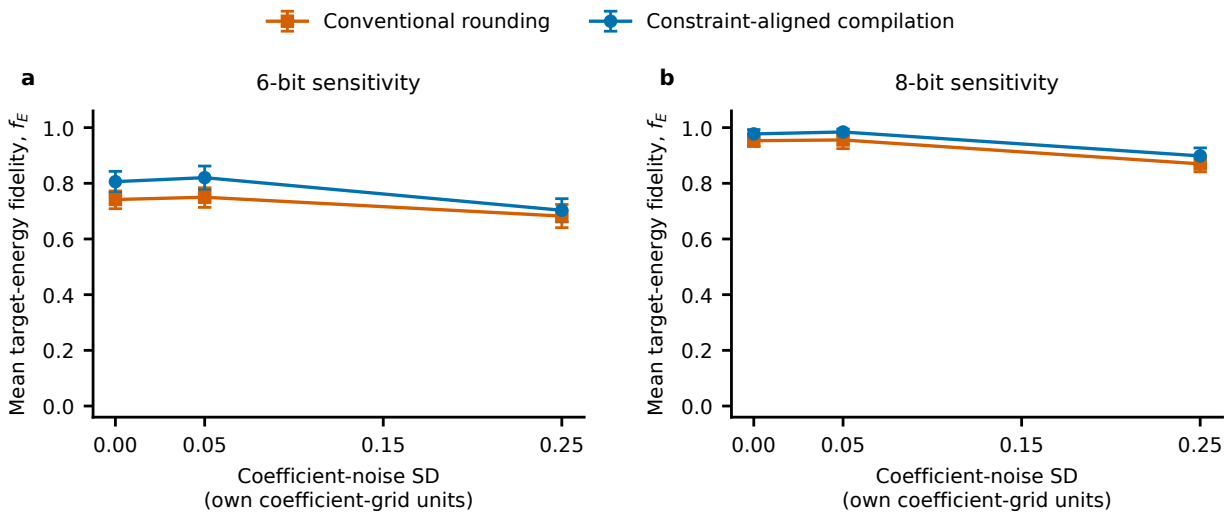


Figure 5: Sensitivity to synthetic coefficient errors. Target-energy fidelity f_E at locked **a**, 6-bit and **b**, 8-bit settings under independent Gaussian errors in every linear and quadratic coefficient slot, including zeros. Fidelity uses physical-energy tolerance $10^{-3}L$ (interval length L) and exact composition. Zero-error points are the deterministic results. At each nonzero standard deviation (SD), four noise draws shared between methods are averaged within each of 96 instances; bars show descriptive 95% stratified instance-bootstrap intervals. Errors use each compiler's coefficient-grid units, so equal plotted SD can represent different physical-energy perturbations.

space, we evaluate

$$\mathcal{F}[\phi] = \int_0^L \left[\frac{(\phi^2 - 1)^2}{4} + \frac{\kappa}{2}(\partial_x \phi)^2 + v(x)\phi \right] dx, \qquad \int_0^L \phi \, dx = 0. \tag{19}$$

This continuum functional integrates over the whole domain; the Fourier reference below imposes no pointwise bound on ϕ. A separate continuous nodal reference retains the grid and F_h but allows any $\phi_i \in [-1, 1]$ (Methods), isolating the restriction to four local values. Eight fixed audit cases use the four κ values and $A \in \{0, 0.08\}$, with one shared phase pair distinct from the primary instances. Fourier–Galerkin minimization represents the field by sine and cosine waves and uses 15, 31 and 63 nonzero modes, and respectively 64, 128 and 256 quadrature nodes. Multistart optimization with the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with bounds (L-BFGS-B) [18], with prolongation between resolutions, produced 208 retained local optimization runs. The resulting fields are stationary references. The largest change in best energy from 31 to 63 modes was 9.66×10^{-14}, and the largest finest-reference chemical-potential residual was 1.47×10^{-7}. These resolution checks do not generally certify a global continuum minimum.

For every feasible four-level eight-site field, we also evaluated the continuum functional of its periodic piecewise-linear interpolation ϕ_{P1}, formed by joining neighboring grid values with straight lines, exactly, integrating the quartic cell potential and smooth forcing analytically. Nodal and piecewise-linear-functional minimizer sets intersected in all eight audit cases. Thus these cases show no observed minimizer reversal from nodal quadrature. They do show value discrepancies: at the selected nodal optimum, $F_h - \mathcal{F}[\phi_{\mathrm{P1}}]$ ranged from -0.280294 to -0.009877 (Fig. 6).

The high-κ unforced case gives an exact representation counterexample. The second variation

measures energy curvature under a small field change. For a mean-zero perturbation η, which preserves total composition, it obeys

$$\delta^2\mathcal{F}[\phi](\eta,\eta) \geq \left[\kappa(2\pi/L)^2 - 1\right] \int_0^L \eta^2 dx. \tag{20}$$

Hence the fixed-mean continuum functional is strictly convex for $\kappa > (L/2\pi)^2 = 1.621139$. At $\kappa = 1.8$ and $A = 0$, its unique minimum is $\phi = 0$, corresponding to equal local fractions of A and B everywhere, with energy 2. The eight-site continuous nodal model is also strictly convex there, since its threshold is $h^2/[4\sin^2(\pi/8)] = 1.707107$. Yet the four-level alphabet excludes zero. Its exact nodal minimum is 2.380247 and its interpolated continuum energy is 2.449383. Even a perfect unquantized Ising solution cannot recover the excluded uniform field. This establishes a finite-alphabet representation limit.

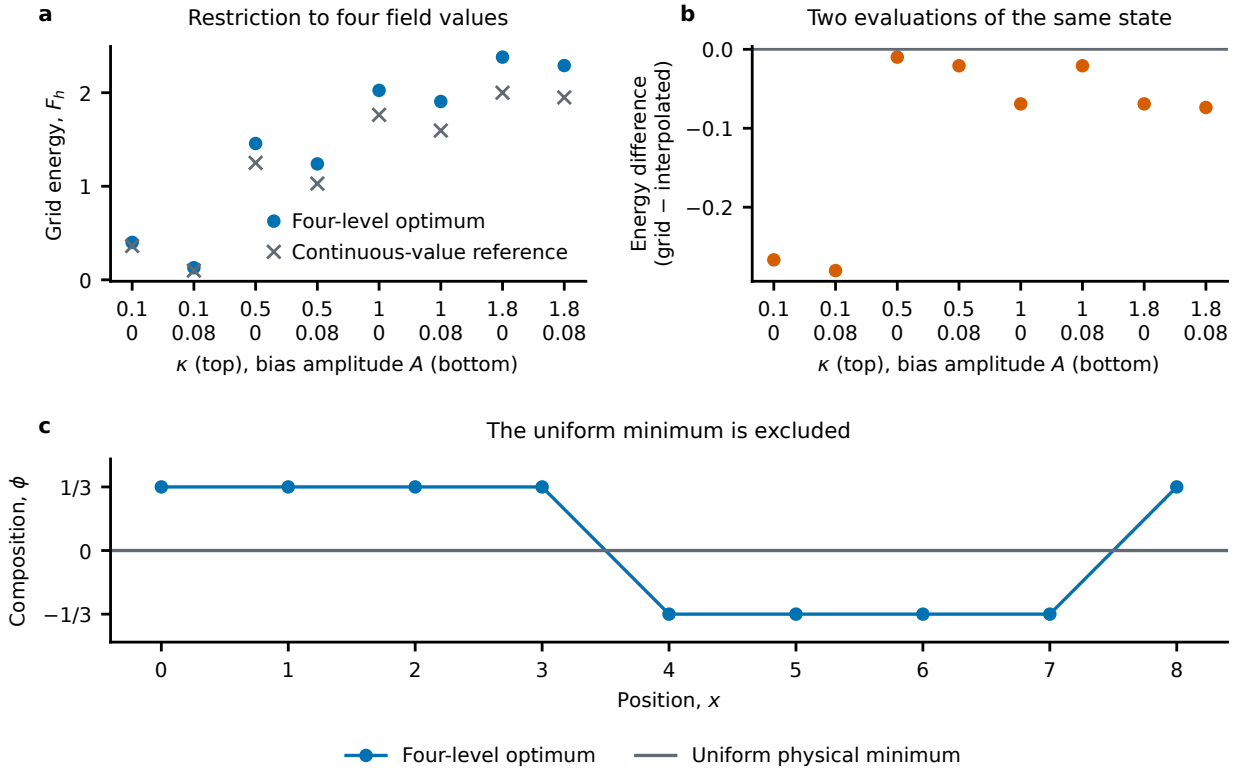


Figure 6: Field-value restrictions and energy evaluation are distinct comparisons. a, Exact four-level minima and continuous-value references for eight fixed cases, evaluated using the same nodal energy F_h. Continuous-value references allow grid-point values in $[-1, 1]$ and are the best of nine local starts, including a discrete-optimum warm start (Methods); global optimality is not generally certified. **b**, Nodal energy minus the continuum functional of the piecewise-linear interpolation of the same selected discrete minimum. **c**, At $\kappa = 1.8$ and zero bias, the unique continuous mean-zero minimum is uniform and excluded by the four-level alphabet. The displayed discrete state has nodal energy 2.38025 and interpolated energy 2.44938, compared with 2 for the uniform field. All values are dimensionless deterministic calculations.

Discussion

Coefficient quantization changes physical-state selection through artificial coupling preferences and lost energy distinctions. The pure composition penalty favors unequal populations of saturated field levels, while the complete objective adds physical rounding and a joint-rounding interaction. The validation example traces a reversal to specific on-site and neighboring encoding-spin couplings. All conventional strict reversals under the locked policies occur between states with identical residual penalty energy, showing why the full objective is needed to explain which structures are selected.

Protecting the penalty changes both strict reversals and ties. Integer alignment reduces the reversal contribution to loss, but its coarser physical grid admits enough inferior ties to lower mean target-energy fidelity across the prescribed precision mixture. Fidelity counts the proportion of acceptable minima; physical regret describes energy-error size among feasible minima. Their rankings can differ when a compilation admits more inferior minima with smaller energy errors. The eight-bit results further show why separately calibrated policies and same-allocation comparisons answer different questions. Applications at a fixed precision should evaluate the corresponding policy, tie multiplicity and physical-energy errors.

Relative to optimum-preserving compression and the analysis of quantization-induced ties [8, 12], the present results identify the field populations and couplings that drive state selection and quantify how reversal, tie and infeasibility losses combine. This physical account complements penalty selection and coefficient conditioning [7, 10]. Finite resolution in optical feedback [19] and digital solver states and updates [20] affects dynamics; the present analysis isolates changes in the static programmed objective.

The continuum comparison identifies a further source of error: the four-level alphabet excludes the uniform zero field favored at high gradient cost. Accurate compilation preserves the finite-state target, while representation accuracy depends on which physical fields that target can express. In all eight reference cases, the grid-based energy and the continuum energy of the interpolated field share at least one minimizing configuration, although their energy values differ. Each comparison therefore requires its own physical objective and admissible set.

Complete enumeration makes the state-selection analysis exact in these small one-dimensional systems. The policy comparison uses one development-stage calibration, and the sensitivity tests use synthetic coefficient-grid errors. These settings define the present evidence; hardware predictions additionally depend on programming errors and sampling weights over tied states. For finite-precision formulations, the practical assessment is to verify the encoding, compare composition feasibility with target-energy fidelity and energy-error size, and check whether the alphabet contains the required physical states.

Methods

Instance generation and exact objectives

Equation (1) discretizes the continuum energy in Eq. (19). Spatial integration of the local and bias terms uses $\int g(x)\,dx \approx h\sum_i g(x_i)$, while $\partial_x\phi$ is approximated by $(\phi_{i+1}-\phi_i)/h$. Squaring this difference and multiplying by the integration weight h gives the coefficient $\kappa/(2h)$. On a uniform

grid, the integrated composition constraint $h\sum_i \phi_i = 0$ is equivalent to the sum constraint in Eq. (1).

Two independent angles θ_1, θ_2, uniform on $[0, 2\pi)$, shift the spatial components of the bias:

$$v(x) = \frac{A}{1.5}\left[\cos(2\pi x/L + \theta_1) + \frac{1}{2}\cos(4\pi x/L + \theta_2)\right], \tag{21}$$

Here A bounds the bias magnitude, $|v(x)| \le A$. We sample at $x_i = (i-1)h$, $i = 1, \ldots, n$, matching the zero-based indexing in the implementation. Random streams derive from SeedSequence namespace $(42, 20261014, 3)$ with distinct split and instance identifiers. The parameter combinations and independent sample counts are specified in Results. Every one of the 2^{2n} binary configurations is enumerated: 4,096 for six sites and 65,536 for eight sites. Direct evaluation of Eq. (1), separately from its Ising coefficients, verifies the mapping. Ground sets retain all ties; the integer compiled energies are exactly representable in the employed floating-point summations at these sizes. A tolerance 10^{-9} is used to collect equal encoded minima, and 10^{-10} to identify physical degeneracy. No instance is removed, and no candidate is repaired after decoding.

Compilation, calibration and cost

The prespecified requested allocation fractions are $\alpha \in \{0.04, 0.08, 0.16, 0.32, 0.64, 0.9\}$. For each precision, ℓ is the nearest-even integer to $\alpha q/4$, clipped to $[1, \lfloor (q-1)/4 \rfloor]$. Duplicate values are removed, yielding 1, 3, 4, 6, 6 and 6 candidates at 4, 5, 6, 8, 10 and 12 bits. Both compilers use the same candidate list and the same physical instance. They are evaluated at matched effective $P = 2\ell M/(q - 4\ell)$ in the full grid.

Training candidates are ranked by mean $1 - f_E$, breaking ties by smaller ℓ. Only the best up to three training candidates are eligible for selection on validation loss, followed by training loss and smaller ℓ as tie breakers. All validation candidates were evaluated, although only the training finalists were eligible for selection. There is no test-based allocation choice. The selected conventional/aligned allocations are respectively 1/1, 1/1, 2/2, 20/10, 41/41 and 164/164 as precision increases. This difference at eight bits means the locked policies do not use identical effective penalties there.

There are 52 method–precision–allocation configurations per instance, producing 832 rows in each development split and 4,992 in the full test grid. The locked primary table contains 1,152 rows (96 instances, six precisions, two methods). The residual-error study adds 3,072 configurations. These counts refer to exhaustive evaluations of compiled objectives, not repeated runs of a physical solver. The $O(2^{2n})$ evaluation cost and dense penalty connectivity preclude a scalability claim. Range normalization, coefficient rounding, objective enumeration and local continuum minimization are reported as distinct operations; software elapsed times are not interpreted as device timings.

Ground-set and energy diagnostics

Fidelity and feasibility follow Eq. (18). Conditional physical regret averages $(F_h - F_*)/L$ only over feasible encoded ground states; when none is feasible it remains missing rather than being replaced by zero. Worst feasible regret, tie multiplicity, interfaces, field root-mean-square amplitude and composition-sum error are retained in the source tables. The Supplementary Informa-

tion summarizes mean and worst feasible regret by precision alongside fidelity. The raw variable named `mean_absolute_mass` denotes mean $|\sum_i \phi_i|$ over the ground set; integrated mass error would multiply this by h. An optional field-distance diagnostic uses the first enumerated physical optimum as its reference, not the minimum distance over a degenerate optimum set.

The relative feasible-sector energy RMSE is the standard deviation of $\widetilde{H} - F_h$ over all feasible configurations, divided by L. Subtracting the mean removes the arbitrary additive constant. Tolerance sensitivities at 0, 10^{-4} and 10^{-2} are exploratory reconstructions from saved ground sets, with the primary 10^{-3} threshold and parameters unchanged. Residual errors add independent Gaussian coefficients in grid units to all linear and quadratic slots. Four noise draws per instance, precision and SD are shared between methods. The deterministic rounding certificate is not applied to these subsequently perturbed coefficients.

For each global ground set, the loss $1 - f_E$ is the sum of its infeasible fraction and its feasible fraction outside the physical-energy tolerance. We assign the latter to strict reversal or tie expansion according to the feasible-sector ordering class. Averaging these fractions over the original 96 instances and six precisions gives the loss contributions in Table 2; event counts alone would give different weights to tied states. The Supplementary Information states the per-configuration identity and reports the unrounded contributions.

State-level mechanism reanalysis

We reanalyzed the saved 16 validation and 96 test instances without generating new instances, selecting new allocations or training a model. Each has 52 method–precision–allocation configurations, giving 5,824 reconstructions. We independently evaluated the direct physical energy, the fixed-grid terms in Eq. (11), feasible-sector minima and global minima. Classification was repeated for exact physical optima (absolute numerical tolerance 10^{-10}) and for the primary energy threshold $10^{-3}L$. Every compiled tie was retained. The feasible-sector ordering classes were crossed with global feasibility (all, some or no global minima feasible), rather than treating constraint violation as an alternative ordering class.

Before full reanalysis, we fixed the illustration rule: sort validation configurations by precision, instance identifier and allocation, and take the first strict reversal, tie expansion and matched-allocation case in which alignment lowers fidelity. Within each selected case, we display the physically worst feasible-sector minimum, breaking equal-energy ties by state identifier, alongside the exact physical optimum with the lowest compiled energy, again breaking ties by state identifier. Test outcomes did not select the displayed cases. All test allocations remain available, but the main summary uses the original locked policies. Reconstructed coefficients, ground-state identifiers and original metrics were checked against the archive. Supplementary Information gives the coefficient-level definitions, state-energy components and numerical residuals. The original nearest-even implementation is retained. Reanalysis of its archived half-grid stabilization check reproduced unchanged locked ground sets, ordering classifications and fidelity; the Supplementary Information distinguishes ideal half-grid arithmetic from the implemented coefficient values.

Independent continuum and nodal references

The continuum audit uses phases $(4.862909272689599, 2.757554564287996)$ generated from seed 42. Zero-mean fields are represented as sums of sine and cosine waves; omitting the constant

term imposes zero total composition. At mode cutoffs $K = 15, 31, 63$, quadrature counts $N = 64, 128, 256$ obey $N > 4K$, resolving the quartic integrand. Gradient energy is integrated analytically. The L-BFGS-B implementation uses analytic gradients, tolerance 10^{-10} on the gradient criterion, function tolerance 10^{-15}, at most 12,000 iterations, 50 line-search steps and 30 correction pairs. Eight prescribed initial fields include first- and second-harmonic profiles and smooth random Fourier fields; each finer level also starts from the prolonged previous best. All 208 optimization runs are retained. Success flags are supplemented by achieved gradient and chemical-potential residuals because the function criterion can terminate a run first.

For a linear interpolation between neighboring grid values a, b, the exact quartic potential integral over one interval of length h is

$$h\left[\frac{a^4 + a^3b + a^2b^2 + ab^3 + b^4}{20} - \frac{a^2 + ab + b^2}{6} + \frac{1}{4}\right]. \tag{22}$$

The gradient integral is $\kappa(b - a)^2/(2h)$ and the product with each cosine bias is integrated analytically. Independent Gaussian quadrature verifies this evaluator. All 8,092 feasible eight-site four-level fields are scored under both the nodal and interpolated continuum objectives, retaining minimizer-set ties.

The continuous nodal diagnostic minimizes the same F_h with $-1 \le \phi_i \le 1$ and zero composition, using sequential least-squares programming (SLSQP) and analytic derivatives. An initial eight-start calculation missed a lower known discrete solution in one case. A documented diagnostic amendment therefore added the exact discrete optimum as a ninth start in every case. All 72 final nodal runs and the original eight-start subsets are retained. This warm-start amendment changes neither primary compiler selection nor the Fourier-reference protocol. These nodal references remain numerical local references except where the stated convexity argument applies.

Statistical analysis

Primary outcomes are first averaged equally over six precisions within each instance. The aligned-minus-conventional differences are then averaged across 96 paired instances. Confidence intervals use 10,000 percentile-bootstrap draws, sampling six instances with replacement within each of the 16 site–gradient–bias strata; root seed is 42. Two-sided random-sign tests use 99,999 draws and a plus-one correction, with Holm adjustment over the fidelity and feasibility contrasts. The null assumes sign exchangeability of paired differences, not randomized method assignment. The bootstrap preserves the declared stratum mixture; it and the random-sign test need not yield identical inferential summaries. No prospective power calculation was specified. Precision-specific, noise and tolerance results are descriptive, without additional primary tests. Noise draws are averaged within instance before estimating intervals. Pure-penalty counts and the eight continuum cases receive no population p-values.

The local protocol was saved before primary simulations but was not externally preregistered. A pretest clarification records validation computational accounting and the interpretation of raw diagnostic fields without changing parameters or selection. Half-boundary stabilization is a post hoc numerical sensitivity analysis with unchanged selected allocations, penalties and grids; it does not replace the primary nearest-even implementation or introduce another primary test. Independent checks reconstruct physical energies, every encoded ground-set tie and selection choices. Source

code, complete coefficient arrays, ground-state identifiers, calibration records, analysis tables and reference optimizations are archived. All observations are simulated.

Use of generative artificial intelligence

OpenAI Codex assisted with code editing and debugging, numerical consistency checks, manuscript language revision, preparation of supplementary materials, and typesetting. The authors remain responsible for the analyses, interpretations and completeness of this disclosure.

Data availability

All synthetic data supporting this study are provided in the accompanying reproducibility package, `Reproducibility_Package.zip`. These include phase-field parameters, compiled coefficient vectors, complete encoded ground-state identifiers, instance-level metrics, calibration records, analytical-example data, Fourier coefficients and nodal-reference results. All data were generated computationally; no measurements or personal data were used.

Code availability

The accompanying reproducibility package contains the Python code, dependency information and execution instructions needed to reproduce the encoding, exhaustive evaluation, calibration, statistical analysis and continuum-reference calculations. It also contains the coupling analysis, complete-ground-set loss decomposition and same-allocation comparison. The reproducibility package additionally contains the current figure-rendering code and display inputs.

References

[1] Cahn, J. W. & Hilliard, J. E. Free Energy of a Nonuniform System. I. Interfacial Free Energy. *J. Chem. Phys.* **28**, 258–267 (1958). URL https://doi.org/10.1063/1.1744102.

[2] Endo, K., Matsuda, Y., Tanaka, S. & Muramatsu, M. A phase-field model by an Ising machine and its application to the phase-separation structure of a diblock polymer. *Sci. Rep.* **12**, 10794 (2022). URL https://doi.org/10.1038/s41598-022-14735-4.

[3] Aoki, S. *et al.* Formulation of Correction Term in QUBO Form for Phase-Field Model. *Int. J. Numer. Methods Eng.* **126**, e70019 (2025). URL https://doi.org/10.1002/nme.70019.

[4] Kitai, K. *et al.* Designing metamaterials with quantum annealing and factorization machines. *Phys. Rev. Res.* **2**, 013319 (2020). URL https://doi.org/10.1103/physrevresearch.2.013319.

[5] Sandt, R., Le Bouar, Y. & Spatschek, R. Quantum annealing for microstructure equilibration with long-range elastic interactions. *Sci. Rep.* **13**, 6036 (2023). URL https://doi.org/10.1038/s41598-023-33232-w.

[6] Verma, A. & Lewis, M. Penalty and partitioning techniques to improve performance of QUBO solvers. *Discrete Optim.* **44**, 100594 (2022). URL https://doi.org/10.1016/j.disopt.2020.100594.

[7] Alessandroni, E., Ramos-Calderer, S., Roth, I., Traversi, E. & Aolita, L. Alleviating the quantum Big-M problem. *npj Quantum Inf.* **11**, 125 (2025). URL https://doi.org/10.1038/s41534-025-01067-0.

[8] Mücke, S., Gerlach, T. & Piatkowski, N. Optimum-preserving QUBO parameter compression. *Quantum Mach. Intell.* **7**, 1 (2025). URL https://doi.org/10.1007/s42484-024-00219-3.

[9] Gerlach, T. & Piatkowski, N. Dynamic Range Reduction via Branch-and-Bound. In *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 119–129 (2025). URL https://doi.org/10.1109/qce65121.2025.00023.

[10] Mirkarimi, P., Hoyle, D. C., Williams, R. & Chancellor, N. Experimental demonstration of improved quantum optimization with linear Ising penalties. *New J. Phys.* **26**, 103005 (2024). URL https://doi.org/10.1088/1367-2630/ad7e4a.

[11] Shirai, T. & Togawa, N. Spin-Variable Reduction Method for Handling Linear Equality Constraints in Ising Machines. *IEEE Trans. Comput.* **72**, 2151–2164 (2023). URL https://doi.org/10.1109/tc.2023.3239539.

[12] Zeng, Z. *et al.* Extractive Summarization on a CMOS Ising Machine. *IEEE Access* **14**, 56724–56735 (2026). URL https://doi.org/10.1109/access.2026.3681908.

[13] Chen, L. Q. & Shen, J. Applications of semi-implicit Fourier-spectral method to phase field equations. *Comput. Phys. Commun.* **108**, 147–158 (1998). URL https://doi.org/10.1016/s0010-4655(97)00115-x.

[14] Chancellor, N. Domain wall encoding of discrete variables for quantum annealing and QAOA. *Quantum Sci. Technol.* **4**, 045004 (2019). URL https://doi.org/10.1088/2058-9565/ab33c2.

[15] Chen, J., Stollenwerk, T. & Chancellor, N. Performance of Domain-Wall Encoding for Quantum Annealing. *IEEE Trans. Quantum Eng.* **2**, 1–14 (2021). URL https://doi.org/10.1109/tqe.2021.3094280.

[16] Berwald, J., Chancellor, N. & Dridi, R. Understanding domain-wall encoding theoretically and experimentally. *Phil. Trans. R. Soc. A* **381**, 20210410 (2023). URL https://doi.org/10.1098/rsta.2021.0410.

[17] Albash, T., Martin-Mayor, V. & Hen, I. Analog errors in Ising machines. *Quantum Sci. Technol.* **4**, 02LT03 (2019). URL https://doi.org/10.1088/2058-9565/ab13ea.

[18] Byrd, R. H., Lu, P., Nocedal, J. & Zhu, C. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM J. Sci. Comput.* **16**, 1190–1208 (1995). URL https://doi.org/10.1137/0916069.

[19] Sevenants, T., Van der Sande, G. & Verschaffelt, G. Requirements on bit resolution in optical Ising machine implementations with analog spin variables. *Commun. Phys.* **8**, 11 (2025). URL https://doi.org/10.1038/s42005-024-01919-9.

[20] Yao, B. *et al.* Precision meets speed through an FPGA-based natively sparse Ising machine for combinatorial optimization. *Nat. Commun.* **17**, 8549 (2026). URL https://doi.org/10.1038/s41467-026-75119-0.

Acknowledgements

All figures were prepared by the authors using Matplotlib 3.10.7 (https://matplotlib.org/).

Funding

This work received no specific funding.

Author contributions

Z.L. conceived the study, wrote the initial draft, supervised the research and provided resources. A.L. conducted the numerical experiments and analyzed the results. Z.Z. developed the supporting software, conducted the literature and analytical investigations, and substantially revised and restructured the manuscript. All authors contributed to review and editing.

Additional Information

Competing interests. The authors declare that they have no conflict of interest.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.pdf
- ReproducibilityPackage.zip