ADAPTIVE CONFIDENCE NEURAL LOGIC MACHINES

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantum annealing offers a promising approach to Boolean satisfiability (SAT); however, existing encodings frequently require large numbers of qubits and produce energy landscapes with narrow spectral gaps, increasing susceptibility to spurious minima. This work introduces Adaptive Confidence Neural Logic Machines (AdaptCNLM), a neural-symbolic framework designed to jointly reduce the representation size of SAT instances and enhance the resulting optimization landscape. AdaptCNLM initially employs neural structural-equivalence encoding (NSEE), a label-free graph neural network that identifies and aggregates structurally equivalent variables while provably preserving satisfying assignments within the induced subcube. Subsequently, dual-confidence scaling is applied, integrating structural and search-derived signals to recover information lost during aggregation and to guide the search toward satisfying assignments. Across five quantum-annealing encodings, NSEE significantly reduces qubit requirements and achieves the highest solve rate with fewer logical qubits. As a complete SAT solver, AdaptCNLM surpasses 16 data-driven and non-data-driven baselines across six SATLIB benchmark families. These findings indicate that learned structural compression, combined with confidence-guided search, enhances both resource efficiency and solving performance in quantum-annealing-based SAT.

1 INTRODUCTION

The Boolean satisfiability (SAT) problem, grounded in classical propositional logic, is among the longest-studied problems in formal reasoning. This problem class enables automated problem proving (Holden et al., 2021) and a diversity of industrial applications, including formal verification (Ganai & Gupta, 2007; Maric, 2008), electronic design automation (Marques-Silva & Sakallah, 2000), automated planning (Ernst et al., 1997), cryptography & security (Soos et al., 2009; Mironov & Zhang, 2006; Massacci & Marraro, 2000; Lafitte, 2018), bioinformatics & computational biology (Lynce & Marques-Silva, 2006; Corblin et al., 2010), and optimization & management (Gu, 1994; Alagarsamy & Gopalakrishnan, 2016). Studying SAT is also essential for advancing neural symbolic artificial intelligence (Selsam et al., 2018; Amizadeh et al., 2018; Ozolins et al., 2022; Colelough & Regli, 2025).

Quantum computing has emerged as a promising platform for SAT solvers (Douglass et al., 2015; Bian et al., 2020; Campos et al., 2021; Oshiyama & Ohzeki, 2022; Mirkarimi et al., 2023; Zhang et al., 2025; Fong et al., 2025), with most solvers developed on quantum annealers (Bian et al., 2017; 2020; Douglass et al., 2015; Pudenz et al., 2017; Mehta et al., 2022; Nüßlein et al., 2023; Zielinski et al., 2024; Kanatbekova et al., 2025; Mehta et al., 2025; Ansótegui & Levy, 2024). In this framework, an instance is mapped to an Ising Hamiltonian or Quadratic Unconstrained Binary Optimization (QUBO) model (Lucas, 2014) via penalty functions that separate satisfying from non-satisfying assignments by an energy gap. Because annealer hardware supports only quadratic interactions, higher-order clauses must be reduced into hardware-compatible quadratic forms (Bian et al., 2020; Pudenz et al., 2017; Ansótegui & Levy, 2024; Bybee et al., 2023).

Two critical challenges confront quantum SAT solvers. The first is qubit overhead. The encoded Hamiltonian must be physically realized through problem encoding, in which the problem's interaction graph is mapped onto the sparse, fixed-topology hardware graph of the annealer (Bunyk et al., 2014; Boothby et al., 2020; 2021). State-of-the-art encoders already require $(n + m)$ logical qubits (Chancellor et al., 2016; Nüßlein et al., 2023), with one qubit allocated per Boolean variable and one per clause, where n and m denote the numbers of variables and clauses in the input formula,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

respectively. This raises the question of how the qubit count can be reduced while preserving the energy separation between satisfying and non-satisfying assignments.

The second challenge is the complex structure of the search landscape, which may contain numerous local optima, saddle points, and flat regions that hinder effective exploration of optimisers. Once embedded, annealing performance is governed by the minimum spectral gap between the ground state and the first excited state of the evolving Hamiltonian (Bybee et al., 2023). A small gap forces the device to evolve more slowly to maintain adiabaticity, drastically increasing runtime, and even then reduces the probability of ending in the ground state at the conclusion of the schedule (Farhi et al., 2000; Kirkpatrick et al., 1983). It is therefore essential to design the energy function so as to both widen this gap and suppress spurious local minima that can trap the dynamics in metastable excited states.

In this work, we address these research gaps by proposing Adaptive Confidence Neural Logic Machines (AdaptCNLM), a neural SAT solver that jointly reduces qubit overhead through learned problem compression and improves annealing landscape quality through confidence-based energy reweighting. Our contributions are as follows:

1. We propose neural structural-equivalence encoding (NSEE), a label-free graph neural network that learns to group structurally equivalent variables, compressing the original CNF formula into a reduced instance with fewer variables and directly lowering the logical-qubit footprint of quantum SAT solving.

2. We introduce dual-confidence scaling of the problem-induced energy landscape, a reweighting scheme that combines static structural signals with search-derived evidence to compensate for the information discarded by aggregation, coupling with NSEE to form a single end-to-end solver.

In the numerical evaluation, NSEE requires substantially fewer logical and physical qubits than five established quantum-annealer encodings while achieving the highest solve rate, and AdaptCNLM outperforms sixteen data-driven and non-data-driven SAT solvers across six SATLIB families and two SAT Competition instances.

2 Related Works

Data-driven SAT solvers. Learning-based approaches to SAT can be categorised according to the type of output they are designed to produce. The first category predicts satisfiability or an assignment end to end. NeuroSAT (Selsam et al., 2018) trains a message-passing network over the literal-clause graph from single-bit supervision and recovers assignments by clustering literal embeddings, and subsequent architectures replace the supervision with a differentiable relaxation of the formula (Amizadeh et al., 2018; 2019) or with an explicitly goal-aware query mechanism (Ozolins et al., 2022). G4SATBench (Li et al., 2023) standardises this family and reports that the graph architecture matters less than the training distribution. The second category learns to guide a solver rather than to replace it, supplying branching heuristics, restart policies or initial assignments to a CDCL or local-search backbone. The third category learns an energy model, either as a restricted Boltzmann machine sampled thermally (Warde-Farley et al., 2023) or as a differentiable surrogate of the clause count (Amin et al., 2018).

Energy-based approaches for SAT. Casting SAT as energy minimization underlies both classical and quantum solvers. On the classical side, simulated annealing (Kirkpatrick et al., 1983) remains the canonical thermal baseline, accepting uphill moves probabilistically via the Metropolis criterion. The quantum route encodes clause penalties as an Ising Hamiltonian (Lucas, 2014), whose ground state corresponds to a satisfying assignment. Higher-order clauses must first be reduced to quadratic form (Bybee et al., 2023; Tseitin, 1983), and performance is governed by the minimum spectral gap of the evolving Hamiltonian (Farhi et al., 2000). Quantum-inspired heuristics such as simulated quantum annealing (Kadowaki & Nishimori, 1998) and digital QUBO solvers (Oshiyama & Ohzeki, 2022; Fong et al., 2025) approximate tunnelling dynamics classically but remain sensitive to landscape geometry.

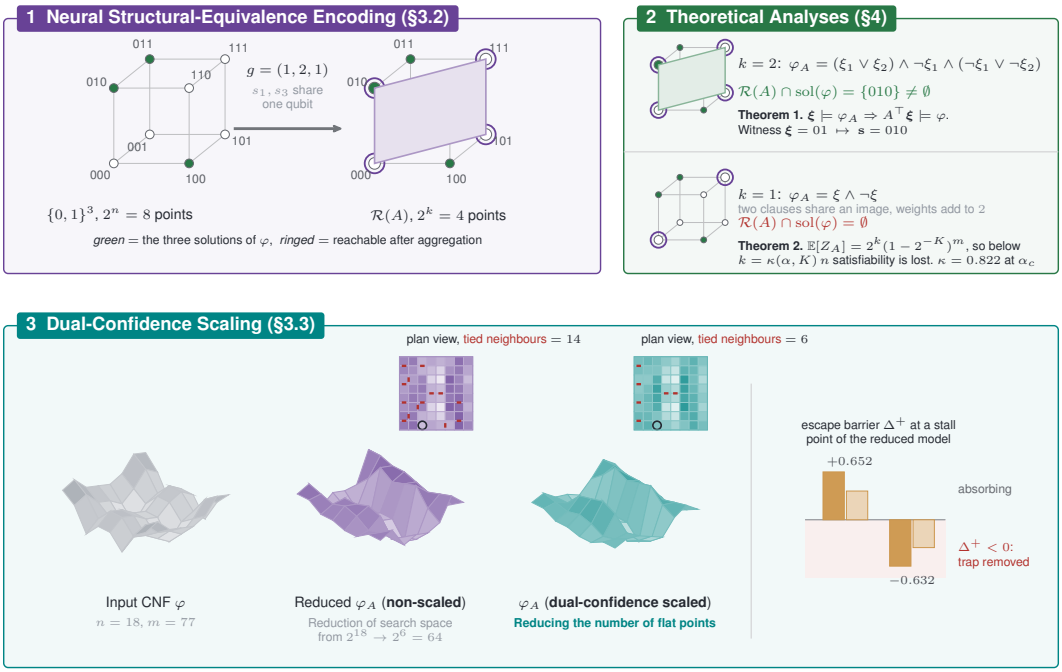


Figure 1: **Illustration of AdaptCNLM. (1)** Aggregation by NSEE restricts the search space from $\{0,1\}^n$ to the subcube $\mathcal{R}(A) = \{A^\top \boldsymbol{\xi}\}$. On $\varphi = (s_1 \vee s_2) \wedge (\neg s_1 \vee \neg s_3) \wedge (\neg s_1 \vee \neg s_2)$, the grouping merges s_1 and s_3, taking 8 points to 4. **(3)** Surfaces are exact enumerations laid on a gray grid, so every horizontal and vertical step is a single-bit move on the cube. The non-scaled reduced surface (purple) is terraced, with 14 of its 112 grid steps flat and only 11 distinct levels. Scaling by the dual confidence resolves it into 52 levels and 6 flat steps. Both surfaces have every minimiser decoding to a satisfying assignment of φ. The plan views above each surface show the same grid from directly overhead, shaded by energy, with every tied pair of neighbours drawn in red and the stall point ringed. Those red segments are exactly the 14 and 6 flat steps. The right-hand panel is why the weight is scaled this way. The reduced model has a stall point at which no single flip improves the objective, so the barrier $\Delta^+ = 0.652$ is positive and a descent is stuck. One bump from equation 7 adds δ_w to each of the 4 clauses violated there, lowering Δ^+ by at least $\delta_w \eta_{\min} = 0.622$ per bump. The barrier of this local trap reaches -0.632 after 1 bump, within the bound $\lceil \Delta^+ / \delta_w \eta_{\min} \rceil = 2$, and the stall point is eliminated. Global confidence ε rescales the same barrier without changing its sign, the light bars against the dark, so ε tunes exploration while the local weight is what removes the trap.

Existing problem encoding methods. For SAT problems on quantum annealers, minor embedding maps the SAT-derived Ising/QUBO graph onto the sparse hardware graph. Limited qubit connectivity often forces logical variables to be represented as chains of physical qubits, making embedding a key scalability and performance bottleneck (Bian et al., 2017; Vinci et al., 2015; Bernal et al., 2020; Cattelan et al., 2025). For an instance with n variables and m clauses, the state-of-the-art constructions are Choi's maximum-independent-set encoding at $3m$ qubits (Choi, 2010), Chancellor's direct clause-wise mapping at $n + m$ (Chancellor et al., 2016), and the two constructions of Nüßlein at $2n + m$ and $n + m$ (Nüßlein et al., 2023). Of note, these methods use template-based constructions that explicitly encode literals and clauses as qubits, resulting in resource requirements that scale with both n and m. Recent learning-based embedding methods using reinforcement learning have also been proposed (Ngo et al., 2025; Nembrini et al., 2025). However, they require extensive training to extract instance-specific patterns.

3 The Proposed Framework

AdaptCNLM has two components. First, we propose neural structural-equivalence encoding (NSEE) module, which learns and imposes the partition by grouping structurally equivalent variables (§3.2). Second, dual-confidence scaling assigns clause weights within the induced partition, incorporat-

ing constraints according to their relative confidence (§3.3). The next section, §3.1, reviews the preliminaries underlying the proposed framework.

3.1 PRELIMINARIES

A CNF formula is a conjunction $\varphi = \bigwedge_{j=1}^{m} C_j$ of m clauses over n Boolean variables $s_1, \ldots, s_n$. We write assignments as vectors $\mathbf{s} \in \{0,1\}^n$, where $s_i = 1$ means that variable i is TRUE. Each clause is a disjunction of literals. A literal is a variable or its negation. We record clause C_j as the pair of index sets (C_j^+, C_j^-), where $i \in C_j^+$ if the positive literal s_i occurs in C_j and $i \in C_j^-$ if the negative literal $\neg s_i$ occurs. The width of the clause is $k_j = |C_j^+| + |C_j^-|$. The total literal count is $\Lambda = \sum_j k_j$. The variable set of a clause is $\mathrm{var}(C_j) = C_j^+ \cup C_j^-$. The occurrence count of variable i is $d_i = |\{j : i \in \mathrm{var}(C_j)\}|$. A clause C_j is violated by an assignment $\mathbf{s}$ if and only if all of its literals evaluate to FALSE under $\mathbf{s}$. The violation indicator is the monomial $\chi_j(\mathbf{s}) = \prod_{i \in C_j^+}(1 - s_i) \prod_{i \in C_j^-} s_i$, with $\chi_j(\mathbf{s}) \in \{0,1\}$, which equals 1 exactly when C_j is violated. We minimise the number of violated clauses $U_{\boldsymbol{\omega}}(\mathbf{s}) = \sum_{j=1}^{m} \omega_j \chi_j(\mathbf{s})$, where $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_m)$ is non-negative clause weights. For strictly positive weights, φ is satisfiable iff $\min_{\mathbf{s} \in \{0,1\}^n} U_{\boldsymbol{\omega}}(\mathbf{s}) = 0$, and we assume $\boldsymbol{\omega} > 0$ throughout, which every weight the method produces satisfies.

The monomial $\chi_j(\mathbf{s})$ has degree k_j. A QUBO problem admits degree two only. Reducing degree costs extra binary variables (ancillas); i.e., $\mathbf{s}_a \in \{0,1\}^a$. The QUBO vector is $\mathbf{v} = (\mathbf{s}, \mathbf{s}_a) \in \{0,1\}^N$ with $N = n + a$. A QUBO instance is a symmetric matrix $Q \in \mathbb{R}^{N \times N}$ (or binary quadratic model, BQM) and a constant $c \in \mathbb{R}$, and its energy is $H(\mathbf{v}) = \mathbf{v}^\top Q\, \mathbf{v} + c = \sum_{i=1}^{N} Q_{ii} v_i + 2 \sum_{1 \le i < l \le N} Q_{il} v_i v_l + c$, where the linear term arises from $v_i^2 = v_i$, an identity that holds on $\{0,1\}$.

3.2 NEURAL STRUCTURAL-EQUIVALENCE ENCODING

NSEE learns a partition of the variables into structurally equivalent groups and enforces a shared value among the variables within each group. As a result, the problem is reduced from n variables to $k \ll n$ aggregated variables.

Formalism. We construct a *surjection* $g : [n] \to [k]$ with $k \le n$, that maps variable i to an index $g(i)$, where $[n] = \{1, \ldots, n\}$. Surjectivity means every index in $[k]$ is attained, so k counts the groups that actually occur and no group is empty. Variable groups are the preimages $G_r = g^{-1}(r)$ for $r \in [k]$, which partition $[n]$. The NSEE learns an aggregation matrix

$$A = A(g) \in \{0,1\}^{k \times n}, \quad A_{ri} = \mathbb{I}[g(i) = r]. \tag{1}$$

Every column of A holds exactly one 1 and row r is G_r's indicator vector. Immediately, A's row sums give group sizes ($(A\mathbf{1}_n)_r = |G_r|$) and its column sums give the partition property $A^\intercal \mathbf{1}_k = \mathbf{1}_n$. For example, Fig. 1 shows $g = (1, 2, 1)$; thus, NSEE groups s_1 and s_3 into a group and a singleton s_2. The matrix A is $\left(\begin{smallmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{smallmatrix}\right)$, whose first row is the indicator of $G_1 = \{1, 3\}$ and whose second row is the indicator of $G_2 = \{2\}$.

The group assignment is a vector $\boldsymbol{\xi} \in \{0,1\}^k$ that gives one bit to each group and $\mathbf{s} = A^\top \boldsymbol{\xi}$ so $s_i = \xi_{g(i)}$ for every $i \in [n]$. The map $\boldsymbol{\xi} \mapsto \mathbf{s}$ copies the bit of a group to all variable and the aggregated subcube is $\mathcal{R}(A) = \left\{A^\top \boldsymbol{\xi} : \boldsymbol{\xi} \in \{0,1\}^k\right\}$ with $|\mathcal{R}(A)| = 2^k$.

We define $\varphi_A = \alpha_g(\varphi) = \bigwedge_p D_p$ where $D_{(.)}$ is the image of $C_{(.)}$ under the aggregation mapping. Specifically, Specifically, the image of C_j replaces every literal by the literal of its group, so $D_p = \bigvee_{i \in C_j^+} \xi_{g(i)} \;\vee\; \bigvee_{i \in C_j^-} \neg \xi_{g(i)}$, and the resulting literal set is then reduced..

Substitution has three effects on a clause: (1) repeated literals of one polarity collapse, so the image is narrower than its source, (2) a clause that acquires a variable in both polarities becomes a tautology, and such an image is dropped, since $\chi \equiv 0$ makes it inert, and (3) distinct clauses may share one image, in which case the images are merged and their weights are added. The correspondence between C_j and D_p is recorded by $\Pi \in \{0,1\}^{m_A \times m}$ where $\Pi_{pj} = \mathbb{I}\big[\text{clause } C_j \text{ has image } D_p\big]$, the index p running over the m_A surviving, non-tautological images. The transformed problem is then the weighted MaxSAT problem on φ_A with weights $\Pi\boldsymbol{\omega}$, and it agrees with the original problem at

every lifted point $U_{\boldsymbol{\omega}}\big(A^\top \boldsymbol{\xi}\big) = U^A_{\Pi\boldsymbol{\omega}}(\boldsymbol{\xi}) := \sum_{p=1}^{m_A} (\Pi\boldsymbol{\omega})_p \, \chi^A_p(\boldsymbol{\xi}) \, \forall \boldsymbol{\xi} \in \{0,1\}^k$, with χ^A_p the violation indicator of D_p. The weights $\Pi\boldsymbol{\omega}$ are not a modelling choice. A group assignment that violates an image D_p simultaneously violates every clause mapped to D_p; counting D_p only once would therefore undercount the original objective by the multiplicity of that image. Summing the weights of merged clauses preserves objective equivalence at every lifted point, which is the key identity underlying Theorem 1. The same additivity permits tautological images to be discarded rather than merged: their violation indicators are identically zero, and hence they contribute no weight.

GNN architecture. The CNF φ is represented via a variable-clause incident graph $G_\varphi = (\mathcal{V} \cup \mathcal{C}, \mathcal{E})$, where node set $\mathcal{V}$ has one node per variable and $\mathcal{C}$ has one node per clause. An edge $(i,j) \in \mathcal{E}$ is assigned the polarity $\sigma_{ij} = +1$ if $i \in C_j^+$, $\sigma_{ij} = -1$ if $i \in C_j^-$, and $\sigma_{ij} = 0$ if variable i does not occur in clause C_j.

Let d_i^+ and d_i^- be the number of positive and negative occurrences of variable i. Here, $d_i = d_i^+ + d_i^-$ and $k_j^\pm = |C_j^\pm|$. With $d_{\max} = \max_i d_i$ and $k_{\max} = \max_j k_j$, we define node features

$$\mathbf{x}_i^{\mathcal{V}} = \Big(\tfrac{d_i}{d_{\max}},\ \tfrac{d_i^+}{d_{\max}},\ \tfrac{d_i^-}{d_{\max}},\ \tfrac{d_i^+ - d_i^-}{\max(d_i,1)}\Big) \in \mathbb{R}^4, \qquad \mathbf{x}_j^{\mathcal{C}} = \Big(\tfrac{k_j}{k_{\max}},\ \tfrac{k_j^+}{k_{\max}},\ \tfrac{k_j^-}{k_{\max}}\Big) \in \mathbb{R}^3. \tag{2}$$

The features are scale-free by construction: each degree is normalised by the instance maximum, allowing a single network to generalise across n and m without rescaling or retraining. The variable features capture constraint density, polarity frequencies, and polarity bias, while the clause features encode width and polarity balance. All seven features depend only on the incidence graph, ensuring that the encoder captures structural information without access to candidate assignments.

The encoder appends four statistics of a pool $\mathcal{P}$ to $\mathbf{x}_i^{\mathcal{V}}$, namely the pool mean of variable i, its pool variance, its agreement rate with the pool incumbent and its normalised flip frequency. The message passing is initiated with affine encoding $\mathbf{h}_i^{(0)} = W_{\mathcal{V}}^{(0)} \mathbf{x}_i^{\mathcal{V}} + \mathbf{b}^{(\mathbf{0})}{}_{\mathcal{V}}$ and $\mathbf{h}_j^{(0)} = W_{\mathcal{C}}^{(0)} \mathbf{x}_j^{\mathcal{C}} + \mathbf{b}_{\mathcal{C}}^{(0)}$. Each of L layers performs one variable-to-clause sweep followed by one clause-to-variable sweep:

$$\hat{\mathbf{h}}_j^{(l)} = \frac{1}{k_j} \sum_{i \in \mathrm{var}(C_j)} \psi_{\mathcal{VC}}^{(l)}\big([\mathbf{h}_i^{(l-1)}, \sigma_{ij}]\big), \quad \mathbf{h}_j^{(l)} = \mathrm{LN}\Big(\mathbf{h}_j^{(l-1)} + \mathrm{MLP}_{\mathcal{C}}^{(l)}\big([\mathbf{h}_j^{(l-1)}, \hat{\mathbf{h}}_j^{(l)}]\big)\Big), \tag{3a}$$

$$\hat{\mathbf{h}}_i^{(l)} = \frac{1}{d_i} \sum_{j\,:\, i \in \mathrm{var}(C_j)} \psi_{\mathcal{CV}}^{(l)}\big([\mathbf{h}_j^{(l)}, \sigma_{ij}]\big), \quad \mathbf{h}_i^{(l)} = \mathrm{LN}\Big(\mathbf{h}_i^{(l-1)} + \mathrm{MLP}_{\mathcal{V}}^{(l)}\big([\mathbf{h}_i^{(l-1)}, \hat{\mathbf{h}}_i^{(l)}]\big)\Big). \tag{3b}$$

Here, $[\,\cdot\,,\cdot\,]$ is concatenation, each ψ and MLP is a two-layer perceptron with SiLU activation, and LN is layer normalisation. We write $\mathbf{h}_i = \mathbf{h}_i^{(L)} \in \mathbb{R}^d$ for the final variable embedding and $H \in \mathbb{R}^{n \times d}$ for the stacked matrix. The grouping head is a soft nearest-prototype classifier. It holds k learnable prototypes $\mathbf{c}_1, \ldots, \mathbf{c}_k \in \mathbb{R}^d$ and a learnable temperature $\tau > 0$, stored as $\log \tau$ so it stays positive. The soft assignment is $P_{ri} = \frac{\exp S_{ri}}{\sum_{q=1}^k \exp S_{qi}}$ where $S_{ri} = -\frac{\|\mathbf{h}_i - \mathbf{c}_r\|_2^2}{\tau}$. Here, $P \in [0,1]^{k \times n}$ and $\sum_r P_{ri} = 1$. The deployed grouping is the hard assignment $g(i) = \arg\max_r P_{ri}$, whose matrix is A as in equation 1.

Training objective. The encoder minimises $\mathcal{L}(P) = \lambda_B \, \mathcal{L}_{\mathrm{bal}}(P) + \lambda_H \, \mathcal{L}_{\mathrm{ent}}(P)$, with the two terms defined from the soft group sizes $s_r = \sum_{i=1}^n P_{ri}$, which satisfy $\sum_r s_r = n$, by

$$\mathcal{L}_{\mathrm{bal}}(P) = \frac{k}{n^2} \sum_{r=1}^k \Big(s_r - \frac{n}{k}\Big)^2 \in [0, k-1], \quad \mathcal{L}_{\mathrm{ent}}(P) = \frac{-1}{n \log k} \sum_{r=1}^k \sum_{i=1}^n P_{ri} \log P_{ri}, \tag{4}$$

Here, $\mathcal{L}_{\mathrm{bal}}(P)$ forces an uniform distribution of variables in to k groups and $\mathcal{L}_{\mathrm{ent}}$ forces the model to make more decisive assignment of a variable to a group. The encoder is fitted per instance and without labels.

3.3 Dual-confidence scaling

Aggregation changes the formula and, consequently, the corresponding objective. This subsection first establishes that the clause weights are determined by the encoding, then shows that these weights are exactly as expressive as permitted by the reduced problem, and finally describes the schedule used to assign them.

Local confidence. Local confidence is defined per clause and comprises a static component fixed prior to the search and a dynamic component updated throughout the search.

The *static confidence* $\eta_j \in (0, 1]$ quantifies the apparent difficulty of clause C_j based solely on the formula structure. It combines two structural signals: clause size and variable rarity. A narrow clause is harder to satisfy because it contains fewer literals, while a clause involving rarely occurring variables is harder to repair because those variables participate in fewer other clauses. Let d_i denote the occurrence count of variable i. We define

$$f_j^{\text{size}} = \frac{1}{1+\gamma_s(k_j-1)}, \quad f_j^{\text{rar}} = \frac{1}{k_j}\sum_{\ell \in C_j} \frac{1}{1+\gamma_r d_{\text{var}(\ell)}}, \quad \eta_j = \omega_{\text{size}} f_j^{\text{size}} + \omega_{\text{rar}} f_j^{\text{rar}} \in (0, 1]. \tag{5}$$

The vector $\boldsymbol{\eta}$ is computed once in a single pass over the clause arrays and remains fixed throughout the search.

The *dynamic confidence* $w_j(t) \geq 1$ represents search-derived evidence. It is initialized as $w_j(0) = 1$ for every clause and increases for clauses that remain difficult to satisfy. Let t denote the current flip index, t_{impr} the flip at which the number of violated clauses was most recently reduced, and t_{bump} the flip at which the weights were most recently updated. A weight update is triggered when

$$t - t_{\text{impr}} > \tau_{\text{bump}} \quad \text{and} \quad t - t_{\text{bump}} > \tau_{\text{bump}}, \tag{6}$$

at which point $t_{\text{bump}} \leftarrow t$ and every currently violated clause is assigned an increased weight

$$w_j \leftarrow \min\big(w_j + \delta_w, w_{\max}\big), \forall j \in \mathcal{U}(\mathbf{s}), \delta_w = 1. \tag{7}$$

The two conditions in equation 6 serve distinct purposes. The first makes weight updates responsive to search stagnation rather than to a fixed clock. The second condition imposes a refractory period, limiting weight growth during prolonged stagnation to at most one update every τ_{bump} flips and thereby controlling the rate at which the landscape is reshaped.

Repeated weight increases can nevertheless produce excessive dispersion in $\mathbf{w}$ and concentrate the search on a small number of clauses. We therefore apply a probabilistic smoothing sweep within the update branch. With probability $p_{\text{smooth}} = p_0 \cdot \min\left(1, \frac{m_{\text{ref}}}{m}\right)$, all weights are contracted towards their mean according to $w_j \longleftarrow \rho_s w_j + (1-\rho_s)\bar{w}$ with $\bar{w} = \frac{1}{m}\sum_{l=1}^{m} w_l$.

Global confidence. Global confidence is a single scalar $\varepsilon(t) > 0$ shared by every clause. It multiplies the whole energy difference, so it acts as an inverse-temperature multiplier applied on top of the geometric annealing temperature of the descent kernel, which decays from $T_{\text{init}} = 0.6$ to $T_{\min} = 0.05$. A small ε flattens the landscape and admits more uphill moves. A large ε sharpens it and enforces descent. The schedule has two parts. A linear ramp raises ε from an exploratory value to an exploitative one,

$$\varepsilon_{\text{ramp}}(t) = \min\big(\varepsilon_{\text{tgt}},\ \varepsilon_0 + t\, r\big), \qquad r = \frac{\varepsilon_{\text{tgt}} - \varepsilon_0}{T_{\text{ramp}}}. \tag{8}$$

A stagnation cut lowers ε when the incumbent has not improved. Let t_{last} be the flip of the last improvement to the best count. The implemented rule is the two-branch alternative

$$\varepsilon(t) = \begin{cases} \max\big(\varepsilon_{\min},\ \varepsilon(t-1) - \delta_\varepsilon\big), & t - t_{\text{last}} > \tau_{\text{stag}}, \text{ and then } t_{\text{last}} \leftarrow t, \\ \varepsilon_{\text{ramp}}(t), & \text{otherwise.} \end{cases} \tag{9}$$

4 Theoretical Analysis

Our proposed neural-network-based aggregation constitutes a lossy mapping. This section characterises the information that is preserved and discarded under this mapping, and explains why the encoder must explicitly incorporate information from the input instance.

Preserved information. We state the soundness of aggregation map as follows:

Theorem 1. *Let $A = A(g)$ and let Π be the image incidence of §3.2. Then $U_{\boldsymbol{\omega}}(A^\top \boldsymbol{\xi}) = U^A_{\Pi\boldsymbol{\omega}}(\boldsymbol{\xi})$ for every $\boldsymbol{\xi} \in \{0,1\}^k$. Consequently, $\boldsymbol{\xi} \models \varphi_A \iff A^\top \boldsymbol{\xi} \models \varphi$. The map $\boldsymbol{\xi} \mapsto A^\top \boldsymbol{\xi}$ is a bijection from $\text{sol}(\varphi_A)$ onto $\text{sol}(\varphi) \cap \mathcal{R}(A)$, and $\min_{\boldsymbol{\xi}} U^A_{\Pi\boldsymbol{\omega}} \geq \min_{\mathbf{s}} U_{\boldsymbol{\omega}}$. In particular, φ_A satisfiable implies φ satisfiable.*

Table 1: Comparison of NSEE with existing problem encodings for quantum annealers in terms of logical qubits (k), physical qubits (P) after minor embedding, and solve rate (SR$^{\uparrow}$, higher is better). Physical qubits are obtained by embedding into Pegasus P16 using `minorminer`, with a 90-second timeout; "*" indicates that no embedding was found within the budget.

Encoding method	$n = 40$			$n = 60$			$n = 100$			$n = 140$		
	k	P	SR$^{\uparrow}$	k	P	SR$^{\uparrow}$	k	P	SR$^{\uparrow}$	k	P	SR$^{\uparrow}$
Choi^{3m} (Choi, 2010)	500	3538	0.50	748	*	0.17	1248	*	0.00	1748	*	0.00
Chancellor^{n+m} (Chancellor et al., 2016)	206	650	0.92	310	1211	0.58	516	3022	0.00	722	*	0.00
Nüsslein^{n+m} (Nüßlein et al., 2023)	206	589	0.58	310	1234	0.25	516	2966	0.00	722	*	0.00
Nüsslein^{2n+m} (Nüßlein et al., 2023)	246	1040	0.75	370	2358	0.42	616	*	0.08	862	*	0.00
Rosenberg (Rosenberg, 1975)	373	1010	0.33	559	2144	0.08	932	4577	0.00	1305	*	0.00
NSEE (Ours)	**33**	**44**	**1.00**	**54**	**89**	**1.00**	**95**	**296**	**0.83**	**108**	**408**	**0.92**

Table 2: Benchmarking results of AdaptCNLM against data-driven SAT solvers (first block) and non-learning-based solvers (second block) across six SATLIB families. The **SAT** columns report the solve rate (SR$^{\uparrow}$), defined as the percentage of satisfiable instances for which the returned assignment satisfies all clauses. The **UNSAT** columns report the Clause Satisfaction Score (CSS$^{\uparrow}$), defined as the fraction of clauses satisfied by the returned assignment.

Solver	**Aim**		**Beijing**		**Induc. Infer**	**Jnh**		**QC**		**SSA**		**All**	
	SAT (48)	UNSAT (24)	SAT (9)	UNSAT (1)	SAT (41)	SAT (16)	UNSAT (34)	SAT (5)	UNSAT (8)	SAT (4)	UNSAT (3)	SAT (123)	UNSAT (70)
SATNet (Wang et al., 2019)	0.0	.8971	0.0	.9374	0.0	0.0	.9528	0.0	.9692	0.0	.7552	0.0	.9269
SGAT-MS (Li et al., 2023)	0.0	.9744	44.4	.8172	0.0	0.0	.9775	0.0	.9851	0.0	.9592	3.3	.9742
GCN-LCG (Li et al., 2023)	0.0	.9636	44.4	.9845	19.5	0.0	.9901	0.0	.9865	0.0	.9828	9.8	.9802
GGNN-LCG (Li et al., 2023)	2.1	.9639	44.4	.9845	9.8	0.0	.9758	0.0	.9857	0.0	.9812	7.3	.9732
GIN-LCG (Li et al., 2023)	0.0	.9601	33.3	.9845	2.4	0.0	.9625	0.0	.9834	0.0	.9521	3.3	.9639
GMS-E (Liu, 2023)	4.2	.9926	44.4	.9655	12.2	0.0	.9489	0.0	.9834	0.0	.9889	8.9	.9697
GMS-N (Liu, 2023)	6.2	.9628	44.4	.9845	19.5	0.0	.9894	0.0	.9862	0.0	.9915	12.2	.9799
NeuroSAT (Selsam et al., 2018)	2.1	.9642	44.4	.9845	19.5	0.0	.9731	0.0	.9850	0.0	.9850	10.6	.9721
PDP (Amizadeh et al., 2019)	18.8	.9932	44.4	.9845	4.9	0.0	.9895	0.0	.9859	0.0	.9899	12.2	.9903
QuerySAT (Ozolins et al., 2022)	25.0	**.9934**	44.4	.9845	19.5	0.0	.9861	0.0	.9727	0.0	.9883	19.5	.9871
Parallel Tempering (Swendsen & Wang, 1986)	64.6	**.9934**	44.4	**.9993**	92.7	**100.0**	**.9980**	0.0	**.9990**	0.0	.9985	72.4	**.9966**
Survey Propagation (Braunstein et al., 2005)	68.8	**.9934**	33.3	.9986	78.0	93.8	**.9980**	0.0	.9977	75.0	.9984	69.9	.9964
Simulated Annealing (Kirkpatrick et al., 1983)	60.4	**.9934**	44.4	**.9993**	65.9	**100.0**	**.9980**	0.0	.9986	0.0	.9982	61.8	.9965
RbmSAT (Warde-Farley et al., 2023)	10.4	**.9934**	11.1	.9972	4.9	6.2	.9916	0.0	.9164	0.0	.9845	7.3	.9834
WalkSAT (Selman et al., 1994)	66.7	**.9934**	**88.9**	**.9993**	**100.0**	**100.0**	**.9980**	20.0	.9974	**100.0**	**.9994**	82.9	.9964
probSAT (Balint & Schöning, 2012)	**70.8**	**.9934**	33.3	.9986	78.0	**100.0**	**.9980**	0.0	.9977	75.0	.9984	71.5	.9964
AdaptCNLM (Ours)	68.8	**.9934**	66.7	**.9993**	**100.0**	**100.0**	**.9980**	**100.0**	**.9990**	**100.0**	**.9994**	**85.4**	**.9966**

Trade-off of the problem reduction. The converse of theorem 1 is generally false; i.e., φ satisfiable does not imply φ_A satisfiable. A counterexample is $\varphi = (s_1 \vee s_2) \wedge (\neg s_1 \vee \neg s_2)$ with $k = 1$, for which $\varphi_A = \xi \wedge \neg\xi$. This theoretical result raises the following question: given a random K-SAT instance, what is the probability that the aggregated formula φ_A remains satisfiable? We formalise this question as follows:

Theorem 2. *Let φ be uniform random K-SAT with n variables and m clauses, let g be any surjection chosen independently of φ, and let $Z_A = |\mathrm{sol}(\varphi) \cap \mathcal{R}(A)|$. Then $\mathbb{E}[Z_A] = 2^k(1 - 2^{-K})^m$ exactly, so $\mathbb{P}(\varphi_A$ satisfiable$) \le 2^k(1 - 2^{-K})^m$. Writing $m = \alpha n$ and $\kappa(\alpha, K) = \alpha \log_2(1 - 2^{-K})^{-1}$, if $k \le (1 - \delta)\kappa(\alpha, K)n$ for some $\delta > 0$ then $\mathbb{P}(\varphi_A$ satisfiable$) \le 2^{-\delta\kappa(\alpha,K)n}$.*

5 Experiments

Experimental design. The evaluation of NSEE uses uniform random 3-SAT at $n \in \{40, 60, 100, 140, 200\}$ and clause density $\alpha/\alpha_c \in \{0.95, 1.0\}$ with $\alpha_c = 4.267$, six instances per cell, so twelve per size and 120 instances in all. The encoder uses $L = 3$ message-passing layers with hidden width $d = 48$. It is fitted independently for each instance and is never pretrained. Every instance is filtered by the complete solver CaDiCaL153 through `pysat`. The evaluation of AdaptCNLM uses six SATLIB families (Hoos, 2000–2011) comprising 193 instances, 123 satisfiable and 70 unsatisfiable, together with two SAT Competition instances. The factorial ablation uses 72 random K-SAT instances with $K \in \{3, 4, 5\}$ and $\alpha/\alpha_c \in \{0.8, 0.9, 1.0\}$ at n sampled from $[200, 400]$, $[100, 200]$ and $[60, 120]$, respectively.

Effectiveness of NSEE. We compare NSEE against five established problem encodings for 3SAT on quantum annealers. Table 1 reports the number of logical qubits and verified SR. For each cell, we evaluate six random 3SAT instances at $\alpha \in \{0.95, 1.0\}\alpha_c$, using the same simulated annealing solver implemented in `neal` with 1000 sweeps; solutions are verified against ground truth obtained from CADICAL153. NSEE consistently requires substantially fewer logical qubits while achieving the highest SR. For $n = 40$, NSEE uses 33 logical qubits and achieves an SR of 1.00, compared with 206–500 logical qubits and SRs of 0.50–0.92 for existing encodings. At $n = 100$, NSEE uses 95 logical qubits with an SR of 0.83, whereas the closest alternatives require at least 616 logical qubits and achieve at most 0.08 SR. Regarding physical qubits, we embed each encoding into the Pegasus P16 topology using `minorminer` (Cai et al., 2014), with one instance per cell and a 90-second timeout. NSEE requires substantially fewer physical qubits than existing encodings: 44, 89, 296, and 408 qubits for $n = 40, 60, 100, 140$, respectively, while the best competing encodings require 589, 1211, 2966, and no embedding is found within the 90-second budget at $n = 140$.

Effectiveness of AdaptCNLM. We compare AdaptCNLM against 16 methods, stratified into two categories, (1) data-driven solvers (first block) and (2) non-data-driven (second block), shown in Table 2. The former are neural-network-based methods, each trained on 35,280 instances with the best model selected by validation on a further 8,820 instances. AdaptCNLM outperforms every data-driven solver on the solve rate in all six families. The data-driven methods fail catastrophically on Jnh, QC and SSA, where all data-driven solvers return an SR of 0% over the 25 satisfiable instances of those three families, while AdaptCNLM solves every one of them. Aggregated over the benchmark, AdaptCNLM returns a verified satisfying assignment on 105 of the 123 satisfiable instances, that is 85.4%, against 24 of 123 or 19.5% for QuerySAT, the strongest data-driven competitor, a margin of 81 instances. CSS on unsatisfiable instances gives the same ordering. AdaptCNLM is strictly above the best data-driven value on four of the five families that contain unsatisfiable instances. Beyond the six benchmarking datasets, we evaluate AdaptCNLM on two competition instances `sgen1-sat-180-100` ($n = 180, m = 432$) and `sgen3-n260-s62321009-sat` ($n = 260$,$m = 624$) from SATCOM24-25 (sat, 2024; 2025). AdaptCNLM solves both instances, whereas the data-driven SAT solvers fail on both instances. Against the non-data-driven solvers, the aggregate margin is narrower, and the more informative difference is one of consistency rather than of peak performance. AdaptCNLM attains the highest aggregate solve rate at 85.4%, ahead of WalkSAT at 82.9%, Parallel Tempering at 72.4%, probSAT at 71.5%, Survey Propagation at 69.9%, Simulated Annealing at 61.8% and RbmSAT at 7.3%, and it attains the highest aggregate CSS at .9966, which it shares with Parallel Tempering.

Ablation study. We design an ablation study to isolate the contributions of two components: A, denoting the presence of NSEE, and B, denoting the application of dual-confidence scaling. The baseline configuration, denoted by A^-B^-, solves the energy function induced directly from the original CNF formula, whereas the proposed end-to-end configuration, A^+B^+, incorporates both NSEE and dual-confidence scaling. We generate 72 random K-SAT instances with $K \in \{3, 4, 5\}$ and $\alpha/\alpha_c \in \{0.8, 0.9, 1\}$, with 8 instances per $(K, \alpha/\alpha_c)$ cell. The numbers of variables for 3-, 4-, and 5-SAT are sampled from $[200, 400]$, $[100, 200]$, and $[60, 120]$, respectively, with the ranges selected such that the optimiser cannot identify a satisfying assignment within the 20-second runtime budget. The A^+B^- configuration improves the SR by 0.472, while reducing the number of logical qubits from 3785 to 764 ($p = 7.3 \times 10^{-9}$). In contrast, the A^-B^+ configuration does not yield a measurable improvement in solver performance. The complete end-to-end configuration, A^+B^+, achieves the largest SR improvement of 0.486, with a 95% confidence interval of $[0.375, 0.597]$, while using approximately the same number of physical qubits as A^+B^-. These results demonstrate that NSEE provides an effective problem encoding that substantially improves SR within a fixed runtime budget, while dual-confidence scaling helps compensate for the information loss introduced by problem compression.

6 DISCUSSION

Rationale of the proposed framework. Table 1 demonstrates that our approach achieves a substantial reduction in qubit requirements. This reduction stems from the fact that competing techniques typically employ template-based problem encodings, for which the required number of qubits scales with the problem size (n, m). In contrast, NSEE learns structural patterns from the variable–clause

graph and aggregates variables into groups. This aggregation effectively reduces the dimensionality and structural complexity of the original problem, thereby decreasing the number of qubits required for its quantum representation. This qubit-overhead advantage, however, comes at the cost of exactness, discussed in §4.

Limitations. We identify two limitations of the current work. First, the method is capable of certifying satisfiability but does not provide certificates of unsatisfiability. A returned assignment is verified against the original formula, whereas failure to find a satisfying assignment does not constitute a refutation. Moreover, the reduced model does not retain sufficient information to construct an unsatisfiability certificate. Second, our evaluation does not include experiments on actual quantum hardware due to limited hardware access. All experimental results are therefore obtained using a classical simulated-annealing sampler, which serves as a uniform proxy across the different encodings. Nevertheless, the physical qubit requirements are computed using a commercial off-the-shelf package, and we therefore expect the corresponding qubit counts to remain the same when the encodings are deployed on actual quantum hardware.

7 Conclusion

We introduced AdaptCNLM, which couples neural structural-equivalence encoding with dual-confidence scaling to compress CNF formulas into compact aggregated problems without sacrificing solution quality. NSEE reduces the logical- and physical-qubit requirements of quantum-annealer encodings by an order of magnitude while attaining the highest solve rate among five compared encodings, and theoretical analyses characterise exactly which satisfying assignments survive aggregation and how reduction size trades against feasibility. As a standalone solver, AdaptCNLM outperforms sixteen data-driven and non-data-driven baselines across six SATLIB families.

AI use statement

In this work, we used generative AI tools for grammar and language editing, code and theorem auditing, and simulated peer review of the original draft. We did not use generative AI tools for method formulation, implementation planning, or the formulation or proof of theorems. We have carefully reviewed all AI-assisted work and take full responsibility for the final content of this work, including all text, claims, analyses, and other artifacts produced with the assistance of generative AI.

Ethics statement

We declare that there are no conflicts of interest.

Reproducibility Statement

We provide detailed descriptions of the proposed method and its implementation in the main text and appendices. The source code implementing our method is also submitted with the manuscript to facilitate reproducibility.

References

Sat competition 2024. https://satisfiability.org/SAT24/competitions.php, 2024. SAT Conference 2024, Pune, India.

Sat competition 2025. https://satcompetition.github.io/2025/index.html, 2025. Competition of Boolean Satisfiability Solvers.

Aravindhan Alagarsamy and Lakshminaraynan Gopalakrishnan. Sat: a new application mapping method for power optimization in 2d—noc. In *2016 20th International Symposium on VLSI Design and Test (VDAT)*, pp. 1–6. IEEE, 2016.

Mohammad H Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Physical Review X*, 8(2):021050, 2018.

Saeed Amizadeh, Sergiy Matusevych, and Markus Weimer. Learning to solve circuit-sat: An unsupervised differentiable approach. In *International conference on learning representations*, 2018.

Saeed Amizadeh, Sergiy Matusevych, and Markus Weimer. Pdp: A general neural framework for learning constraint satisfaction solvers. *arXiv preprint arXiv:1903.01969*, 2019.

Carlos Ansótegui and Jordi Levy. Sat, gadgets, max2xor, and quantum annealers. *arXiv preprint arXiv:2403.00182*, 2024.

Adrian Balint and Uwe Schöning. Choosing probability distributions for stochastic local search and the role of make versus break. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 16–29. Springer, 2012.

David E Bernal, Kyle EC Booth, Raouf Dridi, Hedayat Alghassi, Sridhar Tayur, and Davide Venturelli. Integer programming techniques for minor-embedding in quantum annealers. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pp. 112–129. Springer, 2020.

Zhengbing Bian, Fabian Chudak, William Macready, Aidan Roy, Roberto Sebastiani, and Stefano Varotti. Solving sat and maxsat with a quantum annealer: Foundations and a preliminary report. In *International Symposium on Frontiers of Combining Systems*, pp. 153–171. Springer, 2017.

Zhengbing Bian, Fabian Chudak, William Macready, Aidan Roy, Roberto Sebastiani, and Stefano Varotti. Solving sat (and maxsat) with a quantum annealer: Foundations, encodings, and preliminary results. *Information and computation*, 275:104609, 2020.

Kelly Boothby, Paul Bunyk, Jack Raymond, and Aidan Roy. Next-generation topology of d-wave quantum processors. *arXiv preprint arXiv:2003.00133*, 2020.

Kelly Boothby, Andrew D. King, and Jack Raymond. Zephyr topology of d-wave quantum processors. Technical report, D-Wave Systems Inc., 2021. URL `https://www.dwavequantum.com/media/2uznec4s/14-1056a-a_zephyr_topology_of_d-wave_quantum_processors.pdf`.

Alfredo Braunstein, Marc Mézard, and Riccardo Zecchina. Survey propagation: An algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201–226, 2005.

Paul I Bunyk, Emile M Hoskinson, Mark W Johnson, Elena Tolkacheva, Fabio Altomare, Andrew J Berkley, Richard Harris, Jeremy P Hilton, Trevor Lanting, Anthony J Przybysz, et al. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Transactions on Applied Superconductivity*, 24(4):1–10, 2014.

Connor Bybee, Denis Kleyko, Dmitri E Nikonov, Amir Khosrowshahi, Bruno A Olshausen, and Friedrich T Sommer. Efficient optimization with higher-order ising machines. *Nature Communications*, 14(1):6033, 2023.

Jun Cai, William G. Macready, and Aidan Roy. A practical heuristic for finding graph minors. 2014.

Ernesto Campos, Salvador E Venegas-Andraca, and Marco Lanzagorta. Quantum tunneling and quantum walks as algorithmic resources to solve hard k-sat instances. *Scientific Reports*, 11(1): 16845, 2021.

Michele Cattelan, Jemma Bennett, Sheir Yarkoni, and Wolfgang Lechner. Scalable embedding of parity constraints in quantum annealing hardware. *Physical Review A*, 111(1):012435, 2025.

Nicholas Chancellor, Stefan Zohren, Paul A Warburton, Simon C Benjamin, and Stephen Roberts. A direct mapping of max k-sat and high order parity checks to a chimera graph. *Scientific reports*, 6 (1):37107, 2016.

Vicky Choi. Adiabatic quantum algorithms for the np-complete maximum-weight independent set, exact cover and 3sat problems. *arXiv preprint arXiv:1004.2226*, 2010.

Brandon C Colelough and William Regli. Neuro-symbolic ai in 2024: A systematic review. *arXiv preprint arXiv:2501.05435*, 2025.

Fabien Corblin, Lucas Bordeaux, Eric Fanchon, Youssef Hamadi, and Laurent Trilling. Connections and integration with sat solvers: A survey and a case study in computational biology. In *Hybrid Optimization: The Ten Years of CPAIOR*, pp. 425–461. Springer, 2010.

Adam Douglass, Andrew D King, and Jack Raymond. Constructing sat filters with a quantum annealer. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 104–120. Springer, 2015.

Michael D Ernst, Todd D Millstein, and Daniel S Weld. Automatic sat-compilation of planning problems. In *IJCAI*, volume 97, pp. 1169–1176, 1997.

Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint quant-ph/0001106*, 2000.

Robert Simon Fong, Yanming Song, and Alexander Yosifov. Efficient digital quadratic unconstrained binary optimization solvers for sat problems. *New Journal of Physics*, 27(1):013027, 2025.

Malay K Ganai and Aarti Gupta. *SAT-based scalable formal verification solutions*. Springer, 2007.

David J Griffiths and Darrell F Schroeter. *Introduction to quantum mechanics*. Cambridge university press, 2018.

Jun Gu. Global optimization for satisfiability (sat) problem. *IEEE Transactions on Knowledge and Data Engineering*, 6(3):361–381, 1994.

Sean B Holden et al. Machine learning for automated theorem proving: Learning to solve sat and qsat. *Foundations and Trends® in Machine Learning*, 14(6):807–989, 2021.

Holger H. Hoos. SATLIB benchmark problems. `https://www.cs.ubc.ca/~hoos/SATLIB/benchm.html`, 2000–2011. Accessed: 2026-02-07.

Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58(5):5355, 1998.

Meerzhan Kanatbekova, Vincenzo De Maio, and Ivona Brandic. Qubit-efficient qubo formulation for constrained optimization problems. *arXiv preprint arXiv:2509.08080*, 2025.

Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.

Frédéric Lafitte. Cryptosat: a tool for sat-based cryptanalysis. *IET Information Security*, 12(6): 463–474, 2018.

Zhaoyu Li, Jinpei Guo, and Xujie Si. G4satbench: Benchmarking and advancing sat solving with graph neural networks. *arXiv preprint arXiv:2309.16941*, 2023.

Minghao Liu. GMS: Solve the maxsat problem with graph neural networks. `https://github.com/minghao-liu/GMS`, 2023. GitHub repository.

Andrew Lucas. Ising formulations of many np problems. *Frontiers in physics*, 2:74887, 2014.

Inês Lynce and Joao Marques-Silva. Sat in bioinformatics: Making the case with haplotype inference. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 136–141. Springer, 2006.

Filip Maric. Formal verification of modern sat solvers. *Archive of formal proofs*, 2008.

João P Marques-Silva and Karem A Sakallah. Boolean satisfiability in electronic design automation. In *Proceedings of the 37th Annual Design Automation Conference*, pp. 675–680, 2000.

Fabio Massacci and Laura Marraro. Logical cryptanalysis as a sat problem. *Journal of Automated Reasoning*, 24(1):165–203, 2000.

Vrinda Mehta, Fengping Jin, Hans De Raedt, and Kristel Michielsen. Quantum annealing for hard 2-satisfiability problems: Distribution and scaling of minimum energy gap and success probability. *Physical Review A*, 105(6):062406, 2022.

Vrinda Mehta, Hans De Raedt, Kristel Michielsen, and Fengping Jin. Performance of quantum annealing for 2-satisfiability problems with multiple satisfying assignments. *Physical Review A*, 112(1):012405, 2025.

Puya Mirkarimi, Adam Callison, Lewis Light, Nicholas Chancellor, and Viv Kendon. Comparing the hardness of max 2-sat problem instances for quantum and classical algorithms. *Physical Review Research*, 5(2):023151, 2023.

Ilya Mironov and Lintao Zhang. Applications of sat solvers to cryptanalysis of hash functions. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 102–115. Springer, 2006.

Riccardo Nembrini, Maurizio Ferrari Dacrema, and Paolo Cremonesi. Minor embedding for quantum annealing with reinforcement learning. *arXiv preprint arXiv:2507.16004*, 2025.

Hoang Ngo, Nguyen Do, Minh Vu, Tre Jeter, Tamer Kahveci, and My Thai. Charme: A chain-based reinforcement learning approach for the minor embedding problem. *ACM Transactions on Quantum Computing*, 7(1):1–28, 2025.

Jonas Nüßlein, Sebastian Zielinski, Thomas Gabor, Claudia Linnhoff-Popien, and Sebastian Feld. Solving (max) 3-sat via quadratic unconstrained binary optimization. In *International Conference on Computational Science*, pp. 34–47. Springer, 2023.

Hiroki Oshiyama and Masayuki Ohzeki. Benchmark of quantum-inspired heuristic solvers for quadratic unconstrained binary optimization. *Scientific reports*, 12(1):2146, 2022.

Emils Ozolins, Karlis Freivalds, Andis Draguns, Eliza Gaile, Ronalds Zakovskis, and Sergejs Kozlovics. Goal-aware neural sat solver. In *2022 International joint conference on neural networks (IJCNN)*, pp. 1–8. IEEE, 2022.

Kristen L Pudenz, Gregory S Tallant, Todd R Belote, and Steven H Adachi. Quantum annealing and the satisfiability problem. In *New Frontiers in High Performance Computing and Big Data*, pp. 253–260. IOS Press, 2017.

Ivo G Rosenberg. Reduction of bivalent maximization to the quadratic case. 1975.

Bart Selman, Henry A Kautz, Bram Cohen, et al. Noise strategies for improving local search. In *AAAI*, volume 94, pp. 337–343, 1994.

Daniel Selsam, Matthew Lamm, Benedikt Bünz, Percy Liang, Leonardo de Moura, and David L Dill. Learning a sat solver from single-bit supervision. *arXiv preprint arXiv:1802.03685*, 2018.

Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending sat solvers to cryptographic problems. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 244–257. Springer, 2009.

Robert H Swendsen and Jian-Sheng Wang. Replica monte carlo simulation of spin-glasses. *Physical review letters*, 57(21):2607, 1986.

Grigori S Tseitin. On the complexity of derivation in propositional calculus. In *Automation of reasoning: 2: Classical papers on computational logic 1967–1970*, pp. 466–483. Springer, 1983.

Walter Vinci, Tameem Albash, Gerardo Paz-Silva, Itay Hen, and Daniel A Lidar. Quantum annealing correction with minor embedding. *Physical Review A*, 92(4):042310, 2015.

Po-Wei Wang, Priya L. Donti, Bryan Wilder, and J. Zico Kolter. SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *International Conference on Machine Learning (ICML)*, pp. 6545–6554, 2019.

David Warde-Farley, Vinod Nair, Yujia Li, Ivan Lobov, Felix Gimeno, and Simon Osindero. Solving maxsat with matrix multiplication. *arXiv preprint arXiv:2311.02101*, 2023.

Zewen Zhang, Roger Paredes, Bhuvanesh Sundar, David Quiroga, Anastasios Kyrillidis, Leonardo Duenas-Osorio, Guido Pagano, and Kaden RA Hazzard. Grover-qaoa for 3-sat: Quadratic speedup, fair-sampling, and parameter clustering. *Quantum Science and Technology*, 10(1):015022, 2025.

Sebastian Zielinski, Jonas Nüßlein, Michael Kölle, Thomas Gabor, Claudia Linnhoff-Popien, and Sebastian Feld. Solving max-3sat using qubo approximation. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 1, pp. 681–691. IEEE, 2024.

A IMPLEMENTATION OF NSEE

Elite pool and backbone. A confidence-weighted WALKSAT pass retains the 48 best distinct assignments found over twelve restarts of four thousand flips. Call this the elite pool $\mathcal{P} \subset \{0,1\}^n$. A variable taking the same value in every member of $\mathcal{P}$ is fixed to that value and leaves the model, and the rule is withheld unless $\mathcal{P}$ holds at least six distinct assignments. Write $F \subseteq [n]$ for the variables that remain free. A fixed variable costs no group and no qubit, and it satisfies clauses outright, which removes their images and their ancillas as well. The network is applied to the residual formula $\varphi|_F$, so the features of equation 2 are those of the residual and the surjection returned has domain F.

Merge hypothesis. Let $\Pi_{\mathrm{co}} = \{\{a,b\} \subseteq F \mid a \neq b,\ \exists j,\ a,b \in \mathrm{var}(C_j)\}$ be the pairs of free variables sharing a clause. These are the only pairs a merge can act on, because joining two variables with no clause in common changes no image. Rather than commit to the hard assignment, we read the head's own distribution. Under P the two variables of a pair coincide with probability

$$\alpha_{ab} \;=\; \sum_{r=1}^{k} P_{ra} P_{rb} \;\in [0,1], \tag{10}$$

which is the network's graded opinion about that merge and comes from the same forward pass that yields g. A threshold $\rho \in (0,1]$ admits the $\lceil \rho |\Pi_{\mathrm{co}}| \rceil$ pairs of largest affinity and forbids the rest, giving the admitted set $\Pi_\rho \subseteq \Pi_{\mathrm{co}}$. We use $\rho = 0.6$ throughout, and $\rho = 1$ forbids nothing and so removes the network's influence entirely.

Feasibility certificate. For a partition g and a pool member $\mathbf{p} \in \mathcal{P}$, define the projection $\lambda_g(\mathbf{p}) = A^\top \boldsymbol{\xi}$ with $\xi_r = \mathbb{I}[|G_r|^{-1} \sum_{i \in G_r} p_i \geq \frac{1}{2}]$, the majority vote inside each group lifted back to $\{0,1\}^n$. The certificate is

$$\mathrm{cert}(g) \;=\; \min_{\mathbf{p} \in \mathcal{P}} \; U_{\mathbf{1}}\big(\lambda_g(\mathbf{p})\big). \tag{11}$$

The quantity reads directly. If $\mathrm{cert}(g) = 0$ then a satisfying assignment of φ lies in $\mathcal{R}(A)$, so φ_A is satisfiable and Theorem 1 carries its optimum back to a solution of φ. A positive value bounds the damage from above and is not a proof of infeasibility, since $\mathcal{P}$ visits only finitely many points of $\mathcal{R}(A)$. Evaluating equation 11 costs one matrix product and one pass over φ, which is negligible beside one annealing batch.

Guarded ascent. The partition is built upward rather than taken from the network outright. Every free variable starts alone, which is feasible because the identity partition has $\mathcal{R}(A) = \{0,1\}^n$ and therefore $\mathrm{cert}(\mathrm{id}) = \min_{\mathbf{p} \in \mathcal{P}} U_{\mathbf{1}}(\mathbf{p})$. A move relocates one variable into the group of a variable it shares a clause with, and it is accepted only when the pair lies in Π_ρ, the move strictly lowers the number of logical qubits. The search runs at most ten passes in random order with at most six candidate groups per variable and stops when a pass accepts nothing.

Model construction and selection. The aggregated formula is simplified to a fixed point by moves sound for the weighted objective, which pays because aggregation manufactures short images and short images are cheap or free under t. Surviving images are quadratised with the compact gadget, so a clause of width κ spends $t(\kappa)$ ancillas rather than $\kappa - 1$. The network is queried at three group budgets $k \in \{\lceil 0.06|F| \rceil, \lceil 0.10|F| \rceil, \lceil 0.18|F| \rceil\}$, each answer is guarded, and among the results attaining the smallest certificate the model with the fewest qubits is kept. A round ending with violated clauses records every pair of free variables that a violated clause holds inside one group, removes those pairs from Π_ρ and queries the network again at the next budget up.

B HYPERPARAMETERS

Table 3: Hyperparameter setting of AdaptCNLM

Stage	Symbol	Value	Meaning/Note
Static confidence	$\omega_{\text{size}},\ \omega_{\text{rar}}$	$0.7,\ 0.3$	sum to one, giving $\eta_j \in (0, 1]$
	$\gamma_s,\ \gamma_r$	$0.2,\ 0.1$	size and rarity decay
Dynamic confidence	τ_{bump}	30	stagnation and refractory window
	δ_w	1.0	bump increment
	$w_{\max}$	50.0	weight cap
	$p_0,\ m_{\text{ref}}$	$0.05,\ 2000$	smoothing probability and reference size
	ρ_s	0.8	smoothing contraction
Global confidence	$\varepsilon_0, \varepsilon_{\text{tgt}}, \varepsilon_{\min}$	$0.6,\ 1.0,\ 0.1$	ramp endpoints and floor
	$\delta_\varepsilon,\ \tau_{\text{stag}}$	$0.1,\ 400$	stagnation cut
	T_{ramp}	10^6	*fixed*, independent of instance size
Descent kernel	$T_{\text{init}},\ T_{\min}$	$0.6,\ 0.05$	geometric over 2×10^5 flips
	p_{noise}	0.35	unconditional-accept branch
	tabu tenure	3	flips
	chains	1	single chain in all reported results
Elite pool	$\lvert\mathcal{P}\rvert$	48	best distinct assignments retained
	restarts, flips	$12,\ 4000$	per restart
	backbone minimum	6	distinct assignments required before fixing
Encoder	L, width d	3, 48	message-passing layers and hidden width
	$\lambda_B,\ \lambda_H$	$0.5,\ 0.05$	balance and entropy weights
	optimiser	Adam, 10^{-3}, 10 steps	fitted per instance, unsupervised
Guard	ρ	0.6	affinity threshold
	passes, candidates	$\leq 10, \leq 6$	guarded ascent
	budget ladder	$\{0.06, 0.10, 0.18\}\,\lvert F\rvert$	rounded up
Sampling	`neal` sweeps, reads	1000, 25	identical for every encoding
	embedding budget	90 s	`minorminer` into Pegasus P_{16}

C SUPPLEMENTAL DETAILS OF RESULTS

Sampler and budget. Every encoding in Table 1 is sampled by the same simulated annealer, `neal` at 1000 sweeps per read, so the sampler is held fixed and only the encoded model varies. Wall clock is matched. The proposed method runs first and every other arm receives exactly the time it used, encoding included, so the cost of running the elite pool, the network and the guard is charged against the method and not hidden. Measured end-to-end encoding time is 1.7, 2.4, 5.0, 4.8 and 5.8 seconds at $n = 40$ to 200, of which the elite pool takes 23 to 46 per cent and the guarded encode 39 to 56 per cent. Table 2 uses a 180 second budget per instance for every printed method, AdaptCNLM included, on a single search chain.

Preprocessing. Unit propagation and pure-literal elimination run before every stage, so the n and m of Table 1 and Table 2 are post-preprocessing counts.

Descent kernel. The elite pool and the classical solver share one kernel. A violated clause is drawn with probability proportional to $\eta_j w_j |C_j|$ by rejection sampling against a maintained bound. With probability $p_{\text{noise}} = 0.35$ a uniformly random variable of the selected clause is flipped and the move is accepted unconditionally, bypassing the Metropolis test, which is the only escape mechanism that survives once the temperature has decayed. Otherwise the move is accepted when $\Delta E \leq 0$ or with probability $\exp(-\Delta E/T)$, where $\Delta E = \varepsilon\, \bar{d}$ is the confidence-scaled break count and T follows the geometric schedule of §3.3.

List of benchmarking instances. The encoding study of Table 1 uses generated instances rather than a fixed corpus. For each $n \in \{40, 60, 100, 140, 200\}$ and each $\alpha/\alpha_c \in \{0.95, 1.0\}$ we draw six uniform random 3-SAT instances with $m = \text{round}(\alpha n)$, filter by CADICAL153 and retain only satisfiable ones, giving 120 instances and 1380 annealing runs in total. Clause counts run from $m = 162$ at the smallest size and density to $m = 853$ at the largest. Generator seeds are recorded with the results. The solver study of Table 2 uses six SATLIB families (Hoos, 2000–2011), namely AIM with 48 satisfiable and 24 unsatisfiable instances, BEIJING with 9 and 1, INDUCTIVE

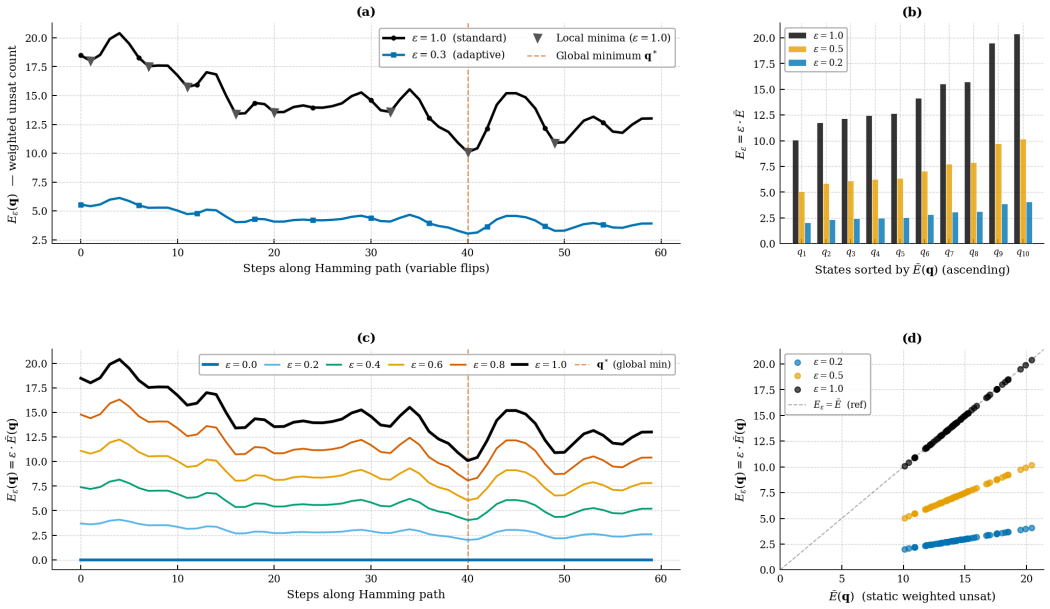


Figure 2: Topology preservation under confidence scaling. **(a)** Hamming-path energy landscape at $\varepsilon = 1.0$ and $\varepsilon = 0.3$: local minima and the global minimum (red dashed) occupy identical positions under both scalings. A Hamming path is a sequence $\mathbf{q}^{(0)}, \mathbf{q}^{(1)}, \ldots, \mathbf{q}^{(n)}$ where $d_H(\mathbf{q}^{(t)}, \mathbf{q}^{(t+1)}) = 1$ for every step; i.e., exactly one variable is flipped at each step. **(b)** Bar chart of ten states sorted by $\bar{E}(\mathbf{q})$: the ascending order is identical for all $\varepsilon \in \{0.2, 0.5, 1.0\}$. **(c)** Full sweep $\varepsilon \in \{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$: the global minimum index (red dashed) is invariant across all curves. **(d)** Scatter of E_ε versus $\bar{E}$: the strictly linear relationship $E_\varepsilon = \varepsilon \cdot \bar{E}$ preserves all pairwise inequalities.

INFERENCE with 41 satisfiable, JNH with 16 and 34, the quasigroup family QG, displayed as QC, with 5 and 8, and SSA with 4 and 3. The totals are 123 satisfiable and 70 unsatisfiable. One further instance is covered by every method and set aside, because AdaptCNLM refutes it outright and therefore returns no assignment for either metric to score. The two competition instances are `sgen1-sat-180-100` and `sgen3-n260-s62321009-sat` from the SAT Competition 2024 and 2025 benchmark sets (sat, 2024; 2025).

Qubit-overhead advantages. The qubit overhead is not a constant factor that hardware growth will absorb. It is a structural property of how SAT instances are encoded. Every compared CNF-to-QUBO construction are template-based methods. It walks the clause list, emits a fixed gadget per clause and per literal, and therefore spends a number of logical qubits that is an affine function of (n, m) and of nothing else. On random 3-SAT at the satisfiability threshold $\alpha_c = 4.267$, the cheapest published constructions spend $n + m \approx 5.27n$ logical qubits, and the classic construction of Choi (2010) spends $3m \approx 12.8n$. Minor embedding then multiplies that count again, because a logical qubit of degree higher than the hardware degree must be realised as a chain of physical qubits. The compared constructions are affine in n with slope $1 + \alpha \approx 5.3$ logical qubits per variable, whereas the measured NSEE count grows at roughly 0.5 qubits per variable over this range, because the backbone removes 20 to 47 per cent of the variables before any grouping occurs and the surviving images are narrow enough for the compact gadget to charge little. The advantage therefore widens with n, from $6.2\times$ on logical qubits at $n = 40$ to $6.7\times$ at $n = 140$, and from $13.4\times$ to $10.0\times$ on physical qubits over the range where a competing embedding exists at all.

Topology preservation of energy landscape. The first geometric guarantee is that the adaptive schedule cannot inadvertently destroy a global minimum or fabricate a new one. Specifically, for any fixed set of clause-level confidences $\{\eta_j\}_{j=1}^m \subset (0, \infty)$ and any global confidence $\varepsilon(t) \in (0, 1]$, these following claims held: **(i)** $\arg\min_{\mathbf{q} \in \{0,1\}^Q} E_\varepsilon(\mathbf{q}) = \arg\min_{\mathbf{q} \in \{0,1\}^Q} \bar{E}_\varepsilon(\mathbf{q})$;

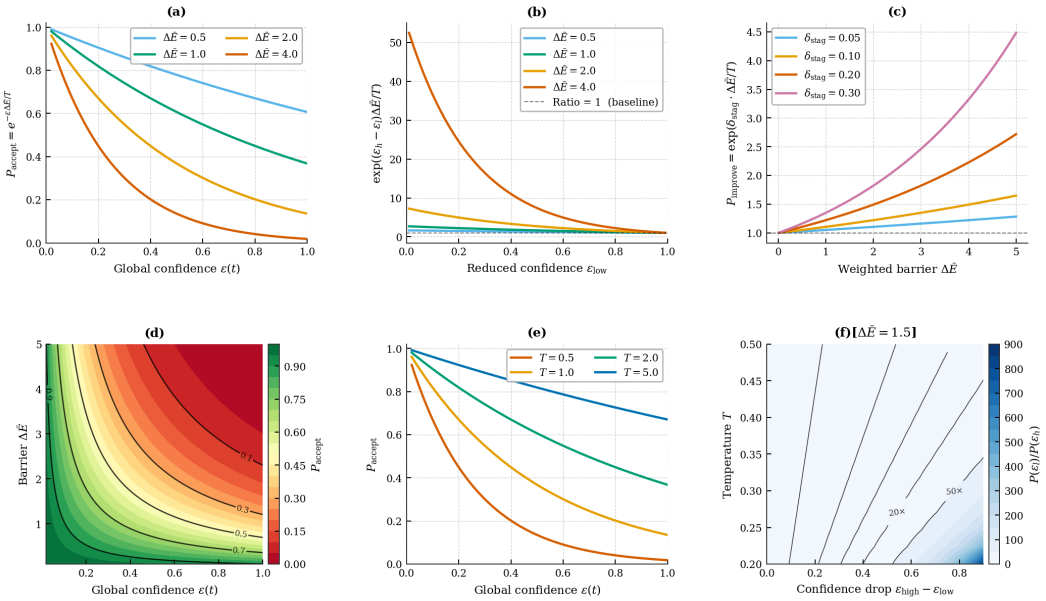


Figure 3: Escape probability improvement and stagnation response. **(a)** $P_{\text{accept}} = e^{-\varepsilon \Delta \bar{E}/T}$ is strictly decreasing in $\varepsilon(t)$ for all barrier heights $\Delta \bar{E}$. **(b)** Acceptance ratio $P(\varepsilon_{\text{low}})/P(\varepsilon_{\text{high}}) > 1$ for all $\varepsilon_{\text{low}} < \varepsilon_{\text{high}} = 1.0$, with larger $\Delta \bar{E}$ yielding greater gain. **(c)** Multiplicative improvement $P_{\text{improve}} = \exp(\delta_{\text{stag}} \cdot \Delta \bar{E}/T)$ grows with both δ_{stag} and $\Delta \bar{E}$. **(d)** Heatmap of P_{accept} over $(\varepsilon, \Delta \bar{E})$ space; contours at $0.1 - 0.9$ intervals. **(e)** P_{accept} at four temperatures with $\Delta \bar{E} = 2.0$ fixed; ε and T act as dual controls. **(f)** Heatmap of acceptance ratio over (δ_ε, T) space with $\Delta \bar{E} = 1.5$; contours label multiplicative gains.

and **(ii)** For any two states $\mathbf{q}$ and $\mathbf{q}'$, $E_\varepsilon(\mathbf{q}) < E_\varepsilon(\mathbf{q}') \Leftrightarrow \bar{E}_\varepsilon(\mathbf{q}) < \bar{E}_\varepsilon(\mathbf{q}')$; where $\bar{E}(\mathbf{q}) = \sum_{j=1}^{m} \eta_j \mathbf{1}[C_j \text{ is unsatisfied by } \mathbf{q}] \geq 0$. We demonstrate this property in Fig. 2.

Escape probability improvement. Topology preservation rules out one failure mode but says nothing about the rate at which the solver escapes a local minimum. The second guarantee establishes that lowering $\varepsilon(t)$ is provably beneficial for escape, with the gain quantified explicitly. Let $\mathcal{N}(\mathbf{q})$ denote the Hamming-1 neighbourhood of $\mathbf{q}$, and let $P_{\text{accept}}(\varepsilon, T_t, \Delta \bar{E}) = \exp\left(-\frac{\varepsilon \cdot \Delta \bar{E}}{T_t}\right)$ be the Metropolis acceptance probability of an uphill move with weighted barrier $\Delta \bar{E} > 0$ at temperature $T_t > 0$ (Fig. 3).

Theorem 3. *Fix any local minimum $\mathbf{q}^*$, any neighbouring state $\mathbf{q}' \in \mathcal{N}(\mathbf{q}^*)$, and any temperature $\mathit{Temp}_t > 0$. Let $\Delta \bar{E} = \Delta \bar{E}(\mathbf{q}^*, i) > 0$ be the fixed weighted energy barrier of the move $\mathbf{q}^* \to \mathbf{q}'$. Then, **(i)** $P_{\text{accept}}(\varepsilon(t), \mathit{Temp}_t, \Delta \bar{E})$ is strictly decreasing in $\varepsilon(t)$; and **(ii)** Reducing $\varepsilon(t)$ from $\varepsilon_{\text{high}}$ to $\varepsilon_{\text{low}} < \varepsilon_{\text{high}}$ yields*

$$\frac{P_{\text{accept}}(\varepsilon_{\text{low}}, \mathit{Temp}_t, \Delta \bar{E})}{P_{\text{accept}}(\varepsilon_{\text{high}}, \mathit{Temp}_t, \Delta \bar{E})} = \exp\left(\frac{(\varepsilon_{\text{high}} - \varepsilon_{\text{low}}) \cdot \Delta \bar{E}}{\mathit{Temp}_t}\right) > 1. \tag{12}$$

Proof. Consider a single bit-flip move from state $\mathbf{q}$ to $\mathbf{q}' = \text{flip}(\mathbf{q}, i)$ at qubit i, we define the *signed clause-change set* as the pair $(\mathcal{J}^+(\mathbf{q}, i), \mathcal{J}^-(\mathbf{q}, i))$, where

$$\mathcal{J}^+(\mathbf{q}, i) = \{j : C_j \text{ is unsatisfied by } \mathbf{q} \text{ but satisfied by } \mathbf{q}'\}, \tag{13a}$$

$$\mathcal{J}^-(\mathbf{q}, i) = \{j : C_j \text{ is satisfied by } \mathbf{q} \text{ but unsatisfied by } \mathbf{q}'\}. \tag{13b}$$

The change in the weighted unsatisfied energy under a single flip is

$$\Delta \bar{E}(\mathbf{q}, i) = \bar{E}(\mathbf{q}') - \bar{E}(\mathbf{q}) = \sum_{j \in \mathcal{J}^-(\mathbf{q}, i)} \eta_j - \sum_{j \in \mathcal{J}^+(\mathbf{q}, i)} \eta_j, \tag{14}$$

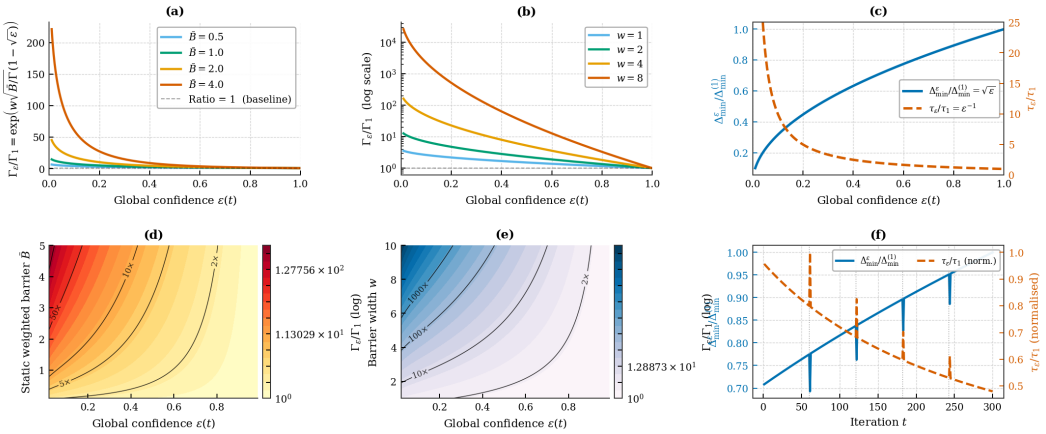


Figure 4: Tunneling rate amplification, exponential amplification in barrier geometry. **(a)** Ratio $\Gamma_\varepsilon/\Gamma_1 > 1$ for all $\varepsilon(t) \in (0,1)$, increasing with barrier height $\bar{B}$ ($w = 3$, $\Gamma = 1$). **(b)** Log-scale plot of $\Gamma_\varepsilon/\Gamma_1$ for varying barrier width w at $\bar{B} = 2.0$: even moderate ε reductions yield exponentially large gains for wide barriers. **(c)** Dual-axis: spectral gap ratio $\Delta^\varepsilon_{\min}/\Delta^{(1)}_{\min} \approx \sqrt{\varepsilon}$ (left) and annealing time ratio $\tau_\varepsilon/\tau_1 = \varepsilon^{-1}$ (right) versus ε. **(d)** Heatmap of $\Gamma_\varepsilon/\Gamma_1$ over $(\varepsilon, \bar{B})$ space ($w = 3$); log-scale colorbar. **(e)** Heatmap of $\Gamma_\varepsilon/\Gamma_1$ over (ε, w) space ($\bar{B} = 2.0$); log-scale colorbar. **(f)** Spectral gap ratio and normalised annealing time ratio along an adaptive $\varepsilon(t)$ schedule; vertical dotted lines mark stagnation events.

and the change in the full two-level energy is

$$\Delta E_\varepsilon(\mathbf{q}, i) \;=\; E_\varepsilon(\mathbf{q}') - E_\varepsilon(\mathbf{q}) \;=\; \underbrace{\varepsilon(t)\bar{E}(\mathbf{q}') - \varepsilon(t)\bar{E}(\mathbf{q})}_{\text{Claim 1}} = \varepsilon(t) \cdot \underbrace{\Delta\bar{E}(\mathbf{q}, i)}_{\text{by Eq. 14}}. \tag{15}$$

Lemma 1. *Let $\mathcal{N}(\mathbf{q}) = \{\mathbf{q}' \in \{0,1\}^Q : \|\mathbf{q} - \mathbf{q}'\|_1 = 1\}$ be the Hamming-1 neighbourhood. A state $\mathbf{q}^*$ is a local minimum of $\bar{E}$ i.f.f $\bar{E}(\mathbf{q}^*) \le \bar{E}(\mathbf{q}')$ for all $\mathbf{q}' \in \mathcal{N}(\mathbf{q}^*)$. Here, $\bar{E}(\mathbf{q}^*) \le \bar{E}(\mathbf{q}')$ i.f.f $E_\varepsilon(\mathbf{q}^*) \le E_\varepsilon(\mathbf{q}')$. When $\mathbf{q}$ is a local (non-global) minimum of $\bar{E}$ and $\mathbf{q}' \in \mathcal{N}(\mathbf{q})$, we have $\Delta\bar{E}(\mathbf{q}, i) > 0$ for every neighbour $\mathbf{q}'$ (by definition of local minimum).*

Definition 1. *The Metropolis–Hastings acceptance probability for the escape move $\mathbf{q} \to \mathbf{q}'$ at iteration-temperature $Temp_t$ is:*

$$P_{\text{accept}}\big(\varepsilon(t),\, Temp_t,\, \Delta\bar{E}\big) = \exp\!\left(-\frac{\varepsilon(t) \cdot \Delta\bar{E}(\mathbf{q}, i)}{Temp_t}\right). \tag{16}$$

Note that $\Delta\bar{E}(\mathbf{q}, i) = \sum_{j \in \mathcal{J}^-(\mathbf{q},i)} \eta_j - \sum_{j \in \mathcal{J}^+(\mathbf{q},i)} \eta_j > 0$ is a weighted sum of the static clause confidences η_j, and remains constant as $\varepsilon(t)$ changes. Only the global scalar $\varepsilon(t)$ is modulated by the adaptive schedule.

We define $f(s) = \exp(-s \cdot \Delta\bar{E}/\text{Temp}_t)$ for $s > 0$. Clearly, $f(s) \in [0,1]$ due to the axiom of probability functions. Since $\Delta\bar{E} > 0$ and $\text{Temp}_t > 0$, the derivative

$$\frac{df}{ds} = f'(s) = -\left(\frac{\Delta\bar{E}}{\text{Temp}_t}\right) f(s) < 0, \quad (\forall f(s) > 0) \tag{17}$$

so f is strictly decreasing. Thus, **(i)** holds.

For $\varepsilon_{\text{low}} < \varepsilon_{\text{high}}$:

$$\frac{f(\varepsilon_{\text{low}})}{f(\varepsilon_{\text{high}})} = \frac{\exp\big(-\varepsilon_{\text{low}}\, \Delta\bar{E}/\text{Temp}_t\big)}{\exp\big(-\varepsilon_{\text{high}}\, \Delta\bar{E}/\text{Temp}_t\big)} = \exp\!\left(\frac{(\varepsilon_{\text{high}} - \varepsilon_{\text{low}})\, \Delta\bar{E}}{\text{Temp}_t}\right) > 1 \tag{18}$$

since $\varepsilon_{\text{high}} - \varepsilon_{\text{low}} > 0$, $\Delta\bar{E} > 0$ and $\text{Temp}_t > 0$ by assumption. Thus, reducing $\varepsilon(t)$ yields $P_{\text{accept}}(\varepsilon_{\text{low}}) > P_{\text{accept}}(\varepsilon_{\text{high}})$. □

Tunneling rate amplification. The escape-probability result is purely classical, governing the acceptance of bit-flip moves under the Metropolis rule. The third guarantee transports the same intuition to the quantum setting: on physical annealers, the same reduction in $\varepsilon(t)$ that sharpens classical escape also accelerates quantum tunneling through the barrier between local minima (Fig. 4).

Theorem 4 (Theoretical prediction under WKB approximation)**.** *Let Γ_ε and Γ_1 denote the quantum tunneling rates under energies E_ε with global confidence $\varepsilon(t) \in (0,1)$ and $\bar{E}$ with global confidence 1, respectively. For the same static weighted barrier $\bar{B} > 0$, barrier width $w \geq 1$, and transverse field $\Gamma > 0$. Then:*

$$\frac{\Gamma_\varepsilon}{\Gamma_1} = \exp\left(w\sqrt{\bar{B}/\Gamma}\left(1 - \sqrt{\varepsilon(t)}\right)\right) > 1 \quad \textit{for all } \varepsilon(t) \in (0,1). \tag{19}$$

Proof. The Hamiltonian interpolates between a transverse-field driver and the problem Hamiltonian encoding E_ε. The tunneling rate between two local minima $\mathbf{q}^*$ and $\mathbf{q}^\dagger$ separated by an energy barrier of height B and width w (the minimum number of simultaneous qubit flips on any path between $\mathbf{q}^*$ and $\mathbf{q}^\dagger$) obeys the WKB-type estimate see (Chapter 13 Griffiths & Schroeter (2018)):

$$\Gamma_{\text{tunnel}} \propto \exp\left(-w\sqrt{B/\Gamma}\right), \tag{20}$$

where $\Gamma > 0$ is the transverse field strength. Under the two-level energy E_ε, the barrier height between $\mathbf{q}^*$ and $\mathbf{q}^\dagger$ along the minimum-barrier path $\mathbf{q}^* = \mathbf{q}^{(0)}, \mathbf{q}^{(1)}, \ldots, \mathbf{q}^{(w)} = \mathbf{q}^\dagger$ is

$$B_\varepsilon \;=\; \max_{0\leq s\leq w} E_\varepsilon(\mathbf{q}^{(s)}) - E_\varepsilon(\mathbf{q}^*) \;=\; \varepsilon(t)\cdot\left(\max_{0\leq s\leq w} \bar{E}(\mathbf{q}^{(s)}) - \bar{E}(\mathbf{q}^*)\right) \;=:\; \varepsilon(t)\cdot\bar{B}, \tag{21}$$

where

$$\bar{B} = \max_{0\leq s\leq w} \sum_{j=1}^{m} \eta_j \cdot \mathbf{1}[C_j \text{ unsat by } \mathbf{q}^{(s)}] - \sum_{j=1}^{m} \eta_j \cdot \mathbf{1}[C_j \text{ unsat by } \mathbf{q}^*] \tag{22}$$

is *the static weighted barrier*, determined entirely by the fixed clause confidences $\{\eta_j\}$ and the SAT instance structure, independently of $\varepsilon(t)$.

Substituting $B_\varepsilon = \varepsilon(t)\cdot\bar{B}$ in Eq. 22 and $B_1 = 1\cdot\bar{B} = \bar{B}$ into Eq. 20 with common prefactor $C > 0$:

$$\Gamma_\varepsilon = C\exp\left(-w\sqrt{\varepsilon(t)\,\bar{B}/\Gamma}\right), \qquad \Gamma_1 = C\exp\left(-w\sqrt{\bar{B}/\Gamma}\right). \tag{23}$$

Therefore:

$$\frac{\Gamma_\varepsilon}{\Gamma_1} = \exp\left(w\sqrt{\bar{B}/\Gamma}\left(1 - \sqrt{\varepsilon(t)}\right)\right). \tag{24}$$

Since $\varepsilon(t) \in (0,1)$, we have $\sqrt{\varepsilon(t)} \in (0,1)$, so $1 - \sqrt{\varepsilon(t)} > 0$. Combined with $w, \bar{B}, \Gamma > 0$, the exponent is strictly positive. Hence $\Gamma_\varepsilon/\Gamma_1 > 1$. □

D Proofs of Theorems

Every statement below was additionally checked against the deployed implementation. The identity of Theorem 1 was verified exhaustively over $\{0,1\}^k$ and $\{0,1\}^n$ on 600 instances, giving 58,042 identity checks with a worst relative deviation of 3.4×10^{-16}.

Proof of Theorem 1. Fix $\boldsymbol{\xi} \in \{0,1\}^k$ and put $\mathbf{s} = A^\top\boldsymbol{\xi}$, so that $s_i = \xi_{g(i)}$. Let C_j be a clause with image D_p, that is $\Pi_{pj} = 1$. Substituting into the violation monomial $\chi_j(\mathbf{s}) = \prod_{i\in C_j^+}(1 - \xi_{g(i)})\prod_{i\in C_j^-}\xi_{g(i)}$. Two variables of C_j with $g(i) = g(i')$ and equal polarity contribute the same factor twice, and $v^2 = v$ on $\{0,1\}$ collapses it to one factor, which is exactly the narrowing that defines D_p. Two variables with $g(i) = g(i')$ and opposite polarity contribute $\xi_r(1-\xi_r) = 0$, so the monomial vanishes identically, which is exactly the tautology rule that drops the image. In either case $\chi_j(A^\top\boldsymbol{\xi}) = \chi_p^A(\boldsymbol{\xi})$, with $\chi_p^A \equiv 0$ for a dropped image. Summing against $\boldsymbol{\omega}$ and collecting the clauses that share an image $U_{\boldsymbol{\omega}}(A^\top\boldsymbol{\xi}) = \sum_{j=1}^m \omega_j\chi_j(A^\top\boldsymbol{\xi}) = \sum_{p=1}^{m_A}\left(\sum_{j:\Pi_{pj}=1}\omega_j\right)\chi_p^A(\boldsymbol{\xi}) = U^A_{\Pi\boldsymbol{\omega}}(\boldsymbol{\xi})$,

which is the identity, and the inner sum is where the additivity of merged weights is used. Since $\boldsymbol{\omega} > 0$, the identity gives $\boldsymbol{\xi} \models \varphi_A \iff U^A_{\Pi\boldsymbol{\omega}}(\boldsymbol{\xi}) = 0 \iff U_{\boldsymbol{\omega}}(A^\top \boldsymbol{\xi}) = 0 \iff A^\top \boldsymbol{\xi} \models \varphi$. The map $A^\top$ is injective on $\{0,1\}^k$, because the groups G_r are non-empty and disjoint, so $\boldsymbol{\xi}$ is recovered from $A^\top \boldsymbol{\xi}$ by reading any coordinate of each group, and its image is $\mathcal{R}(A)$ by definition. Restricting to solutions gives the bijection onto $\mathrm{sol}(\varphi) \cap \mathcal{R}(A)$. Finally $\min_{\boldsymbol{\xi}} U^A_{\Pi\boldsymbol{\omega}} = \min_{\mathbf{s} \in \mathcal{R}(A)} U_{\boldsymbol{\omega}}(\mathbf{s}) \geq \min_{\mathbf{s} \in \{0,1\}^n} U_{\boldsymbol{\omega}}(\mathbf{s})$, since $\mathcal{R}(A) \subseteq \{0,1\}^n$. □

Proof of Theorem 2. Because g is chosen independently of φ, each of the 2^k points $A^\top \boldsymbol{\xi}$ is a fixed assignment with respect to the randomness of the formula. A uniform random K-clause is violated by a fixed assignment with probability 2^{-K}, so it is satisfied with probability $1 - 2^{-K}$, and the m clauses are independent. Hence, $\mathbb{P}(A^\top \boldsymbol{\xi} \models \varphi) = (1 - 2^{-K})^m$ for each $\boldsymbol{\xi}$ and, by linearity, $\mathbb{E}[Z_A] = 2^k (1 - 2^{-K})^m$, with equality and not merely a bound. Since Z_A is a non-negative integer, Markov's inequality gives $\mathbb{P}(\varphi_A \text{ satisfiable}) = \mathbb{P}(Z_A \geq 1) \leq \mathbb{E}[Z_A]$. Writing $m = \alpha n$ and $\kappa = \alpha \log_2 (1 - 2^{-K})^{-1}$,

$$2^k (1 - 2^{-K})^{\alpha n} = 2^{\,k - \kappa n},$$

so $k \leq (1-\delta)\kappa n$ gives the exponent $-\delta \kappa n$. At $K = 3$ and $\alpha = \alpha_c = 4.267$, $\kappa = 4.267 \log_2(8/7) = 0.822$. □