Spectral–Entropic Rigidity Certificates for Structured NP Instance Families

Fourier Diagnostics, Wasserstein–KL Control, and Flow-and-Extract Bounds

Version v1.0

Lee Byoungwoo

June 2, 2026

Abstract

Many NP-complete problems contain structured instance families that are substantially easier than worst-case complexity suggests. This manuscript develops a spectral–entropic rigidity framework for diagnosing and certifying such effective subfamilies through four separate gates: structural admissibility, spectral or reference-law control, entropy–metric or finite-closure contraction, and extractor validity.

The paper deliberately does not address the global P versus NP problem and does not introduce new solvers for classical tractable families. Its contribution is a certificate calculus that makes explicit how known tractability mechanisms correspond to non-circular reference-law construction, obstruction detection, finite or metric collapse, and verified extraction. Four calibration mechanisms are developed: product decomposition for block-factorized CSPs, junction-tree decomposition for bounded-treewidth CSPs, algebraic duality for affine systems over $\mathbb{F}_2$, and implication closure for 2-SAT.

The manuscript also includes two theorem-level refinements to the structure-to-spectrum layer. First, for affine feasible sets over $\mathbb{F}_2$, the entire Boolean spectral heat-decay profile is expressed exactly in terms of the dual weight enumerator. Second, a low-overlap obstruction theorem shows that even disjoint monotone 2-CNF clauses can yield feasibility indicators whose normalized Fourier mass escapes every fixed low-degree cutoff. Thus low overlap alone is not a valid structure-to-spectrum certificate for feasibility indicators.

The analytic layer gives Wasserstein–KL contraction with spectral, covariance, discretization, statistical, modeling, and bounded-variation correctors; the finite-closure layer handles exact graph or algebraic obstruction procedures. A conditional flow-and-extract theorem supplies explicit time budgets once the relevant gates and a positive extraction margin are certified. A predictive structural classifier is included only as a diagnostic and reproducibility protocol, not as a proof oracle.

Contents

1 Introduction

NP-complete problems are defined by their worst-case complexity, but many concrete instances encountered in practice are substantially more structured than arbitrary worst-case inputs. Such instances may exhibit bounded interaction width, low overlap, rank degeneracy, local product structure, symmetry, or other forms of redundancy. Classical algorithmic theory already exploits many of these structures through dynamic programming, local consistency, message passing, or decomposition methods. The purpose of this manuscript is not to replace those theories, nor to address the global P versus NP problem. Rather, it develops a spectral–entropic certificate language for describing when such structured instances admit effective tractability.

The guiding principle is that an instance should not be called effectively rigid merely because a smoothing flow suppresses high-frequency information. Heat-type smoothing alone always reduces high-degree components and therefore cannot by itself constitute an algorithmic certificate. A useful rigidity certificate must include four logically separate gates:

(G1) a *structural gate*, certifying that the instance has controlled combinatorial complexity;

(G2) a *spectral or reference-law gate*, certifying low-degree concentration, covariance control, or a non-circular reference law;

(G3) an *entropy–metric gate*, certifying contraction toward a reference rigidity basin under a declared flow, Markov semigroup, or message evolution;

(G4) an *extractor gate*, certifying that proximity to the rigidity basin yields an actual witness or verified solution.

This separation is the central safeguard of the paper. Entropy collapse is treated as a diagnostic and geometric phenomenon. Algorithmic meaning enters only after a non-circular reference law and a valid extractor are specified.

The manuscript develops four calibration mechanisms. Block-factorized CSPs give a product reference law and exact local extraction. Bounded-treewidth CSPs give a polynomial-size junction-tree representation, infeasibility certificate, and exact backward sampler. Affine systems over $\mathbb{F}_2$ give an algebraic spectral calibration through the dual weight profile. Finally, 2-SAT gives a non-treewidth implication-closure mechanism. Although these families are classically tractable, they are useful as calibration cases because the relevant chains from structure to representation, finite or metric collapse, and verified extraction close without hidden solution information.

The manuscript further consolidates two tests of the structure-to-spectrum gate. On the positive side, the affine calibration is strengthened from support control to an exact heat-decay formula: the centered spectral heat profile of an affine indicator is the dual weight enumerator evaluated along the Boolean heat parameter. On the negative side, a disjoint-clause CNF construction proves that low overlap alone does not imply low-degree Fourier concentration for feasibility indicators. This obstruction explains why the framework treats additive local energies, reference-law representations, and global feasibility indicators as distinct objects.

The resulting framework should be read as a conditional theory of structured NP instance families. It does not prove $P = NP$, $P \neq NP$, or a polynomial-time algorithm for arbitrary NP instances. Its contribution is a certificate architecture connecting known structural tractability, spectral concentration, entropy contraction, and extraction stability in a common language.

1.1 Contributions

The main contributions are as follows.

(C1) A four-gate rigidity certificate separating structural admissibility, spectral/reference-law control, entropy–metric contraction, and extractor validity.

(C2) A block-factorized CSP anchor theorem showing complete non-circular closure of the four gates.

(C3) A bounded-treewidth CSP certificate theorem recasting junction-tree tractability as reference-law construction plus exact extraction.

(C4) A structure-to-spectrum gate that carefully distinguishes additive bounded-arity energies from global feasibility indicators.

(C5) An affine low-weight dual theorem giving an exact Fourier-tail formula for affine feasible indicators over $\mathbb{F}_2$.

(C6) An affine heat-profile theorem expressing the whole Boolean spectral heat-decay curve in terms of the dual weight enumerator.

(C7) A 2-SAT implication-graph rigidity theorem showing a non-treewidth finite-closure certificate with exact extraction, with the finite closure gate explicitly distinguished from the Wasserstein–KL entropy-flow gate.

(C8) A low-overlap CNF obstruction theorem showing that disjoint monotone 2-CNF clauses need not yield low-degree Fourier concentration of the feasibility indicator.

(C9) A Wasserstein–KL and perturbed rigidity-contraction layer with explicit error envelopes.

(C10) A flow-and-extract conditional theorem converting certified contraction and extractor validity into an explicit time budget.

(C11) A diagnostic PSC and reproducibility protocol for exploratory rigid/transitional/nonrigid classification.

1.2 Relation to standard methods

This manuscript does not replace classical tractability theory. Bounded-treewidth CSPs are polynomial-time solvable by dynamic programming, tree-decomposition, hypertree-width, and junction-tree methods [12, 13, 14, 6, 7]; 2-SAT is solved by implication-graph strongly connected components and closure methods going back to Krom and later linear-time graph algorithms [10, 11, 5]; Boolean Fourier terminology follows the standard Walsh–Fourier analysis of Boolean functions [8]; graphical-model language is standard in probabilistic inference [17, 18]; and the Wasserstein–KL layer is based on the transport–entropy and log-Sobolev framework developed in optimal transport, Markov-chain, and gradient-flow theory [19, 21, 1, 2, 3, 20, 4]. The broader complexity and CSP background is classical [22, 23, 24, 15, 16]. The contribution here is organizational and certifying: these ingredients are separated into structural, spectral/reference-law, entropy–metric, and extractor gates so that smoothing, contraction, and witness recovery are not conflated. The affine dual-weight section also uses standard coding-theoretic weight-enumerator language [25].

2 Rigidity Certificates

Let I be a Boolean constraint instance on variables $x_1, \dots, x_n$, and let $\mathcal{A}(I) \subseteq \{0,1\}^n$ denote its feasible set. We write $V_I(y) \in \{0,1\}$ for the polynomial-time verifier associated with I.

Definition 2.1 (Rigidity certificate)**.** *A rigidity certificate for an instance I is a quadruple*

$$\mathsf{Cert}(I) = (\mathsf{Struct}, \mathsf{Spec}, \mathsf{Flow}, \mathsf{Ext})$$

satisfying the following four gates.

Gate 1. **Structural gate.** Struct *certifies that I lies in a declared structured family, such as a block-factorized family, bounded-treewidth family, bounded-overlap family, low-rank affine family, or another explicitly defined class.*

Gate 2. **Spectral/reference-law gate.** Spec *supplies either a spectral concentration statement for an associated observable, a covariance-tail estimate for a probability law, or a non-circular construction of a reference law or rigidity basin.*

Gate 3. **Entropy–metric or finite-closure gate.** Flow *supplies either an entropy–metric contraction estimate for a declared flow, semigroup, or message evolution,*

$$d_{\text{rigid}}(\mu_t, \mathcal{M}_{\text{rigid}}) \le Ae^{-\gamma t} d_{\text{rigid}}(\mu_0, \mathcal{M}_{\text{rigid}}) + B\eta_{[0,t]},$$

or a finite exact closure map with an obstruction or residual certificate that vanishes after a polynomial-time computation. Here $\eta_{[0,t]}$ is an explicit error envelope in the metric-flow case.

Gate 4. **Extractor gate.** Ext *is a randomized or deterministic extraction procedure such that*

$$d_{\text{rigid}}(\mu, \mathcal{M}_{\text{rigid}}) \le \varepsilon_{\text{round}} \quad \Longrightarrow \quad \Pr[V_I(\mathsf{Ext}(\mu)) = 1] \ge 1 - \alpha.$$

In a finite-closure certificate, the extractor gate may instead be triggered by a certified zero residual or obstruction-free closure state, and the extracted object must still be verified by V_I.

The reference law, rigidity basin, and extractor must be constructible without encoding a hidden global solution.

Definition 2.2 (Certified effective families: search and decision forms)**.** *A feasible-instance family $\mathcal{F} = \{I_n\}$ is called* search-certified effective *under the rigidity framework if there exists a polynomial-time procedure that, for every feasible $I_n \in \mathcal{F}$, constructs a rigidity certificate with parameters satisfying*

$$\gamma^{-1}, \qquad (\varepsilon_{\text{round}} - B\eta)^{-1}, \qquad \alpha^{-1}$$

bounded by at most a polynomial in n, and whose extraction step returns a verified witness with probability at least $1 - \alpha$.

A family is called decision-certified effective *if the same procedure either returns a verified witness when the instance is feasible or certifies infeasibility when no witness exists.*

This definition does not assert that all NP instances are effectively rigid. It only defines certified subfamilies in which the relevant gates are verified. The distinction between search and decision forms is useful because product examples are stated for nonempty local feasible sets, whereas bounded-treewidth dynamic programming can also certify infeasibility.

3 Anchor Case: Block-Factorized Bounded-Arity CSPs

We first verify the certificate architecture in a simple but fully closed setting.

Definition 3.1 (Block-factorized bounded-arity CSP)**.** *Let I be a Boolean constraint instance on variables $x_1, \dots, x_n$. We say that I is block-factorized with arity bound q if there exists a partition*

$$[n] = B_1 \sqcup \cdots \sqcup B_m, \qquad |B_j| \le q,$$

and nonempty local feasible sets

$$\mathcal{A}_j \subseteq \{0,1\}^{B_j}$$

such that the global feasible set factors as

$$\mathcal{A}(I) = \mathcal{A}_1 \times \cdots \times \mathcal{A}_m.$$

We assume q is fixed independently of n, and that each $\mathcal{A}_j$ is computed by direct enumeration of the $2^{|B_j|}$ local assignments.

Define the reference law

$$\mu_\infty = \bigotimes_{j=1}^{m} \mathrm{Unif}(\mathcal{A}_j).$$

For the transport statement below we use the normalized product metric

$$d(x,y)^2 = \frac{1}{m}\sum_{j=1}^{m} d_j(x_j,y_j)^2,$$

where each local metric d_j has diameter at most one. With the unnormalized Hamming metric, the transport constant should instead be allowed to scale with m.

Theorem 3.2 (Block-factorized rigidity certificate)**.** *Let $\mathcal{F}_q$ be the family of block-factorized bounded-arity CSP instances with fixed arity bound q. Then $\mathcal{F}_q$ is search-certified effective under the rigidity framework.*

More precisely, for every $I \in \mathcal{F}_q$:

(i) the reference law μ_∞ is constructible in time $O(m2^q)$;

(ii) μ_∞ does not encode hidden global solution information beyond locally enumerated feasible sets;

(iii) the product entropy flow with invariant law μ_∞ satisfies

$$\mathrm{KL}(\mu_t\|\mu_\infty) \le e^{-2\lambda_q t}\mathrm{KL}(\mu_0\|\mu_\infty),$$

for some $\lambda_q > 0$ depending only on q;

(iv) the transport distance is controlled by entropy:

$$W_2^2(\mu_t,\mu_\infty) \le C_q \mathrm{KL}(\mu_t\|\mu_\infty);$$

(v) the extractor obtained by independently sampling $y_j \sim \mathrm{Unif}(\mathcal{A}_j)$ and concatenating $y = (y_1,\dots,y_m)$ returns a valid witness with probability 1.

Consequently, the search certificate is non-circular and runs in polynomial time for fixed q. If empty local feasible sets are also allowed and detected by enumeration, the same construction gives a decision-certified variant.

Proof. For each block B_j, the set $\mathcal{A}_j$ is obtained by checking at most 2^q assignments. Hence all local feasible sets are computed in total time $O(m2^q)$, which is polynomial in n for fixed q. Since the global feasible set is assumed to factor as $\prod_j \mathcal{A}_j$, the product law $\mu_\infty = \bigotimes_j \mathrm{Unif}(\mathcal{A}_j)$ is constructed from local information only.

On each finite local state space $\mathcal{A}_j$, choose the standard reversible heat-bath or complete-graph Markov chain with invariant law $\mathrm{Unif}(\mathcal{A}_j)$. Since $|\mathcal{A}_j| \le 2^q$, the local log-Sobolev constant is bounded below by a positive number depending only on q. Denote this lower bound by λ_q. Tensorization of the product chain gives the same lower bound for the product log-Sobolev constant. Along the corresponding entropy gradient flow or continuous-time Markov semigroup,

$$\frac{d}{dt}\mathrm{KL}(\mu_t\|\mu_\infty) \le -2\lambda_q \mathrm{KL}(\mu_t\|\mu_\infty),$$

and Gronwall's inequality gives the stated entropy decay. With the normalized product metric, tensorization of the finite-product transport–entropy inequality gives the displayed W_2 control, with C_q depending only on the local geometry and the fixed arity bound. With an unnormalized product metric, the same statement holds with the corresponding dimension-dependent metric constant. Finally, independent local sampling produces an element of $\prod_j \mathcal{A}_j = \mathcal{A}(I)$, so the verifier accepts with probability 1. $\square$

Remark 3.3 (Calibrational role)**.** *The theorem is not intended as a new algorithmic result for product CSPs. Its role is calibrational: it shows that the four-gate rigidity certificate can close without circularity in a transparent structured family.*

4 Bounded-Treewidth CSPs and Junction-Tree Rigidity

The block-factorized case verifies the architecture in a product setting. We now move to a less trivial but still classically tractable family: bounded-treewidth CSPs.

4.1 Bounded-treewidth CSPs

Let I be a Boolean CSP instance on variables $x_1, \ldots, x_n$, with constraints $\{C_\alpha\}_{\alpha\in A}$. Each constraint C_α depends on a subset $S_\alpha \subseteq [n]$ of variables, with $|S_\alpha| \leq q$, where q is fixed.

The primal graph G_I has vertex set $[n]$, with an edge between i and j whenever some constraint involves both x_i and x_j. A tree decomposition of G_I is a tree $T = (\mathcal{B}, E_T)$, whose nodes are bags $B \in \mathcal{B}$, satisfying:

(i) every variable $i \in [n]$ belongs to at least one bag;

(ii) for every constraint scope S_α, there exists a bag $B \in \mathcal{B}$ such that $S_\alpha \subseteq B$;

(iii) for every variable i, the set of bags containing i forms a connected subtree of T.

The width is $\max_{B\in\mathcal{B}} |B| - 1$. We say I has bounded treewidth w if a tree decomposition of width at most w is given or computable in polynomial time, with w fixed independently of n.

Let

$$\mathcal{A}(I) = \{x \in \{0,1\}^n : C_\alpha(x_{S_\alpha}) = 1 \text{ for all } \alpha\}$$

be the feasible set.

4.2 Junction-tree reference law

Assume first that $\mathcal{A}(I) \neq \emptyset$. Define the uniform feasible law

$$\mu_\infty = \mathrm{Unif}(\mathcal{A}(I)).$$

This definition is algorithmically meaningful only because bounded treewidth supplies a non-circular construction.

For each bag B, let $\psi_B(x_B) \in \{0,1\}$ be the product of all local constraints assigned to that bag. Sum-product messages are defined along oriented edges $B \to B'$ by

$$m_{B\to B'}(x_{B\cap B'}) = \sum_{x_{B\setminus B'}} \psi_B(x_B) \prod_{A\in N(B)\setminus\{B'\}} m_{A\to B}(x_{A\cap B}).$$

Because $|B| \leq w+1$, each message table has size at most 2^{w+1}. A standard inward–outward pass computes the partition function $Z_I = |\mathcal{A}(I)|$, exact bag/separator marginals, and the conditional tables required for exact backward sampling in time $O(|\mathcal{B}|2^{w+1}\mathrm{poly}(q))$. Thus $\mu_\infty = \mathrm{Unif}(\mathcal{A}(I))$ is not output as an explicit table over all satisfying assignments. It is represented non-circularly by a polynomial-size junction-tree message representation, including the partition function, local marginals, and an exact backward sampler. If $Z_I = 0$, the instance is certified infeasible. If $Z_I > 0$, backward sampling through the tree decomposition gives an exact sample from μ_∞.

Theorem 4.1 (Bounded-treewidth rigidity certificate). *Let $\mathcal{F}_{w,q}$ be a family of Boolean CSP instances with arity at most q and treewidth at most w, where w, q are fixed. Assume that a width-w tree decomposition is given or computable in polynomial time.*

Then $\mathcal{F}_{w,q}$ is decision-certified effective under the rigidity framework. More precisely, for every instance $I \in \mathcal{F}_{w,q}$:

(i) the structural gate is certified by the given tree decomposition;

(ii) the reference law $\mu_\infty = \mathrm{Unif}(\mathcal{A}(I))$, when $\mathcal{A}(I) \neq \emptyset$, is represented non-circularly by a polynomial-size junction-tree message representation, including the partition function, local marginals, and an exact backward sampler;

(iii) infeasibility is certified when the dynamic-programming partition function Z_I vanishes;

(iv) if $Z_I > 0$, an exact sampler for μ_∞ is constructible in polynomial time;

(v) the extractor obtained by exact junction-tree sampling returns a valid witness with probability 1;

(vi) the entropy–metric gate may be realized as a finite message-flow gate: the exact inward–outward schedule reaches the dynamic-programming fixed point after finitely many sweeps, so the message residual vanishes exactly;

(vii) all construction and extraction costs are bounded by $O(|\mathcal{B}|2^{w+1}\mathrm{poly}(q))$, hence polynomial in n for fixed w, q.

Consequently, bounded-treewidth CSPs satisfy a decision-certified four-gate rigidity certificate non-circularly.

Proof. The structural gate is immediate from the assumed tree decomposition. Since w is fixed, every bag contains at most $w + 1$ variables. Assign each constraint to one bag containing its scope. This is possible by the definition of tree decomposition. The product of these assigned local factors yields a decomposable junction-tree representation of the global feasibility indicator.

Run the standard sum-product recursion on the tree. Since messages are indexed by separator assignments $x_{B\cap B'}$, each message has size at most 2^{w+1}, and each message is computed by summing over at most 2^{w+1} bag assignments. The total cost is $O(|\mathcal{B}|2^{w+1}\mathrm{poly}(q))$. The root table gives the partition function $Z_I = |\mathcal{A}(I)|$. If $Z_I = 0$, infeasibility is certified.

If $Z_I > 0$, the inward messages determine exact conditional laws for backward sampling. Starting from a root bag, sample a root assignment according to its marginal. Then recursively sample each neighboring bag conditioned on the already sampled separator assignment. The running-intersection property ensures global consistency. Every constraint is assigned to a bag whose local factor was enforced, so the sampled full assignment satisfies every constraint. The verifier accepts with probability 1. The same dynamic-programming recursion is an exact finite message flow on the acyclic junction tree; after the prescribed inward–outward passes, all messages equal their fixed-point values and the residual defined below is zero. No hidden global solution is inserted into the reference law or sampler. □

4.3 Entropy and message-space interpretation

The preceding theorem already gives a complete certificate. One can also view the junction-tree computation as a finite-dimensional message flow on local tables. Let $M^{(\ell)}$ denote the collection of messages after ℓ scheduled updates, and let M^* be the exact fixed point obtained by dynamic programming. Define a message residual

$$\mathcal{R}(M^{(\ell)}) = \sum_{B \to B'} \| m^{(\ell)}_{B \to B'} - m^*_{B \to B'} \|_1.$$

For an acyclic junction tree, an inward–outward schedule reaches M^* after finitely many passes. Thus rigidity is not merely asymptotic: the message residual vanishes after a finite number of exact dynamic-programming sweeps.

Remark 4.2. *The bounded-treewidth theorem is not presented as a new algorithm for bounded-treewidth CSPs. Its role is foundational and calibrational. It demonstrates that the rigidity-certificate language can reproduce a known tractable family without circularity.*

5 The Structure-to-Spectrum Gate

We now discuss the broader structure-to-spectrum principle that motivated the framework. The purpose of this section is deliberately limited. We do not claim that every structurally simple NP instance has low-degree Fourier concentration in its feasibility indicator. Such a claim is false in general or, at minimum, requires additional hypotheses.

5.1 Fourier tail and spectral concentration

Let $f : \{0,1\}^n \to \mathbb{R}$ be an observable associated with an instance I. Its Walsh–Fourier expansion is

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \chi_S(x), \qquad \chi_S(x) = (-1)^{\sum_{i \in S} x_i},$$

with normalization

$$\widehat{f}(S) = 2^{-n} \sum_{x \in \{0,1\}^n} f(x) \chi_S(x).$$

All Fourier tails in this section use this normalization and the associated Parseval identity. For $k \geq 0$, define the high-degree Fourier tail

$$\mathrm{Tail}_{>k}(f) = \sum_{|S| > k} \widehat{f}(S)^2.$$

If $f \not\equiv 0$, define the normalized tail ratio

$$\tau_k(f) = \frac{\mathrm{Tail}_{>k}(f)}{\sum_{S \subseteq [n]} \widehat{f}(S)^2}.$$

We say that f is (k, ε)-spectrally concentrated if $\tau_k(f) \leq \varepsilon$.

5.2 Safe observable: additive violation energy

For a CSP instance I with constraints $\{C_\alpha\}_{\alpha \in A}$, define the additive violation energy

$$E_I(x) = \sum_{\alpha \in A} \left(1 - C_\alpha(x_{S_\alpha})\right).$$

Proposition 5.1 (Bounded arity gives exact spectral truncation for additive energies)**.** *Let I be a Boolean CSP instance of arity at most q, and let E_I be its additive violation energy. Then*

$$\widehat{E_I}(S) = 0 \qquad \text{for every } |S| > q.$$

Equivalently, $\mathrm{Tail}_{>q}(E_I) = 0$.

Proof. Each local violation term depends on at most q Boolean variables. Its Walsh–Fourier expansion is supported only on subsets of its scope, and therefore only on degrees at most q. Since E_I is a finite sum of such local functions, no coefficient of degree greater than q can appear. $\square$

5.3 Why feasibility indicators are harder

The feasible-set indicator is

$$1_{\mathcal{A}(I)}(x) = \prod_{\alpha \in A} C_\alpha(x_{S_\alpha}).$$

Even if every C_α has arity at most q, this product can generate Fourier terms of degree as large as n. Therefore bounded arity alone controls additive observables, local energies, and message updates, but not necessarily the global feasibility indicator.

5.4 General structure-to-spectrum gate

Hypothesis 5.2 (Structure-to-spectrum gate)**.** *For a declared structured family $\mathcal{F}$, there exist computable structural features*

$$S(I) = (s_1(I), \dots, s_d(I))$$

and thresholds depending on a parameter box PB *such that, whenever $S(I)$ lies in the admissible region, the chosen observable or law satisfies a certified spectral concentration estimate*

$$\mathrm{Tail}_{>k}(f_I) \le \varepsilon \|f_I\|_2^2,$$

or an analogous covariance/spectral-tail bound for the associated probability law.

The observable f_I, the degree cutoff k, and the error ε must be specified as part of the certificate.

This hypothesis is not assumed for arbitrary NP instances. It is a family-level gate that must be proved, estimated, or rejected for each proposed class.

6 Wasserstein–KL Control and Perturbed Rigidity Contraction

We now formulate the entropy–metric gate. We separate transport distance, information distance, and spectral or covariance deviation.

6.1 Reference law and entropy

Let (X, d) be a finite or compact metric space associated with a structured instance family. Let $\mu_\infty \in \mathcal{P}(X)$ be a reference law constructed by one of the non-circular mechanisms described earlier. For $\mu \ll \mu_\infty$, define

$$\mathrm{KL}(\mu \| \mu_\infty) = \int_X \log\left(\frac{d\mu}{d\mu_\infty}\right) d\mu.$$

Let W_2 denote the 2-Wasserstein distance on $\mathcal{P}(X)$ with respect to d.

Remark 6.1 (Support convention for KL)**.** *When μ_∞ is supported only on a feasible set, the KL layer is interpreted either on the already-constructed support, or on a message/sampler representation whose state space has been restricted by the structural certificate. It is not a claim that an arbitrary full-cube initial law has finite KL relative to* $\mathrm{Unif}(\mathcal{A}(I))$*. If a full-cube initialization is desired, one must either use a strictly positive regularized reference law or first apply a certified restriction/projection procedure.*

Assumption 6.2 (T_2 or log-Sobolev input). *The reference law μ_∞ satisfies a transport–entropy inequality*

$$W_2^2(\mu, \mu_\infty) \le C_T \mathrm{KL}(\mu \| \mu_\infty)$$

for all $\mu \ll \mu_\infty$, where $C_T > 0$ is a declared constant. Alternatively, μ_∞ satisfies a log-Sobolev inequality with constant $\lambda_{\mathrm{LS}} > 0$, in which case one may take $C_T \asymp \lambda_{\mathrm{LS}}^{-1}$ up to convention-dependent constants.

Lemma 6.3 (Wasserstein control by KL). *Under the preceding assumption,*

$$W_2^2(\mu, \mu_\infty) \le C_T \mathrm{KL}(\mu \| \mu_\infty).$$

Remark 6.4 (No Fourier-displacement identification). *The inequality above does not require identifying optimal transport displacement with Fourier coefficient differences. Spectral quantities enter only as additional correctors or observables.*

6.2 Entropy decay along the ideal flow

Assume the ideal flow $(\nu_t)_{t \ge 0}$ satisfies

$$\frac{d}{dt} \mathrm{KL}(\nu_t \| \mu_\infty) \le -2\lambda \mathrm{KL}(\nu_t \| \mu_\infty)$$

for some $\lambda > 0$.

Proposition 6.5 (Ideal entropy and metric contraction). *If the displayed entropy dissipation holds, then*

$$\mathrm{KL}(\nu_t \| \mu_\infty) \le e^{-2\lambda t} \mathrm{KL}(\nu_0 \| \mu_\infty),$$

and, under the transport–entropy inequality,

$$W_2(\nu_t, \mu_\infty) \le \sqrt{C_T} e^{-\lambda t} \mathrm{KL}(\nu_0 \| \mu_\infty)^{1/2}.$$

Proof. The entropy estimate follows from Gronwall's inequality. The Wasserstein estimate follows by applying the transport–entropy inequality at time t. □

6.3 Spectral and covariance correctors

Let $\Phi : X \to \mathcal{H}$ be a feature map into a Hilbert space, and define the covariance operator

$$C(\mu) = \int_X \Phi(x) \otimes \Phi(x) \, d\mu(x).$$

For two probability laws μ, ν, define the pairwise covariance deviation

$$\Delta_C(\mu, \nu) = \|C(\mu) - C(\nu)\|_{\mathrm{HS}}.$$

We write $\Delta_C(\mu) = \Delta_C(\mu, \mu_\infty)$ when the reference law is clear.

Definition 6.6 (Rigidity distance). *For $\varepsilon_C \in (0, 1]$, define the pairwise rigidity distance*

$$d_{\mathrm{rigid}}(\mu, \nu)^2 = W_2^2(\mu, \nu) + \varepsilon_C \Delta_C(\mu, \nu)^2.$$

For a rigidity basin $\mathcal{M}_{\mathrm{rigid}}$, set

$$d_{\mathrm{rigid}}(\mu, \mathcal{M}_{\mathrm{rigid}}) = \inf_{\nu \in \mathcal{M}_{\mathrm{rigid}}} d_{\mathrm{rigid}}(\mu, \nu).$$

6.4 Perturbed realized flow

Let $(\mu_t)_{t\geq 0}$ be the realized flow, numerical flow, empirical law, approximate sampler, or learned message evolution. Define an error envelope

$$\eta_{[0,t]} = \sup_{0\leq s\leq t} \eta_s,$$

where

$$\eta_s^2 = \eta_{\mathrm{disc}}(s)^2 + \eta_{\mathrm{model}}(s)^2 + \eta_{\mathrm{stat}}(s)^2 + \eta_{\mathrm{BV}}(s)^2.$$

Theorem 6.7 (Perturbed Wasserstein–KL control)**.** *Assume:*

(i) μ_∞ *satisfies the transport–entropy inequality;*

(ii) the ideal flow ν_t *satisfies entropy decay with rate* $\lambda > 0$*;*

(iii) the ideal covariance corrector satisfies

$$\Delta_C(\nu_t, \mu_\infty) \leq A_C e^{-\lambda_C t} \Delta_C(\nu_0, \mu_\infty)$$

for some $A_C \geq 1$ *and* $\lambda_C > 0$*;*

(iv) the realized flow μ_t *satisfies the pairwise perturbation estimate*

$$W_2(\mu_t, \nu_t) + \Delta_C(\mu_t, \nu_t) \leq K_\eta \eta_{[0,t]}.$$

Let $\gamma_0 = \min\{\lambda, \lambda_C\}$*. Then there exist constants* A *and* B*, depending only on* C_T, A_C, K_η*, and* ε_C*, such that*

$$d_{\mathrm{rigid}}(\mu_t, \mu_\infty) \leq Ae^{-\gamma_0 t}\Big(\mathrm{KL}(\nu_0\|\mu_\infty)^{1/2} + \Delta_C(\nu_0, \mu_\infty)\Big) + B\eta_{[0,t]}.$$

Proof. By ideal contraction,

$$W_2(\nu_t, \mu_\infty) \leq \sqrt{C_T} e^{-\lambda t} \mathrm{KL}(\nu_0\|\mu_\infty)^{1/2}.$$

The triangle inequality gives

$$W_2(\mu_t, \mu_\infty) \leq W_2(\mu_t, \nu_t) + W_2(\nu_t, \mu_\infty).$$

The first term is controlled by the perturbation estimate. For the covariance part,

$$\Delta_C(\mu_t, \mu_\infty) \leq \Delta_C(\mu_t, \nu_t) + \Delta_C(\nu_t, \mu_\infty),$$

and the two terms are controlled by the perturbation estimate and the ideal covariance decay assumption. Combining the transport and covariance components in the pairwise definition of d_{rigid} gives the result. □

Assumption 6.8 (Coercive perturbed Lyapunov inequality)**.** *Let* $V(t) = d_{\mathrm{rigid}}(\mu_t, \mathcal{M}_{\mathrm{rigid}})^2$*. There exist constants* $\gamma > 0$ *and* $K > 0$ *such that*

$$\frac{d}{dt} V(t) \leq -2\gamma V(t) + K\eta_t^2.$$

Theorem 6.9 (Nonasymptotic rigidity radius)**.** *Under the preceding assumption,*

$$V(t) \leq e^{-2\gamma t} V(0) + K \int_0^t e^{-2\gamma(t-s)} \eta_s^2 \, ds.$$

Consequently,

$$d_{\mathrm{rigid}}(\mu_t, \mathcal{M}_{\mathrm{rigid}}) \leq e^{-\gamma t} d_{\mathrm{rigid}}(\mu_0, \mathcal{M}_{\mathrm{rigid}}) + \sqrt{\frac{K}{2\gamma}} \eta_{[0,t]}.$$

Proof. Multiplying the differential inequality by $e^{2\gamma t}$ and integrating from 0 to t gives the first estimate. The second follows from $\eta_s \leq \eta_{[0,t]}$ and taking square roots. □

7 Flow-and-Extract Conditional Theorem

We now combine the four gates into the main conditional algorithmic consequence of the framework.

Definition 7.1 (Valid extractor radius). *Let $\varepsilon_{\mathrm{round}} > 0$ and $\alpha \in [0, 1)$. An extractor Ext_I is $(\varepsilon_{\mathrm{round}}, \alpha)$-valid for I if*

$$d_{\mathrm{rigid}}(\mu, \mathcal{M}_{\mathrm{rigid}}(I)) \le \varepsilon_{\mathrm{round}}$$

implies

$$\Pr[V_I(\mathsf{Ext}_I(\mu)) = 1] \ge 1 - \alpha.$$

Remark 7.2. *Entropy contraction alone does not imply extractor validity. A flow may approach a smooth or low-entropy law that loses witness information. Extractor validity must be proved or certified separately for each instance family.*

The following theorem applies to the entropy–metric branch of Gate 3. In finite-closure certificates, the analogous budget is the polynomial runtime of the closure computation, and extraction is triggered by the certified zero-residual or obstruction-free state.

Theorem 7.3 (Flow-and-extract conditional theorem). *Let $\mathcal{F} = \{I_n\}$ be a declared structured instance family. Suppose that for every $I_n \in \mathcal{F}$, the following four gates are certified:*

(i) a structural certificate is constructible in polynomial time;

(ii) the chosen observable, reference law, or rigidity basin is constructible without hidden global solution information;

(iii) the realized flow satisfies

$$d_{\mathrm{rigid}}(\mu_t, \mathcal{M}_{\mathrm{rigid}}(I_n)) \le Ae^{-\gamma t} d_{\mathrm{rigid}}(\mu_0, \mathcal{M}_{\mathrm{rigid}}(I_n)) + B\eta;$$

(iv) the extractor Ext_{I_n} is $(\varepsilon_{\mathrm{round}}, \alpha)$-valid.

Assume also that $\varepsilon_{\mathrm{round}} > B\eta$. Then it suffices to run the flow until

$$T \ge \frac{1}{\gamma} \log \left(\frac{A d_{\mathrm{rigid}}(\mu_0, \mathcal{M}_{\mathrm{rigid}}(I_n))}{\varepsilon_{\mathrm{round}} - B\eta} \right).$$

At this time,

$$\Pr[V_{I_n}(\mathsf{Ext}_{I_n}(\mu_T)) = 1] \ge 1 - \alpha.$$

If certificate construction, flow simulation, and extraction costs are polynomial in n, and if the displayed parameters are polynomially bounded, the total flow-and-extract procedure runs in polynomial time.

Proof. By the entropy–metric gate,

$$d_{\mathrm{rigid}}(\mu_T, \mathcal{M}_{\mathrm{rigid}}(I_n)) \le Ae^{-\gamma T} d_{\mathrm{rigid}}(\mu_0, \mathcal{M}_{\mathrm{rigid}}(I_n)) + B\eta.$$

The chosen lower bound on T makes the first term at most $\varepsilon_{\mathrm{round}} - B\eta$. Hence

$$d_{\mathrm{rigid}}(\mu_T, \mathcal{M}_{\mathrm{rigid}}(I_n)) \le \varepsilon_{\mathrm{round}}.$$

Extractor validity gives the verifier success bound. The runtime statement follows by adding certificate construction, flow simulation, extraction, and verification costs. ☐

7.1 Gray-zone policy

If $\varepsilon_{\text{round}} \le B\eta$, the certificate is not strong enough to justify extraction. A conservative gray-zone policy is to run the flow up to a declared maximum budget, extract only if the certified radius enters the extraction ball, and otherwise output Inconclusive together with a diagnostic transcript.

8 Diagnostic PSC and Reproducibility Protocol

The predictive structural classifier, or PSC, is not a proof oracle. It estimates whether an instance is likely to satisfy the structural, spectral, entropy–metric, and extractor gates, and decides whether the instance should be accepted, rejected, or assigned to a gray zone.

Let

$$S(I) = (\text{tw}(I), \text{rank}(I), \text{overlap}(I), \text{sym}(I), \text{dualwt}(I), \dots)$$

be a structural feature vector, and let $\mathsf{Diag}(I)$ include quantities such as $\text{Tail}_{>k}(f_I)$, $\text{KL}(\mu_t \| \mu_\infty)$, $W_2(\mu_t, \mu_\infty)$, covariance deviation, and error-envelope estimates.

A PSC is a map

$$\text{PSC} : I \longmapsto \{\text{Rigid}, \text{Transitional}, \text{Nonrigid}, \text{Abstain}\}.$$

It should output Rigid only when structural features pass with margin, spectral or covariance tails are small, contraction is observed or certified, and the extractor gate has positive margin. If evidence is incomplete, it should output Abstain.

8.1 Avoiding data leakage

A pre-flow PSC may use only features computable before running the flow, such as structural graph features, local constraint statistics, and low-cost spectral estimates. A post-flow PSC may also use observed KL decay, covariance decay, and extractor outcomes. If the label is defined using KL decay or extraction success, and the classifier input includes the same quantities, then the classifier is summarizing the outcome rather than predicting it.

8.2 Reproducibility protocol

Any empirical PSC claim should specify instance generators, seeds, size ranges, exact or approximate Fourier computation, flow implementation, extractor implementation, verifier success rate, runtime, and ablation studies. Recommended ablations include no-flow baseline, structure-only PSC, spectrum-only PSC, flow-only PSC, feature-shuffle control, null-instance control, extractor ablation, and error-envelope ablation.

9 Algebraic Calibration: Affine Feasible Sets and Low-Weight Dual Rigidity

We now record a spectral calibration theorem for affine feasible sets over $\mathbb{F}_2$. Algebraic redundancy may imply spectral structure, but low rank alone is not sufficient to guarantee low-degree Fourier concentration. The correct quantity is the Hamming-weight profile of the dual space.

9.1 Affine feasible sets over $\mathbb{F}_2$

Let

$$V = a + L \subseteq \mathbb{F}_2^n$$

be an affine subspace, where $L \leq \mathbb{F}_2^n$ is a linear subspace of dimension d. Let $f = 1_V$. We use Walsh characters

$$\chi_s(x) = (-1)^{s \cdot x}, \qquad s, x \in \mathbb{F}_2^n,$$

and normalized Fourier coefficients

$$\widehat{f}(s) = 2^{-n} \sum_{x \in \mathbb{F}_2^n} f(x)\chi_s(x).$$

The dual subspace is

$$L^\perp = \{s \in \mathbb{F}_2^n : s \cdot \ell = 0 \text{ for all } \ell \in L\}.$$

Lemma 9.1 (Fourier support of affine indicators). *For $f = 1_{a+L}$,*

$$\widehat{f}(s) = \begin{cases} 2^{d-n}(-1)^{s \cdot a}, & s \in L^\perp, \\ 0, & s \notin L^\perp. \end{cases}$$

Proof. Writing $x = a + \ell$,

$$\widehat{f}(s) = 2^{-n}(-1)^{s \cdot a} \sum_{\ell \in L} (-1)^{s \cdot \ell}.$$

If $s \in L^\perp$, the sum equals $|L| = 2^d$. Otherwise the character is nontrivial on L and the sum vanishes. □

Define the dual weight enumerator

$$A_j(L^\perp) = |\{s \in L^\perp : |s| = j\}|, \qquad A_{>k}(L^\perp) = \sum_{j>k} A_j(L^\perp).$$

It is often convenient to package these numbers into the weight-enumerator polynomial

$$W_{L^\perp}(z) = \sum_{j=0}^{n} A_j(L^\perp) z^j.$$

Proposition 9.2 (Exact Fourier-tail formula). *Let $V = a + L \subseteq \mathbb{F}_2^n$, with $\dim L = d$. Then*

$$\mathrm{Tail}_{>k}(1_V) = 2^{2d-2n} A_{>k}(L^\perp).$$

Moreover,

$$\|1_V\|_2^2 = 2^{d-n},$$

and the normalized tail ratio is

$$\tau_k(1_V) = \frac{A_{>k}(L^\perp)}{|L^\perp|}.$$

Proof. By the preceding lemma, all nonzero Fourier coefficients are supported on $L^\perp$, and each has magnitude 2^{d-n}. Summing the squared coefficients over weights $> k$ gives the tail formula. Summing over all of $L^\perp$ gives $\|1_V\|_2^2 = |L^\perp| 2^{2d-2n} = 2^{d-n}$. Dividing gives the normalized ratio. □

Definition 9.3 (Low-weight dual rigidity). *An affine feasible set $V = a + L$ satisfies a (k, ε)-low-weight dual condition if*

$$\frac{A_{>k}(L^\perp)}{|L^\perp|} \leq \varepsilon.$$

It satisfies exact k-dual boundedness if $A_{>k}(L^\perp) = 0$.

Theorem 9.4 (Affine low-weight dual spectral certificate). *Let $V = a + L \subseteq \mathbb{F}_2^n$, and let $f = 1_V$. If V satisfies the (k, ε)-low-weight dual condition, then*

$$\mathrm{Tail}_{>k}(f) \leq \varepsilon \|f\|_2^2.$$

If V satisfies exact k-dual boundedness, then $\mathrm{Tail}_{>k}(f) = 0$.

Proof. This follows immediately from the exact normalized tail formula. □

9.2 Exact heat profile from the dual weight enumerator

The preceding theorem identifies the Fourier tail. The same computation gives a sharper dynamical statement: under the Boolean spectral heat semigroup, the entire decay curve is determined by the dual weight enumerator.

For an observable $f : \mathbb{F}_2^n \to \mathbb{R}$, define the Boolean spectral heat smoothing

$$P_t f = \sum_{s \in \mathbb{F}_2^n} e^{-t|s|} \widehat{f}(s) \chi_s.$$

For $V = a + L$, set

$$\delta_V(t)^2 = \|P_t 1_V - \mathbb{E}[1_V]\|_2^2.$$

Here $\mathbb{E}[1_V] = \widehat{1_V}(0) = |V|/2^n$, and the L^2-norm uses normalized counting measure on $\mathbb{F}_2^n$.

Theorem 9.5 (Affine heat profile from the dual weight enumerator)**.** *Let $V = a + L \subseteq \mathbb{F}_2^n$, with $\dim L = d$. Then*

$$\delta_V(t)^2 = 2^{2d-2n} \sum_{s \in L^\perp \setminus \{0\}} e^{-2t|s|} = 2^{2d-2n} \sum_{j=1}^{n} A_j(L^\perp) e^{-2tj}.$$

Equivalently,

$$\frac{\delta_V(t)^2}{\|1_V\|_2^2} = \frac{1}{|L^\perp|} \sum_{j=1}^{n} A_j(L^\perp) e^{-2tj}.$$

Equivalently, in terms of the weight-enumerator polynomial,

$$\delta_V(t)^2 = 2^{2d-2n} \big(W_{L^\perp}(e^{-2t}) - 1\big).$$

Thus the centered heat-decay profile is exactly the discrete heat transform of the nonzero dual weight enumerator, normalized by $|L^\perp|$.

Proof. By the Fourier support lemma, the nonzero Fourier coefficients of 1_V are supported on $L^\perp$, and each has squared magnitude 2^{2d-2n}. The constant coefficient corresponds to $s = 0$. Therefore Parseval's identity gives

$$\|P_t 1_V - \mathbb{E}[1_V]\|_2^2 = \sum_{s \neq 0} e^{-2t|s|} \widehat{1_V}(s)^2 = 2^{2d-2n} \sum_{s \in L^\perp \setminus \{0\}} e^{-2t|s|}.$$

Grouping dual words by Hamming weight gives the second expression. Since $A_0(L^\perp) = 1$, this is the same as $2^{2d-2n}(W_{L^\perp}(e^{-2t}) - 1)$. Dividing by $\|1_V\|_2^2 = 2^{d-n} = |L^\perp| 2^{2d-2n}$ gives the normalized profile. $\square$

Corollary 9.6 (Minimum dual weight controls heat collapse)**.** *Let*

$$w_{\min}(L^\perp) = \min\{|s| : s \in L^\perp \setminus \{0\}\}.$$

If $L^\perp \neq \{0\}$, then

$$\frac{\delta_V(t)^2}{\|1_V\|_2^2} \le e^{-2t w_{\min}(L^\perp)}.$$

More generally, the multiplicity of low-weight dual words determines the pre-asymptotic heat-decay profile.

Proof. Every nonzero dual word has weight at least $w_{\min}(L^\perp)$. Apply the theorem and use

$$\sum_{j=1}^{n} A_j(L^\perp) = |L^\perp| - 1 \le |L^\perp|.$$

$\square$

Corollary 9.7 (Low-weight dual words slow heat collapse)**.** *For $k \geq 1$, let*

$$A_{1 \leq j \leq k}(L^{\perp}) = \sum_{j=1}^{k} A_j(L^{\perp}).$$

Then

$$\frac{\delta_V(t)^2}{\|1_V\|_2^2} \geq \frac{A_{1 \leq j \leq k}(L^{\perp})}{|L^{\perp}|} e^{-2tk}.$$

Consequently, if a positive fraction of nonzero dual words has Hamming weight at most k, the centered heat profile cannot decay faster than this low-weight contribution.

Proof. For all nonzero dual words with $|s| \leq k$, one has $e^{-2t|s|} \geq e^{-2tk}$. Keeping only those terms in the normalized heat-profile formula proves the bound. □

Remark 9.8 (Rank versus heat profile)**.** *The dimension* $\dim L^{\perp}$*, equivalently the rank of a defining linear system, fixes only the number of dual words. It does not determine their Hamming weights. The heat profile is therefore controlled by the full weight enumerator, not by rank alone.*

9.3 Non-circular extraction for affine systems

Consider an affine constraint system $Ax = b$ over $\mathbb{F}_2$. Its feasible set is either empty or an affine subspace $V = a + \ker A$. Gaussian elimination determines feasibility, constructs a particular solution a, and computes a basis for $\ker A$ in polynomial time. If feasible, an exact uniform sampler for V is obtained by choosing a random coefficient vector and outputting a random affine combination of the kernel basis.

Gaussian elimination constructs $L = \ker A$ and a representation of $L^{\perp}$, but exact certification of the full dual weight profile may require additional structure or a supplied succinct certificate; it is not automatic from rank alone.

Theorem 9.9 (Affine spectral/reference-law certificate with dual profile)**.** *Let $\mathcal{F}_{\mathrm{aff}}$ be a family of affine systems $Ax = b$ over $\mathbb{F}_2$. Suppose that, for each feasible instance, the associated affine space $V = a + \ker A$ has a supplied or efficiently checkable certificate of a (k, ε)-low-weight dual profile. Then each feasible instance satisfies the structural gate, the spectral/reference-law gate, and the extractor gate: the structural gate is the affine representation, the spectral gate is the preceding theorem plus the supplied dual-profile certificate, the reference law is $\mu_\infty = \mathrm{Unif}(V)$, and exact affine sampling gives verifier success probability 1. If finite-step Gaussian elimination is admitted as an algebraic projection flow, this also supplies a degenerate exact flow gate.*

Proof. Gaussian elimination determines feasibility. If feasible, it constructs a and a basis for $\ker A$, hence constructs μ_∞ and an exact sampler without hidden solution information. The spectral gate follows from the low-weight dual theorem together with the assumed dual-profile certificate. Gaussian elimination alone supplies the affine space and its dual representation, not necessarily an efficient exact enumeration of the full dual weight profile. The sampled point satisfies $Ax = b$, so the verifier accepts with probability 1. The optional flow-gate interpretation is finite-step algebraic projection rather than an entropy-mixing claim. □

Remark 9.10 (Low rank is not enough)**.** *The rank of A determines* $\dim L^{\perp}$*, but it does not determine the Hamming weights of vectors in $L^{\perp}$. Thus low rank alone is not a low-degree Fourier certificate. The dual weight distribution is the relevant spectral quantity.*

10 Implication-Graph Rigidity: 2-SAT as a Non-Treewidth Calibration Family

The previous calibration families are based on product structure, tree decompositions, and linear algebra. We now add a fourth calibration mechanism: implication closure. The class of 2-SAT formulas is tractable even when the primal graph has large treewidth, so it provides a useful non-treewidth test of the certificate architecture. The point is not to introduce a new 2-SAT algorithm. Rather, the point is to show that the four-gate language can also represent graph-obstruction and closure-based tractability.

In this section the word "flow" is used only in the generalized finite-closure sense. The 2-SAT closure gate is not a Wasserstein–KL entropy flow of the kind discussed in Section 6; it is an exact polynomial-time graph procedure whose residual vanishes after the SCC computation.

10.1 2-CNF formulas and implication graphs

A 2-CNF formula is a conjunction of clauses

$$F = \bigwedge_{\alpha=1}^{m} (\ell_{\alpha,1} \vee \ell_{\alpha,2}),$$

where each literal ℓ is either x_i or $\neg x_i$. The implication graph G_F has one vertex for each literal $x_i, \neg x_i$, and for every clause $(a \vee b)$ it contains the two directed edges

$$\neg a \to b, \qquad \neg b \to a.$$

The standard satisfiability criterion is:

$$F \text{ is satisfiable} \quad \Longleftrightarrow \quad x_i \text{ and } \neg x_i \text{ lie in distinct strongly connected components for every } i.$$

If the criterion fails, the formula is infeasible. If it passes, a satisfying assignment is obtained by ordering the strongly connected components of the condensation DAG and assigning each variable according to the relative order of the components containing x_i and $\neg x_i$.

10.2 The spectral gate for the additive violation energy

For a 2-CNF formula F, define the additive violation energy

$$E_F(x) = \sum_{\alpha=1}^{m} 1\{(\ell_{\alpha,1} \vee \ell_{\alpha,2}) \text{ is violated by } x\}.$$

Each summand depends on at most two variables. Therefore, by Proposition 5.1,

$$\mathrm{Tail}_{>2}(E_F) = 0.$$

This is the safe spectral gate for 2-SAT. It should not be confused with a claim about the global satisfying-assignment indicator $1_{\mathcal{A}(F)}$, which may have high-degree Fourier terms.

Proposition 10.1 (Degree-two spectral gate for 2-SAT energies)**.** *Let F be a 2-CNF formula, and let E_F be its additive violation energy. With the normalized Walsh–Fourier convention,*

$$\widehat{E_F}(S) = 0 \qquad \text{for all } |S| > 2.$$

Equivalently, E_F is exactly $(2,0)$-spectrally concentrated.

Proof. A violated 2-clause indicator depends only on the two variables appearing in the clause, or on one variable if the clause is degenerate. Its Walsh–Fourier expansion is therefore supported in degree at most two. Summing over clauses cannot create higher-degree Fourier coefficients. Hence all coefficients of degree greater than two vanish. □

10.3 Finite implication-closure gate

The implication graph supplies a finite closure gate. Let $\mathrm{SCC}(G_F)$ denote the strongly connected component partition of G_F, and define the contradiction residual

$$R_{\mathrm{imp}}(F) = \#\{i : x_i \text{ and } \neg x_i \text{ belong to the same SCC of } G_F\}.$$

Then $R_{\mathrm{imp}}(F) = 0$ if and only if the implication obstruction is absent. Computing strongly connected components is a finite graph-closure operation, not an entropy flow. In the rigidity-certificate language, it plays the role of a finite exact collapse map: after the SCC computation, the obstruction residual is known exactly.

10.4 2-SAT rigidity certificate

Theorem 10.2 (Implication-graph rigidity certificate for 2-SAT)**.** *Let $\mathcal{F}_{2\mathrm{SAT}}$ be the family of 2-CNF formulas. Then $\mathcal{F}_{2\mathrm{SAT}}$ is decision-certified effective under the rigidity framework. More precisely, for each formula F:*

(i) the structural gate is the 2-CNF representation and its implication graph G_F;

(ii) the spectral gate holds for the additive violation energy, with $\mathrm{Tail}_{>2}(E_F) = 0$;

(iii) the finite closure gate is the SCC computation, which determines $R_{\mathrm{imp}}(F)$ exactly;

(iv) if $R_{\mathrm{imp}}(F) > 0$, infeasibility is certified;

(v) if $R_{\mathrm{imp}}(F) = 0$, a fixed relative-order rule on the SCC condensation DAG constructs a satisfying assignment: with a topological order ord satisfying $C \to D \Rightarrow \mathrm{ord}(C) < \mathrm{ord}(D)$, set $x_i = \mathrm{true}$ iff $\mathrm{ord}(\mathrm{SCC}(x_i)) > \mathrm{ord}(\mathrm{SCC}(\neg x_i))$;

(vi) the final verifier accepts the constructed assignment with probability 1.

The certificate and extractor run in time linear in the size of the implication graph, hence polynomial in the formula size.

Proof. Given a 2-CNF formula, construct the implication graph by adding two directed edges for each clause. This gives a graph with $2n$ vertices and $2m$ edges. Strongly connected components are computable in linear time in the graph size.

If some variable x_i and its negation $\neg x_i$ belong to the same strongly connected component, then the implication graph contains paths $x_i \Rightarrow \neg x_i$ and $\neg x_i \Rightarrow x_i$. In any satisfying assignment, implications preserve truth. Since exactly one of x_i and $\neg x_i$ is true, whichever literal is true would force the other literal to be true as well, a contradiction. Hence the formula is unsatisfiable.

Conversely, suppose no variable is identified with its negation in the SCC partition. Collapse the SCCs to the condensation DAG and fix a topological order ord with the convention

$$C \to D \quad \Longrightarrow \quad \mathrm{ord}(C) < \mathrm{ord}(D).$$

Set

$$x_i = \mathrm{true} \quad \Longleftrightarrow \quad \mathrm{ord}(\mathrm{SCC}(x_i)) > \mathrm{ord}(\mathrm{SCC}(\neg x_i)).$$

This relative-order rule assigns opposite truth values to paired literal components. It also satisfies every implication. Indeed, if an implication $a \to b$ had a true and b false, then the implication edges $a \to b$ and $\neg b \to \neg a$ would give

$$\mathrm{ord}(\neg a) < \mathrm{ord}(a) < \mathrm{ord}(b) < \mathrm{ord}(\neg b) < \mathrm{ord}(\neg a),$$

a contradiction. Therefore every implication, and hence every original clause, is satisfied. The original verifier accepts the assignment.

The spectral statement follows from the preceding proposition. Thus the structural, spectral, finite-closure, extractor, and verifier components are all supplied non-circularly and in polynomial time. $\square$

10.5 Relation to the four-gate framework

The 2-SAT certificate differs from the bounded-treewidth certificate in an important way. The reference object is not a junction-tree feasible law, and the decisive mechanism is not decomposition width. Instead, the decisive mechanism is implication closure and obstruction detection. Thus the four-gate architecture is satisfied here after replacing the entropy–metric gate by an exact finite closure gate. This broadens the calibration scope of the framework:

$$\text{decomposition} \quad \text{vs.} \quad \text{linear algebra} \quad \text{vs.} \quad \text{implication closure.}$$

The entropy–metric language is therefore not required to be literal in every tractable class. In finite exact classes such as 2-SAT, the flow gate may be realized by a finite closure map whose residual vanishes exactly after a polynomial-time computation. This reinforces the main methodological point: effective rigidity is a certificate architecture, not a single universal dynamical system.

Remark 10.3 (No new 2-SAT algorithm is claimed)**.** *Theorem 10.2 is a recasting theorem. It does not improve the classical 2-SAT algorithm. Its role is to show that the rigidity framework covers a non-treewidth tractability mechanism and separates the safe degree-two spectral statement for additive energies from the much stronger and generally false statement that all 2-SAT feasibility indicators are low-degree.*

11 Low-Overlap CNF Obstruction: Disjoint Clauses Need Not Give Low-Degree Feasibility Indicators

This section turns the warning of Section 5 into an explicit obstruction theorem. The structure-to-spectrum gate is useful only if the chosen structural feature truly controls the chosen spectral object. We now prove a simple obstruction: low clause overlap alone does not imply low-degree Fourier concentration of the global feasibility indicator. The example is deliberately elementary. It shows why the framework distinguishes additive violation energies from feasibility indicators.

The family below is algorithmically trivial and even block-factorized. The obstruction is therefore not an obstruction to tractability. It is specifically an obstruction to the inference that low overlap, or even product structure, forces low-degree concentration of the global feasibility indicator.

11.1 A disjoint monotone 2-CNF family

For $m \geq 1$, define a monotone 2-CNF formula on $2m$ variables by

$$F_m = \bigwedge_{j=1}^{m} (x_{2j-1} \vee x_{2j}).$$

The clauses are pairwise variable-disjoint. Thus any clause-overlap measure based on shared variables between distinct clauses is zero. Let

$$f_m = 1_{\mathcal{A}(F_m)} = \prod_{j=1}^{m} C_j,$$

where $C_j(x_{2j-1}, x_{2j}) = 1_{x_{2j-1} \vee x_{2j}}$.

11.2 One-clause Fourier distribution

Let $C(x, y) = 1_{x \vee y}$. With the normalized Walsh convention on two variables,

$$\widehat{C}(\emptyset) = \frac{3}{4}, \qquad \widehat{C}(\{x\}) = -\frac{1}{4}, \qquad \widehat{C}(\{y\}) = -\frac{1}{4}, \qquad \widehat{C}(\{x, y\}) = -\frac{1}{4}.$$

Hence

$$\|C\|_2^2 = \frac{3}{4}.$$

The normalized Fourier energy by degree is

$$\Pr(D = 0) = \frac{3}{4}, \qquad \Pr(D = 1) = \frac{1}{6}, \qquad \Pr(D = 2) = \frac{1}{12}.$$

Here D is the random degree obtained by selecting a Fourier coefficient with probability proportional to its squared magnitude.

11.3 Product convolution and high-degree escape

Because the clauses use disjoint variable pairs, the Fourier transform of f_m is the tensor product of the one-clause Fourier transforms. Therefore the normalized Fourier degree distribution of f_m is the distribution of

$$S_m = D_1 + \cdots + D_m,$$

where the D_j are independent copies of D. In particular,

$$\mathbb{E}D = \frac{1}{3}.$$

For any fixed k, the probability $\Pr(S_m \le k)$ tends to zero as $m \to \infty$. Consequently, the normalized Fourier mass above degree k tends to one.

Theorem 11.1 (Low-overlap does not imply low-degree concentration of feasibility indicators)**.** *There exists a family of monotone 2-CNF formulas F_m with pairwise disjoint clauses, hence zero clause overlap, such that for every fixed k,*

$$\frac{\mathrm{Tail}_{>k}(1_{\mathcal{A}(F_m)})}{\|1_{\mathcal{A}(F_m)}\|_2^2} \longrightarrow 1 \qquad (m \to \infty).$$

More precisely, the normalized Fourier degree distribution of $1_{\mathcal{A}(F_m)}$ is the distribution of $S_m = \sum_{j=1}^{m} D_j$, where

$$\Pr(D_j = 0) = \frac{3}{4}, \qquad \Pr(D_j = 1) = \frac{1}{6}, \qquad \Pr(D_j = 2) = \frac{1}{12}.$$

Proof. The one-clause Fourier computation gives the stated degree distribution for a single factor. Since the clauses are disjoint, products of coefficients multiply and degrees add. Therefore the normalized degree distribution for f_m is the m-fold convolution of the one-clause distribution, i.e. the law of S_m.

Let $B_m \sim \mathrm{Binomial}(m, 1/4)$. Since $\Pr(D_j > 0) = 1/4$, the random variable S_m stochastically dominates B_m in the sense that $S_m \ge \sum_j 1_{\{D_j > 0\}}$. Hence for fixed k,

$$\Pr(S_m \le k) \le \Pr(B_m \le k) \to 0.$$

The normalized Fourier mass of degree at most k is exactly $\Pr(S_m \le k)$. Therefore the normalized mass above degree k tends to one. □

Corollary 11.2 (A structural warning for the gate)**.** *Low clause overlap, even zero overlap, is not by itself a certificate for low-degree Fourier concentration of the feasibility indicator. This remains true even for an algorithmically trivial product family. Any valid structure-to-spectrum gate for CNF feasibility indicators must include additional hypotheses or must analyze a different observable, such as the additive violation energy.*

12 Limitations, Non-Claims, and Open Gates

This manuscript is deliberately framed as a certificate framework for structured NP instance families. Because the language of rigidity, entropy contraction, and effective tractability can be misread as a claim about the global P versus NP problem, we record the main limitations and non-claims explicitly.

12.1 No claim about the global P versus NP problem

The framework does not prove $P = NP$, $P \neq NP$, or any separation or collapse between standard complexity classes. The term certified effective family refers only to a declared family of structured instances for which the four gates are verified.

12.2 Smoothing is not solving

The spectral heat diagnostic $\widehat{f_t}(S) = e^{-t|S|}\widehat{f}(S)$ suppresses high-degree Fourier components by definition. Therefore, the mere fact that a smoothed observable collapses toward a low-frequency or constant component does not imply that the original instance has been solved.

12.3 Fourier access may be expensive

Fourier coefficients of a Boolean feasibility indicator may be computationally expensive to compute exactly. Full enumeration of the truth table costs 2^n in general. Any spectral gate must therefore state its access model: exact enumeration for small n, low-degree computation, sampling access, oracle access, structural dynamic programming, or covariance/local-energy proxies.

12.4 Reference laws must be non-circular

The reference law μ_∞, rigidity basin $\mathcal{M}_{\mathrm{rigid}}$, and extractor must be constructible without secretly encoding a global solution. If μ_∞ is defined as the uniform law on all satisfying assignments but no polynomial construction, sampler, or marginal oracle is supplied, then the certificate is not algorithmically meaningful.

12.5 Metric and corrector constants depend on conventions

Transport constants depend on the metric normalization. A normalized product metric may give dimension-free constants for product examples, while an unnormalized Hamming metric usually introduces dimension dependence. Similarly, covariance correctors require either a pairwise covariance stability estimate, a feature-Lipschitz comparison, or a separate covariance decay assumption. These dependencies must be declared in any concrete certificate.

12.6 Low rank is not sufficient

Algebraic redundancy and low matrix rank can be useful structural indicators, but low rank alone does not imply low-degree Fourier concentration. For affine feasible sets, the Fourier support is governed by the dual subspace, and the relevant spectral quantity is the Hamming-weight profile of that dual space. Gaussian elimination constructs the affine space and a dual representation, but exact certification of $A_{>k}(L^\perp)$ or the full dual weight enumerator can be computationally expensive without additional structure or a supplied succinct certificate.

12.7 Open gates

The most important open gates are:

(O1) Low-overlap CNF beyond the obstruction theorem: determine which additional hypotheses, beyond low overlap alone, imply useful spectral concentration or message contraction.

(O2) Approximate structure-to-spectrum estimates beyond additive energies and low-weight dual affine systems.

(O3) Extractor stability for partially rigid families.

(O4) Computable error-envelope certification for discretization, model, statistical, and BV errors.

(O5) Runtime correlation between rigidity diagnostics and actual solver performance.

(O6) Families beyond known product, tree/hypertree, affine, and implication-closure tractable classes, together with structural hypotheses that avoid the low-overlap obstruction, that still admit spectral–entropic certificates.

13 Conclusion

This manuscript proposed a spectral–entropic rigidity certificate framework for structured NP instance families. The framework separates effective tractability into four gates: structural admissibility, spectral or reference-law control, entropy–metric contraction or finite closure, and extractor validity. This separation prevents the common mistake of treating smoothing or entropy collapse as solving.

Four calibration mechanisms were developed. Block-factorized bounded-arity CSPs provide a transparent product case. Bounded-treewidth CSPs provide a more substantial structured family, where junction-tree dynamic programming constructs the reference law, certifies infeasibility when appropriate, and samples exact witnesses when feasible. Affine feasible sets over $\mathbb{F}_2$ provide a spectral calibration theorem: low-degree Fourier concentration and the entire Boolean heat-decay profile are governed by the Hamming-weight enumerator of the dual space, not by rank alone. Finally, 2-SAT provides a non-treewidth implication-closure calibration: strongly connected components detect obstruction, a finite closure residual replaces the entropy-flow gate, and a condensation-order extractor yields a verified witness when the obstruction is absent.

The manuscript also adds a negative calibration theorem. Disjoint monotone 2-CNF clauses have zero clause overlap, but their feasibility indicators have normalized Fourier mass escaping every fixed low-degree cutoff. Thus low overlap alone is not a valid structure-to-spectrum certificate for feasibility indicators. This obstruction clarifies why the framework separates additive energies, reference-law representations, and global feasibility indicators.

The analytic core of the framework is a Wasserstein–KL contraction layer with perturbative spectral, covariance, statistical, discretization, and BV correctors. This layer supplies explicit rigidity radii and flow-and-extract time budgets, but only becomes algorithmically meaningful when paired with an extractor gate and final verifier.

The broader structure-to-spectrum principle remains the main mathematical challenge. It is exact for additive bounded-arity energies and can be quantified exactly for affine feasible sets through the dual weight enumerator. At the same time, the disjoint-clause obstruction shows that low overlap alone, even in a tractable product family, does not force low-degree concentration of global feasibility indicators. The predictive structural classifier is therefore positioned as a diagnostic tool rather than a proof engine.

The resulting framework does not address the global P versus NP problem. Its contribution is narrower and more constructive: it gives a disciplined language for certifying when structured NP instances behave effectively, and it identifies the precise gates that must be closed before entropy-rigidity diagnostics can become algorithmic guarantees.

References

[1] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker–Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.

[2] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. *Communications in Partial Differential Equations*, 26(1–2):101–174, 2001.

[3] F. Otto and C. Villani. Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.

[4] C. Villani. *Optimal Transport: Old and New.* Springer, 2009.

[5] B. Aspvall, M. F. Plass, and R. E. Tarjan. A linear-time algorithm for testing the truth of certain quantified Boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.

[6] R. Dechter. *Constraint Processing.* Morgan Kaufmann, 2003.

[7] S. L. Lauritzen. *Graphical Models.* Oxford University Press, 1996.

[8] R. O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, 2014.

[9] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* Wiley, second edition, 2006.

[10] M. R. Krom. The decision problem for a class of first-order formulas in which all disjunctions are binary. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 13:15–20, 1967.

[11] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.

[12] N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.

[13] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996.

[14] G. Gottlob, N. Leone, and F. Scarcello. Hypertree decompositions and tractable queries. *Journal of Computer and System Sciences*, 64(3):579–627, 2002.

[15] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.

[16] P. G. Kolaitis and M. Y. Vardi. A game-theoretic approach to constraint satisfaction. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 175–181, 2000.

[17] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques.* MIT Press, 2009.

[18] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.

[19] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *The Annals of Applied Probability*, 6(3):695–750, 1996.

[20] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times.* American Mathematical Society, 2009.

[21] M. Ledoux. *The Concentration of Measure Phenomenon.* American Mathematical Society, 2001.

[22] C. H. Papadimitriou. *Computational Complexity.* Addison-Wesley, 1994.

[23] M. Sipser. *Introduction to the Theory of Computation.* Cengage Learning, third edition, 2012.

[24] S. Arora and B. Barak. *Computational Complexity: A Modern Approach.* Cambridge University Press, 2009.

[25] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes.* North-Holland, 1977.