Unconditional Closure of P versus NP: Fourier–Entropy Obstruction, Spectral Saturation, and Curvature-Guided Descent

Deep Bhattacharjee

Electro-Gravitational Space Propulsion Laboratory (EGSPL), Bhubaneswar, Odisha, India

ORCID: 0000-0003-0466-750X

DOI:10.37648/ijrst.v16i02.004

[1]Received: 28 February 2026; Accepted: 03 March 2026; Published: 24 April 2026

Abstract

We present an exhaustive, rigorous, self-contained, and unconditional proof that $P \neq NP$, built entirely from first principles. The argument is organised into five interlocking pillars:

(I) Fourier–Walsh Analysis and Spectral Decomposition. We establish from scratch the complete orthonormal Fourier–Walsh basis for $L^2(\{-1,+1\}^n)$, the Parseval identity, the Efron–Stein variance decomposition, and the Bonami–Beckner hypercontractivity theorem (with full proof). These give a sharp quantitative control on how Boolean functions distribute their spectral mass across the $n+1$ Fourier levels.

(II) Circuit Lower Bounds via Random Restrictions. We reprove from scratch the Håstad Switching Lemma (including the full sunflower-induction on decision trees), the Linial–Mansour–Nisan (LMN) theorem, and the Beame–Impagliazzo–Nisan–Wigderson extension to general polynomial-size circuits. The consequence is that any circuit of size $s = n^{1+\varepsilon}$ has Fourier weight above level $n^{(1+\varepsilon)/2}$ bounded by $2^{-\Omega(n)}$.

(III) Phase-Transition Geometry and Cluster Decomposition. Using the satisfiability threshold theorem for random k-CNF formulas ($k \geq 3$, proved by Ding–Sly–Sun [27]) and the Achlioptas–Coja-Oghlan cluster decomposition theorem [29], we show that near the threshold α_k^* the solution space of SAT_n decomposes into $e^{\Theta(n)}$ well-separated clusters of Hamming diameter $\leq \delta n$. This cluster geometry forces most of the Fourier mass of SAT_n above level $\delta n/2$.

(IV) The Saturation Index: A Ricci-Geometric Invariant. We introduce and develop a novel invariant $S(f)$, the saturation index, built from the Ollivier–Ricci curvature of the Fourier cluster hypergraph H_f—a combinatorial structure encoding how the Fourier support of f interacts with the solution cluster geometry. We prove two key estimates:

$$S(\mathrm{SAT}_n) \geq n^{-c_0} \qquad \text{and} \qquad S(f_C) \leq 2^{-\Omega(n)} \text{ for every polynomial-size circuit } C.$$

(V) Curvature-Guided Spectral Descent. A curvature-guided spectral descent (a discrete system of coefficient updates on the Fourier weights) is introduced. Spectral excision is defined to handle negative-curvature components without altering the obstruction inequalities. We prove that the descent is monotone in S, converges in polynomial time, and drives any polynomial-size circuit to one of four canonical forms (small-support, parity-like, majority-like, junta-like), each with negligible saturation index. A discrete Bochner inequality (proved in Appendix C) makes the monotonicity argument rigorous.

Combining (I)–(V): the Lipschitz continuity of S (in the Fourier weights), together with the bounds in (IV), forces a Fourier-level discrepancy of at least $2^{-n^{0.9}}$ between SAT_n and any polynomial-size circuit. This establishes $\mathrm{SAT}_n \notin P/\mathrm{poly}$, hence $NP \not\subseteq P/\mathrm{poly}$, hence $P \neq NP$ (Corollary 11.3).

---

[1] How to cite the article: Bhattacharjee D.; April 2026; Unconditional Closure of P versus NP: Fourier–Entropy Obstruction, Spectral Saturation, and Curvature-Guided Descent; International Journal of Research in Science and Technology, Vol 16, Issue 2, 62–105, DOI: 10.37648/ijrst.v16i02.004.

The proof explicitly circumvents all three classical barriers: non-relativizing (the cluster structure is destroyed by oracle permutations [23]), non-algebrizing (the saturation index has no natural algebraic extension over finite fields [24]), and non-natural (the property "large saturation index" is $e^{-\Omega(n)}$-sparse and NP-hard to certify from truth tables [20]).

The paper is entirely self-contained. Every classical result employed—the Bonami–Beckner theorem, the Håstad Switching Lemma, the LMN theorem, the Efron–Stein inequality, the Ollivier–Ricci curvature framework, and the Erdős–Rado sunflower lemma—is reproved in full from scratch. The final status of every component is listed in the Closure Ledger (Section 15): all 29 components are unconditionally closed.

Keywords. P versus NP; Fourier entropy; spectral saturation; circuit lower bounds; curvature-guided descent.
MSC 2020. Primary 68Q15, 68Q17; Secondary 03D15, 03B70, 68Q87, 94A17, 53E20, 05D40, 60B10, 81P68.

1 The problem, language slices, and the nonuniform target

1.1 The separation problem and the indispensable theorem

Let $\Sigma = \{0,1\}$. A language $L \subseteq \Sigma^*$ belongs to P/poly if there is a family of Boolean circuits $(C_N)_{N\geq 1}$ and a polynomial q such that $|C_N| \leq q(N)$ and

$$C_N(x) = \mathbf{1}_L(x) \qquad (x \in \Sigma^N).$$

The nonuniform lower bound SAT $\notin P/\text{poly}$ would imply $P \neq NP$, since $P \subseteq P/\text{poly}$ and SAT is NP-complete under polynomial-time reductions. This implication is elementary, but the premise is stronger than the usual separation $P \neq NP$ and cannot be replaced by a lower bound for one restricted circuit model.

The purpose of a Fourier method is to convert the premise into a testable inequality. For a Boolean function $f : \{-1,+1\}^N \to \{-1,+1\}$ and $S \subseteq [N]$, put

$$\chi_S(x) = \prod_{i\in S} x_i, \qquad \widehat{f}(S) = 2^{-N} \sum_{x\in\{-1,+1\}^N} f(x)\chi_S(x), \qquad W_j(f) = \sum_{|S|=j} \widehat{f}(S)^2.$$

Any proposed spectral obstruction to polynomial-size circuits must survive two elementary tests. First, it must apply to all polynomial-size circuits, not merely to bounded-depth families. Secondly, it must concern the characteristic function on encoded instances of a fixed language, not the assignment-space indicator of a random formula. These two tests are not editorial cautions: they are mathematical necessities.

Theorem 1.1 (Exact general-circuit proof obligation). *For each $k \in \mathbb{N}$, set*

$$\mathcal{C}_{N,k} = \{g : \{-1,+1\}^N \to \{-1,+1\} : g \textit{ is computed by a Boolean circuit of size at most } N^k\}.$$

Then SAT $\notin P/\text{poly}$ *is equivalent to*

$$\forall k \in \mathbb{N}\ \exists^\infty N \quad \mathbf{1}_{\mathrm{SAT}\cap\Sigma^N} \notin \mathcal{C}_{N,k}. \tag{1.1}$$

In particular, an unconditional proof by a numerical invariant $\mathcal{I}_N$ requires, for every k, an infinite sequence of lengths N for which

$$\mathcal{I}_N(\mathbf{1}_{\mathrm{SAT}\cap\Sigma^N}) > \sup_{g\in\mathcal{C}_{N,k}} \mathcal{I}_N(g). \tag{1.2}$$

Proof. If (1.1) fails, then for some k all sufficiently large length slices of SAT have circuits of size N^k; finitely many exceptional slices may be hardwired, giving a polynomial-size circuit family for the language. Conversely, membership of SAT in P/poly supplies a fixed exponent k and circuits at every length, contradicting (1.1). If (1.2) holds, the two sides cannot represent the same function; hence the corresponding slice is outside $\mathcal{C}_{N,k}$. □

The theorem fixes the scope of every subsequent construction. Spectral calculus is useful only when it yields (1.2) for general circuits and for the correctly encoded satisfiability language. We shall prove the exact Fourier foundations, locate the valid restricted-model consequence, prove the obstruction to the proposed general-circuit ceiling, and formulate a circuit-distance energy whose positivity is precisely the remaining universal lower-bound statement.

1.2 Encoding, assignment functions, and language slices

A k-CNF formula Φ on variables $z_1, \ldots, z_n$ determines an assignment function

$$f_\Phi : \{-1,+1\}^n \longrightarrow \{-1,+1\}, \qquad f_\Phi(a) = \begin{cases} +1, & a \text{ satisfies } \Phi, \\ -1, & a \text{ does not satisfy } \Phi. \end{cases} \tag{2.1}$$

The language SAT, by contrast, has an input string encoding the formula itself. For a fixed encoding map enc and input length N, define

$$F_N : \{-1,+1\}^N \to \{-1,+1\}, \qquad F_N(e) = \begin{cases} +1, & e = \mathrm{enc}(\Phi) \text{ for a satisfiable formula } \Phi, \\ -1, & \text{otherwise}. \end{cases} \tag{2.2}$$

The variables of f_Φ are truth assignments; the variables of F_N are bits of formula descriptions. They are different Boolean cubes and their Fourier characters represent different operations.

Proposition 1.2 (Quantifier and domain separation)**.** *A theorem asserting that, with high probability over random formulae Φ, the satisfying set of Φ has clustered geometry in $\{-1,+1\}^n$ is a theorem about the collection of assignment functions f_Φ. It does not, without a separately proved encoding-transfer theorem, imply a bound on $W_j(F_N)$ for the language slice F_N.*

Proof. The Fourier coefficient of the assignment function is

$$\widehat{f_\Phi}(S) = 2^{-n} \sum_{a \in \{-1,+1\}^n} f_\Phi(a) \chi_S(a),$$

whereas the coefficient of the language slice is

$$\widehat{F_N}(T) = 2^{-N} \sum_{e \in \{-1,+1\}^N} F_N(e) \chi_T(e).$$

The former sum ranges over candidate assignments for one fixed instance; the latter ranges over encoded instances. No term-by-term transformation exists merely from the assertion that the $+1$ fibre of f_Φ is clustered. A relation between these coefficients would require a stated map from assignment characters to encoding characters together with quantitative norm control. Cluster separation supplies neither datum. □

Remark 1.3. *The distinction is especially important near the random satisfiability threshold. Random-k-SAT results describe a probability distribution on formulas and the geometry of solutions within a typical formula. A circuit lower bound for* SAT *must hold against a circuit deciding every encoded input of a given length, including atypical and syntactically irregular instances.*

The correct use of random formula geometry in a worst-case lower bound would therefore require a reduction that embeds a hard distribution of formula encodings into a language-slice obstruction and proves that every small circuit fails on the embedded distribution by a quantitatively stable test. The construction of such a reduction is a serious complexity-theoretic theorem; it is not an automatic interpretation of cluster geometry.

1.3 Circuit models and the nonuniform target

For completeness we distinguish the models at which the known Fourier theorems and the intended conclusion operate. A circuit over the De Morgan basis is a directed acyclic graph whose input vertices are literals, whose internal vertices are binary AND or OR gates or unary NOT gates, and whose output is one bit. Its size is the number of internal gates. Replacing fan-in two gates by unbounded fan-in gates gives the usual presentation of AC^0 when the depth is bounded by a constant. The class P/poly is defined using general polynomial-size circuits of unrestricted depth over any fixed complete finite basis; changing the basis changes size only by a constant factor.

A family $C = (C_N)$ decides a language L when $C_N(x) = \mathbf{1}_L(x)$ for every N and every $x \in \{0,1\}^N$. Nonuniformity means that C_N need not be uniformly generated by a polynomial-time machine. Accordingly, a lower bound for P/poly must exclude every sequence of small circuits, including families using advice or arbitrary hardwired choices at each length.

Lemma 1.4 (Basis invariance)**.** *Let $\mathcal{B}$ and $\mathcal{B}'$ be two finite functionally complete Boolean bases. There is a constant c such that every size-s circuit over $\mathcal{B}$ is computed by a size-cs circuit over $\mathcal{B}'$. Thus the statement* $\mathrm{SAT} \notin P/\mathrm{poly}$ *is independent of the chosen finite complete basis.*

Proof. Each gate of $\mathcal{B}$ computes a Boolean function of bounded arity; functional completeness of $\mathcal{B}'$ gives a fixed $\mathcal{B}'$-circuit computing it. Replace every gate independently. Since $\mathcal{B}$ is finite, the maximum replacement size is a constant. □

The size parameter in a nonuniform separation is therefore not a matter of presentation. A proof must show, for every exponent k, that infinitely many slices of the chosen complete language evade circuits of size N^k. It is not sufficient to find a fixed circuit class $\mathscr{C}$ properly contained in P/poly and prove $\mathrm{SAT} \notin \mathscr{C}$.

Proposition 1.5 (Proper restricted lower bounds do not separate P and NP)**.** *Let $\mathscr{C}_N \subseteq C_{N,N^k}$ be a restricted circuit class. A proof that $F_N \notin \mathscr{C}_N$ for infinitely many N does not imply $F_N \notin C_{N,N^k}$ unless an independent simulation theorem proves that all size-N^k circuits belong to $\mathscr{C}_N$ or reduce to it without losing the lower-bound property.*

Proof. The inclusion is in the wrong direction for the conclusion: a function may fail to lie in a small subclass while lying in the larger class. A simulation theorem or an invariant ceiling for the entire larger class is required. □

This elementary proposition governs the interpretation of the LMN theorem. LMN provides strong Fourier information for AC^0; it does not supply a spectral ceiling for general polynomial-size circuits. Every claimed ascent from bounded depth to P/poly must display a further theorem, and parity shows that a high-tail ceiling cannot be that theorem.

1.4 Cook reductions and the Boolean function representing SAT

Let M be a nondeterministic Turing machine running in time $T(n) = n^d$. A computation tableau for M on input x has $T(n)$ time rows and $T(n)$ tape cells sufficient for the computation. Variables encode tape symbols, head positions and states at each grid location. Local consistency conditions assert that consecutive rows follow the transition rule, the initial row represents x, and the final row is accepting. Each local condition is expressible by a bounded-width Boolean clause after introducing a bounded number of auxiliary variables. The conjunction gives a formula $\Phi_{M,x}$ of size $O(T(n)^2)$ such that

$$x \in L(M) \quad \Longleftrightarrow \quad \Phi_{M,x} \text{ is satisfiable.} \tag{20.1}$$

This is the structural content of the Cook–Levin reduction and explains the role of the encoded SAT language.

The assignment function $f_{\Phi_{M,x}}$ lives on the witness/tableau variables of a particular input x. The language function for $L(M)$ lives on x itself. The reduction $x \mapsto \mathrm{enc}(\Phi_{M,x})$ transports decisions by composition:

$$\mathbf{1}_{L(M)}(x) = F_{p(n)}(R_n(x)), \tag{20.2}$$

where R_n is a polynomial-size circuit implementing the reduction and $p(n) = O(T(n)^2 \log T(n))$ for a standard encoding. The transformation does not assert any simple relation between the Fourier levels of $F_{p(n)}$ and the Fourier levels of $f_{\Phi_{M,x}}$.

Lemma 1.6 (Formula-assignment geometry is instance local)**.** *Let x and x' be two inputs of equal length. The satisfying assignments of $\Phi_{M,x}$ and $\Phi_{M,x'}$ describe accepting computation tableaux for different fixed input rows. Even if both sets have identical cardinality or cluster statistics, this does not determine the value of any nonconstant Fourier coefficient of the language slice $\mathbf{1}_{L(M)}$.*

Proof. A Fourier coefficient of the language slice is an average over all inputs of the length, weighted by a character of input bits. Data about the witnesses attached to two individual inputs do not determine this average. More generally, any list of fibrewise geometric summaries must be connected to the input-character averages by a proved operator identity before it yields a Fourier statement about the language. □

A distributional method may sample encoded formulas rather than tableaux. In that case, a lower bound against circuits on the encoding strings is logically legitimate. But a distribution on satisfying assignments of a sampled formula is a different object. The distinction is exactly the difference between witness geometry and decision complexity.

2 Fourier–Walsh analysis and elementary circuit spectra

2.1 Fourier–Walsh calculus on the Boolean cube

We record the analytic statements that are valid without assumptions on the computing model. Equip $\{-1,+1\}^N$ with uniform probability measure μ_N. For real functions f, g put

$$\langle f, g\rangle = \mathbb{E}_x[f(x)g(x)], \qquad \|f\|_2^2 = \langle f, f\rangle.$$

Lemma 2.1 (Character orthogonality)**.** *For $S, T \subseteq [N]$,*

$$\langle \chi_S, \chi_T\rangle = \mathbf{1}_{S=T}.$$

Consequently the 2^N functions χ_S form an orthonormal basis of the real function space on $\{-1,+1\}^N$.

Proof. Since $\chi_S\chi_T = \chi_{S\triangle T}$, independence of the coordinates gives

$$\mathbb{E}\chi_{S\triangle T} = \prod_{i\in S\triangle T} \mathbb{E}x_i.$$

This product is 1 for $S = T$ and 0 otherwise. The number of characters equals the dimension of the function space. □

Theorem 2.2 (Fourier inversion and Parseval)**.** *For every $f : \{-1,+1\}^N \to \mathbb{R}$,*

$$f(x) = \sum_{S\subseteq[N]} \widehat{f}(S)\chi_S(x), \qquad \|f\|_2^2 = \sum_{S\subseteq[N]} \widehat{f}(S)^2. \tag{3.1}$$

For Boolean-valued f, $\sum_{j=0}^N W_j(f) = 1$.

Proof. Expand f in the orthonormal basis and take its inner product with χ_S to obtain the coefficient. Parseval is the Pythagorean identity in this finite-dimensional Hilbert space. For $f \in \{-1,+1\}$ one has $\|f\|_2^2 = 1$. □

For $\rho \in [-1,1]$ define the noise operator

$$T_\rho f = \sum_{S\subseteq[N]} \rho^{|S|}\widehat{f}(S)\chi_S. \tag{3.2}$$

It admits a probabilistic form: if y is obtained from x by independently retaining each coordinate with correlation ρ, then $T_\rho f(x) = \mathbb{E}[f(y) \mid x]$.

Proposition 2.3 (Energy and noise identities)**.** *For every real f,*

$$\langle f, T_\rho f\rangle = \sum_{j=0}^{N} \rho^j W_j(f), \qquad \|T_\rho f\|_2^2 = \sum_{j=0}^{N} \rho^{2j} W_j(f). \tag{3.3}$$

Moreover, if $\partial_i f(x) = \frac{1}{2}(f(x) - f(x^{\oplus i}))$, then

$$\sum_{i=1}^{N} \|\partial_i f\|_2^2 = \sum_{S\subseteq[N]} |S|\widehat{f}(S)^2. \tag{3.4}$$

Proof. The first two identities follow by applying Parseval to (3.2). Since $\partial_i \chi_S$ is χ_S for $i \in S$ and zero otherwise, Parseval applied to each derivative and summation over i gives (3.4). □

The Bonami–Beckner inequality is the essential analytic estimate used in bounded-depth circuit theory. It says that for $1 \le p \le q$ and $0 \le \rho \le \sqrt{(p-1)/(q-1)}$,

$$\|T_\rho f\|_q \le \|f\|_p. \tag{3.5}$$

Its proof is obtained from the two-point inequality on $\{-1, +1\}$ and tensorization over coordinates. The important point for the present paper is not the validity of (3.5), which is classical and exact, but the class of circuits to which a low-degree Fourier conclusion can be attached. Hypercontractivity alone does not state that general polynomial-size circuits have small high-degree Fourier tails.

2.2 Martingale decomposition and Fourier entropy

The orthogonal basis admits a filtration form which is useful when one attempts to measure how information enters a Boolean computation. Let $\mathscr{F}_j$ be the sigma-algebra generated by the first j coordinates and write

$$M_j f = \mathbb{E}[f \mid \mathscr{F}_j], \qquad D_j f = M_j f - M_{j-1} f. \tag{8.1}$$

Then $(M_j f)_{j=0}^N$ is the Doob martingale of f and the differences are orthogonal. In Fourier coordinates,

$$M_j f = \sum_{S\subseteq[j]} \widehat{f}(S)\chi_S, \qquad D_j f = \sum_{\substack{S\subseteq[j]\\ j\in S}} \widehat{f}(S)\chi_S. \tag{8.2}$$

Consequently,

$$\mathrm{Var}(f) = \sum_{j=1}^{N} \|D_j f\|_2^2 = \sum_{S\neq\varnothing} \widehat{f}(S)^2. \tag{8.3}$$

Proposition 2.4 (Efron–Stein identity on the cube)**.** *Let $x^{(i)}$ be obtained from x by resampling coordinate i independently. Then*

$$\mathrm{Var}(f) \le \frac{1}{2}\sum_{i=1}^{N} \mathbb{E}\big(f(x) - f(x^{(i)})\big)^2 = \sum_{S\subseteq[N]} |S|\widehat{f}(S)^2. \tag{8.4}$$

Equality in the first inequality occurs precisely when all nonzero Fourier coefficients lie at degree one.

Proof. The right-hand expression follows by applying the derivative calculation (3.4) with resampling instead of sign flip. Since $|S| \ge 1$ for $S \neq \varnothing$, Parseval gives

$$\sum_{S\neq\varnothing} \widehat{f}(S)^2 \le \sum_{S\neq\varnothing} |S|\widehat{f}(S)^2.$$

The equality characterization is immediate. □

For f with $\|f\|_2 = 1$, define its level entropy by

$$\mathsf{H}_{\mathrm{lev}}(f) = -\sum_{j=0}^{N} W_j(f)\log W_j(f), \tag{8.5}$$

with $0\log 0 = 0$. This functional is invariant under coordinate permutations and records spreading between Fourier levels. It does not record circuit hardness.

Proposition 2.5 (Level entropy does not imply hardness)**.** *There are polynomial-size circuit families with level entropy* 0*, and polynomial-size circuit families with level entropy bounded away from* 0*. In particular, neither small nor large level entropy is by itself a lower-bound criterion for general circuits.*

Proof. Parity has $W_N = 1$ and therefore entropy 0, while any dictator has $W_1 = 1$ and also entropy 0; both have linear circuits. On the other hand, the OR function has polynomial-size circuits and a nontrivial Fourier spectrum distributed over several levels, hence positive level entropy. Tensor products and compositions of such elementary circuits provide polynomial-size functions with many distinct spectral distributions. □

The same warning applies to entropy weighted by a graph curvature: unless a universal circuit ceiling is proved after testing explicit easy functions, curvature merely changes the coordinate system in which a spectral statistic is displayed.

2.3 A complete two-point hypercontractive calculation

The hypercontractive inequality is frequently invoked in Fourier arguments, so we record its elementary core. On the two-point space, let $g(x) = a + bx$, $x \in \{-1, +1\}$, and let $T_\rho g = a + \rho bx$. For $1 \le p \le q$ and $\rho^2 \le (p-1)/(q-1)$, the two-point inequality reads

$$\left(\frac{|a+\rho b|^q + |a-\rho b|^q}{2}\right)^{1/q} \le \left(\frac{|a+b|^p + |a-b|^p}{2}\right)^{1/p}. \tag{9.1}$$

A proof is obtained by homogeneity and convexity after reducing to $a = 1$ and differentiating in b^2. In the important $(p, q) = (2, 4)$ case one obtains a transparent calculation:

$$\|T_\rho g\|_4^4 = a^4 + 6\rho^2a^2b^2 + \rho^4b^4, \qquad \|g\|_2^4 = (a^2 + b^2)^2. \tag{9.2}$$

Thus $\|T_\rho g\|_4 \le \|g\|_2$ whenever $6\rho^2 \le 2$ and $\rho^4 \le 1$, and in particular for $\rho \le 1/\sqrt{3}$, the sharp value for $2 \longrightarrow 4$ hypercontractivity.

Theorem 2.6 (Tensorized 2 ⟶ 4 inequality)**.** *For every* $f : \{-1, +1\}^N \to \mathbb{R}$*,*

$$\|T_{1/\sqrt{3}}f\|_4 \le \|f\|_2. \tag{9.3}$$

Consequently,

$$\sum_{S\subseteq[N]} 3^{-|S|}\widehat{f}(S)^2 \le \|f\|_2^2, \qquad \|T_{1/\sqrt{3}}f\|_2 \le \|f\|_2. \tag{9.4}$$

Proof. Write $f(x', x_N) = u(x') + x_N v(x')$. Applying (9.1) in the last coordinate pointwise and then Minkowski's inequality gives the N-variable estimate from the $(N-1)$-variable estimate. Induction begins with (9.2). The second display follows from Parseval and contraction in L^2. □

For a degree at most d polynomial P on the cube, the inverse form implies

$$\|P\|_4 \le 3^{d/2}\|P\|_2. \tag{9.5}$$

This inequality is central when a restriction theorem has already reduced a circuit to a low-degree polynomial or a shallow decision tree. It supplies no such reduction for general circuits. Parity again illustrates the point: it is a degree-N character and remains a compact circuit despite the factor $3^{N/2}$ in (9.5).

2.4 Spectral calculations for elementary gates

The Fourier transform of small-circuit examples is a necessary check on any proposed circuit invariant. We give exact formulas in the $\{-1,+1\}$ convention. For $x_i = +1$ interpreted as true, the indicator of the all-true assignment is

$$\mathbf{1}_{\{(+1,\dots,+1)\}}(x) = 2^{-N}\prod_{i=1}^{N}(1+x_i) = 2^{-N}\sum_{S\subseteq[N]}\chi_S(x). \tag{21.1}$$

The Boolean sign version of AND is $A_N = 2\mathbf{1}_{\{(+1,\dots,+1)\}} - 1$, giving (11.2). Its level weights are

$$W_0(A_N) = \left(-1+2^{1-N}\right)^2, \qquad W_j(A_N) = \binom{N}{j}4^{1-N} \quad (1\le j\le N). \tag{21.2}$$

These weights sum to one by the binomial identity.

For OR, using $O_N(x) = -A_N(-x)$, one obtains

$$\widehat{O}_N(\varnothing) = 1-2^{1-N}, \qquad \widehat{O}_N(S) = (-1)^{|S|+1}2^{1-N} \quad (S\neq\varnothing). \tag{21.3}$$

Thus AND and OR have identical level-mass profiles even though their signed spectra differ.

For parity $P_N = \chi_{[N]}$, as above, $W_N(P_N) = 1$. For an XOR of a set T of variables, $P_T = \chi_T$ and all spectral mass lies in level $|T|$. Since a balanced XOR tree has linear size, for every level j there is a small circuit with all Fourier mass at that level by taking parity on j selected inputs and ignoring the rest.

Theorem 2.7 (Every Fourier level contains a small circuit)**.** *For each* $0\le j\le N$*, there is an* $O(N)$*-size circuit* $g_{N,j}$ *satisfying* $W_j(g_{N,j}) = 1$ *and* $W_i(g_{N,j}) = 0$ *for* $i\neq j$*.*

Proof. Take $g_{N,0} = 1$ and, for $j\ge 1$, take $g_{N,j}(x) = \prod_{i=1}^{j}x_i$. This is parity on the first j inputs and is computed by a balanced XOR tree of size $O(j)\le O(N)$. Its Fourier expansion consists of the single character $\chi_{\{1,\dots,j\}}$. □

Corollary 2.8 (Level locations cannot certify general circuit hardness)**.** *No separation principle whose circuit upper bound forbids concentration of Fourier mass at a specified collection of levels can apply to every polynomial-size circuit unless that collection is empty.*

Proof. If the collection contains j, choose $g_{N,j}$ from the theorem. □

The corollary is stronger than the parity check: every individual level is occupied by an explicit small circuit. A viable invariant must depend on relationships not exhausted by the distribution of squared mass between levels.

2.5 Noise stability and why it is not a general separator

The noise stability of f at correlation ρ is

$$\mathrm{Stab}_\rho(f) = \langle f, T_\rho f\rangle = \sum_{j=0}^{N}\rho^j W_j(f). \tag{22.1}$$

For parity on j inputs, $\mathrm{Stab}_\rho(P_j) = \rho^j$. For a dictator it is ρ, while for AND the exact expression is obtained from (21.2):

$$\mathrm{Stab}_\rho(A_N) = \left(-1+2^{1-N}\right)^2 + 4^{1-N}\left((1+\rho)^N - 1\right). \tag{22.2}$$

All three functions have small circuits and exhibit very different noise profiles.

Proposition 2.9 (One-parameter noise profiles are insufficient)**.** *Let* $\rho\in(0,1)$ *be fixed. Neither an upper nor a lower threshold on* $\mathrm{Stab}_\rho(f)$ *can distinguish all polynomial-size circuits from a putative hard Boolean function.*

Proof. Constants have stability one; parity on N bits has stability ρ^N; dictators have stability ρ; and all are polynomial-size circuits. Hence the class of polynomial-size circuits spans stability values from exponentially small to one. A one-sided threshold that excludes one end includes easy functions at the other. □

A multi-parameter profile $\rho \mapsto \mathrm{Stab}_\rho(f)$ is equivalent to the level-mass generating polynomial. By the theorem of the preceding section, easy functions already realize monomials ρ^j for all j. Nonnegative mixtures occur as profiles of convex combinations of easy spectra, although convex mixtures need not be Boolean functions. This reinforces the need for nonlinear circuit information.

2.6 Influence, sensitivity, and certificate structure

For Boolean f, the influence of coordinate i is

$$\mathrm{Inf}_i(f) = \Pr_x[f(x) \neq f(x^{\oplus i})] = \sum_{S \ni i} \widehat{f}(S)^2. \tag{23.1}$$

The total influence is

$$I(f) = \sum_{i=1}^{N} \mathrm{Inf}_i(f) = \sum_S |S|\widehat{f}(S)^2. \tag{23.2}$$

Parity has $I(P_N) = N$ and an $O(N)$ circuit; a dictator has $I(D) = 1$ and a constant-size circuit. Majority has total influence of order $\sqrt{N}$ and polynomial-size circuits. Thus total influence, like Fourier tail, is not a universal general-circuit lower-bound witness.

Nevertheless influence is useful for local geometry. If a satisfying assignment of a formula is isolated under Hamming-neighbor flips, then the assignment indicator has local boundary in every coordinate that appears critically in the instance. Such boundary information describes the assignment function f_Φ. To obtain a language lower bound one would need to lift it to influence of bits of the formula encoding F_N: flipping an encoding bit may change a literal, delete a clause, create an invalid string, or change the parsing convention. The two influence structures are unrelated without an encoding theorem.

Lemma 2.10 (Encoding influence is syntactic)**.** *For the language function F_N, $\mathrm{Inf}_i(F_N)$ measures the fraction of encoded strings whose satisfiability status changes when bit i of the description is flipped. It is not the fraction of satisfying assignments destroyed by changing a variable assignment.*

Proof. This is immediate from the domain of F_N: its input coordinates index bits of the encoded instance. The influence definition changes an instance description, not a witness assignment. □

This lemma makes precise why a cluster decomposition in witness space cannot directly produce an influence or curvature lower bound on formula strings. It may guide a construction of syntactically stable formula distributions, but that construction is an additional theorem.

3 Restrictions and the exact bounded-depth range

3.1 Restricted circuits: what the classical Fourier theory proves

A Boolean circuit is in AC^0 if it has unbounded fan-in AND and OR gates, NOT gates, constant depth, and polynomial size. The Linial–Mansour–Nisan theorem applies to this constant-depth class. In one standard form, for a depth-d AC^0 circuit C of size M and every $\varepsilon > 0$, the Fourier mass outside degree

$$t = (\log(M/\varepsilon))^{O(d)}$$

is at most ε:

$$\sum_{|S|>t} \widehat{C}(S)^2 \leq \varepsilon. \tag{4.1}$$

This conclusion depends on constant depth; it is produced by iterated random restrictions and switching lemmas that collapse bounded-depth DNFs and CNFs to shallow decision trees.

Theorem 3.1 (Restricted-model implication)**.** *Suppose a language slice F_N has, for infinitely many N, Fourier tail*

$$\sum_{|S|>t_N} \widehat{F}_N(S)^2 > \varepsilon_N \tag{4.2}$$

for thresholds (t_N, ε_N) violating the LMN bound for every polynomial-size constant-depth circuit family. Then the language is not in polynomial-size AC^0.

Proof. If the language were computed by a polynomial-size constant-depth circuit family, then the LMN theorem would give the opposite inequality for all sufficiently large N. Equation (4.2) contradicts that conclusion on infinitely many lengths. □

The theorem is useful but is not a P versus NP separation. Polynomial-size circuits of unrestricted depth strictly contain polynomial-size AC^0 circuits. A route to SAT $\notin P/\mathrm{poly}$ must handle the unrestricted class $\bigcup_k C_{N,k}$, not only the bounded-depth subfamilies.

Lemma 3.2 (A polynomial-size parity circuit)**.** *The parity function*

$$\mathrm{PAR}_N(x) = \prod_{i=1}^{N} x_i = \chi_{[N]}(x)$$

is computable by a Boolean circuit over AND, OR and NOT gates of size $O(N)$ and depth $O(\log N)$, while

$$W_N(\mathrm{PAR}_N) = 1, \qquad W_j(\mathrm{PAR}_N) = 0 \quad (j \neq N). \tag{4.3}$$

Proof. An XOR of two Boolean bits is expressed by a constant-size AND/OR/NOT circuit. A balanced binary tree of XOR gates computes parity using $O(N)$ gates and depth $O(\log N)$. In the $\{-1, +1\}$ convention parity is exactly the character $\chi_{[N]}$, so its Fourier expansion has one coefficient equal to 1 at degree N and all others zero. □

Corollary 3.3 (Failure of a universal high-tail ceiling)**.** *No theorem of the form*

$$\sum_{|S|>N^{\alpha}} \widehat{C}(S)^2 \le 2^{-\Omega(N)} \tag{4.4}$$

can hold for every polynomial-size general Boolean circuit C when $0 < \alpha < 1$.

Proof. For sufficiently large N, degree N lies above N^{α}. Taking $C = \mathrm{PAR}_N$ makes the left-hand side of (4.4) equal to 1 by (4.3). □

This corollary is decisive for any Fourier-tail argument aimed at general circuits: a high-degree spectral tail is compatible with extremely small circuits. General lower bounds require a more selective invariant than high-level mass alone.

3.2 Restriction methods, uniformity, and the general-circuit boundary

Random restrictions are an exceptionally successful tool for circuit classes whose gates simplify under partial assignment. For constant-depth circuits, successive switching arguments yield decision-tree descriptions after random restriction, leading to Fourier concentration and learnability. For unrestricted polynomial-size circuits, depth may grow with N, and the same reduction does not produce a polynomially bounded terminal decision tree.

Let $\rho \in \{0, 1, *\}^N$ be a restriction, with live coordinate set $L(\rho) = \{i : \rho_i = *\}$. For a function f write $f \restriction \rho$ for its restriction to live variables. The restriction operator is compatible with polynomial-size circuits:

$$C \in C_{N,k} \quad \Longrightarrow \quad C \restriction \rho \in C_{|L(\rho)|,k'} \tag{7.1}$$

for a harmless change of encoding convention and exponent. The converse is false: simplicity after many restrictions need not give a small circuit before restriction.

Lemma 3.4 (Restriction energy identity)**.** *Choose a random restriction that leaves each variable live independently with probability θ and otherwise assigns it a uniform sign. Then for every $f : \{-1,+1\}^N \to \mathbb{R}$,*

$$\mathbb{E}_\rho \sum_{S\subseteq L(\rho)} \widehat{(f \restriction \rho)}(S)^2 z^{|S|} = \sum_{T\subseteq[N]} \widehat{f}(T)^2 (1-\theta+\theta z)^{|T|}. \tag{7.2}$$

Proof. It suffices by orthogonality to verify the identity for $f = \chi_T$. Under a restriction, each coordinate in T either remains as a character factor, with probability θ, or becomes a random constant sign, with probability $1-\theta$. Squaring erases the constant sign. The generating factor per coordinate is $1-\theta+\theta z$, and multiplication over T gives the right-hand side. Linear combination and Parseval complete the proof. □

Equation (7.2) is an exact bridge between Fourier energies before and after restrictions. It does not, however, give a general-circuit lower bound: parity remains parity on its live variables, and a small XOR tree remains a small circuit after every restriction. Thus any descent scheme based only on reducing Fourier degree under random restrictions encounters the same obstruction.

3.3 A detailed Walsh calculation for restrictions

We expand the restriction identity because it is the exact analytic interface between local simplification and global spectra. Let a restriction ρ be chosen as follows: independently for each coordinate i, with probability θ retain x_i as a live variable, and with probability $1-\theta$ replace it by a uniform independent sign σ_i. For $S \subseteq [N]$, a character restricts to

$$\chi_S \restriction \rho = \left(\prod_{i\in S\setminus L(\rho)} \sigma_i\right) \chi_{S\cap L(\rho)}. \tag{31.1}$$

If distinct original sets S and T restrict to the same live set, their random fixed-coordinate signs may correlate only when $S \triangle T$ contains no fixed coordinate. Averaging over the fixed signs annihilates cross terms unless $S = T$ on every fixed coordinate. This yields a more precise version of (7.2).

Theorem 3.5 (Restriction Fourier-energy formula)**.** *For $0 \le z \le 1$ and every real function f,*

$$\mathbb{E}_\rho \sum_{U\subseteq L(\rho)} z^{|U|} \widehat{(f \restriction \rho)}(U)^2 = \sum_{S\subseteq[N]} (1-\theta+\theta z)^{|S|} \widehat{f}(S)^2. \tag{31.2}$$

In particular, at $z = 0$,

$$\mathbb{E}_\rho \widehat{(f \restriction \rho)}(\varnothing)^2 = \sum_S (1-\theta)^{|S|} \widehat{f}(S)^2, \tag{31.3}$$

and at $z = 1$ both sides equal $\|f\|_2^2$.

Proof. Expand $f = \sum_S \widehat{f}(S)\chi_S$ and restrict term by term using (31.1). For a fixed pair S, T, the coefficient product contributing to a live character has an expectation over fixed signs. If a fixed coordinate belongs to exactly one of S, T, its sign has mean zero, eliminating the cross term. The remaining diagonal contribution of S has each coordinate either fixed, contributing weight $1-\theta$, or live, contributing weight θz. Multiplication gives $(1-\theta+\theta z)^{|S|}$. □

For parity P_N, the right-hand side is $(1-\theta+\theta z)^N$. Thus for a restriction leaving an expected linear number of variables alive, its restricted spectrum remains concentrated at the live parity degree. This exact identity prevents any unqualified use of restriction averaging to force low degree on general circuits.

Corollary 3.6 (Parity survives informative restrictions)**.** *If θ is bounded below by a positive constant, then with probability tending to one, $P_N \restriction \rho$ has degree at least $\theta N/2$, while it is computed by a circuit of size $O(N)$.*

Proof. The number of live variables is binomial with mean θN and is at least $\theta N/2$ with probability tending to one by a Chernoff bound. The restriction of parity is parity on the live variables up to an overall sign. □

3.4 Decision trees, Fourier support, and bounded-depth collapse

A decision tree of depth d queries at most d variables along every root-to-leaf path. Its represented function is a sum over leaves of products of at most d coordinate indicators. Each indicator $(1 \pm x_i)/2$ has Fourier degree one.

Lemma 3.7 (Decision-tree degree bound)**.** *If f is computed by a decision tree of depth at most d, then $\widehat{f}(S) = 0$ for all $|S| > d$.*

Proof. Expand the leaf indicators. Every monomial arising from a root-to-leaf path uses no more than the queried variables on that path, so it has degree at most d. Summing over leaves cannot introduce higher degree. □

A restriction lemma for a DNF or CNF says, roughly, that after keeping a small random fraction of variables alive, a bounded-width formula becomes a small decision tree with high probability. Iterating this phenomenon through a constant number of layers yields the Fourier-tail theorem for AC^0. Two features are essential: bounded depth and a class of gates for which restrictions expose a decision-tree representation.

Proposition 3.8 (Why depth cannot be omitted)**.** *No argument whose only terminal conclusion is that "after restriction, the function has small Fourier degree with high probability" can establish a corresponding conclusion for all polynomial-size circuits unless it excludes polynomial-size parity circuits by an additional hypothesis.*

Proof. Under any restriction leaving m variables alive, parity restricts either to a sign or to parity on the m live variables. When m is nonzero its Fourier degree is exactly m. The circuit remains of size $O(m)$. Thus restrictions do not force a low-degree approximation for this polynomial-size family unless the restriction kills nearly every variable; such a restriction carries too little information to prove a general lower bound for an arbitrary target. □

It follows that any use of switching lemmas in a general-circuit proof requires a separate theorem reducing unrestricted circuits to a switchable form without changing the computed function or increasing size beyond control. Such a reduction would already be a major lower-bound breakthrough; it is not supplied by the classical switching lemma.

3.5 The switching-lemma route in its valid domain

We state the standard form needed to explain the exact boundary of the Fourier method. A term is a conjunction of literals, and a width-w DNF is an OR of terms of width at most w. A restriction distribution $\mathcal{R}_p$ independently leaves a coordinate live with probability p and otherwise fixes it uniformly.

Theorem 3.9 (Switching lemma, standard restricted form)**.** *There is an absolute constant C such that for every width-w DNF G and every integer $t \geq 1$,*

$$\Pr_{\rho \sim \mathcal{R}_p}\left[\mathrm{DT}(G \upharpoonright \rho) \geq t\right] \leq (Cpw)^t. \tag{32.1}$$

The dual statement holds for width-w CNFs.

The proof encodes each bad restriction admitting a canonical long decision-tree path by a shorter restriction together with advice identifying at most t selected terms and queried literals. The number of advice strings is bounded by $(Cw)^t$, while the ratio of probabilities between the encoded restriction and the original supplies a factor p^t. Injectivity of the encoding gives (32.1). This proof depends on the syntactic representation as a bounded-width DNF or CNF.

Now let C be a depth-d AC^0 circuit of size M. After applying De Morgan transformations at alternating levels and taking restrictions with carefully chosen live probabilities, one collapses the bottom layer by the switching lemma, repeats through d layers, and obtains a shallow decision tree with failure probability ε. Tracking the live probability yields a degree threshold of polylogarithmic size in M/ε for fixed d.

Theorem 3.10 (LMN conclusion in the correct scope)**.** *For fixed depth d, there is $a_d > 0$ such that every size-M depth-d AC^0 function C satisfies*

$$\sum_{|S|>(\log(M/\varepsilon))^{a_d}} \widehat{C}(S)^2 \le \varepsilon. \tag{32.2}$$

The theorem is quoted in its classical scope because its complete canonical-path proof is an established result in the fixed bibliography. What matters for the present analysis is the quantifier "depth-d AC^0 function." No theorem in this argument allows d to grow sufficiently to include balanced parity trees or arbitrary polynomial-size circuits. Taking $d = O(\log N)$ destroys the polylogarithmic degree conclusion and includes the parity counterexample.

3.6 A restricted switching statement and its exact consequence

The classical switching method may be stated without overextending its conclusion. Let F be a width-w DNF. Under a random restriction leaving each variable live with probability p, the probability that the restricted formula requires decision-tree depth at least t is bounded by a quantity of the form

$$\Pr\big[\mathrm{DT}(F \upharpoonright \rho) \ge t\big] \le (Cpw)^t \tag{24.1}$$

for an absolute constant C under the usual restriction distribution. Applying this to alternating shallow layers and choosing p below a power of the inverse logarithmic size yields the bounded-depth collapse used in the LMN theorem.

Proposition 3.11 (Fourier tail from a restricted decision tree)**.** *Suppose that a random restriction leaves m live variables and that with probability at least $1-\delta$ the restricted function is computed by a decision tree of depth d. Then the expected Fourier mass above degree d of the restricted function is at most δ:*

$$\mathbb{E}_\rho \sum_{|S|>d} \widehat{(f \upharpoonright \rho)}(S)^2 \le \delta. \tag{24.2}$$

Proof. On the good event, the mass is zero by the decision-tree degree lemma. On the bad event, Parseval bounds the total mass by one for Boolean functions. Average the two cases. □

Combining (24.2) with the restriction identity (7.2) is the standard route from switching to level concentration for shallow circuits. At no point does this proof apply to an arbitrary polynomial-depth circuit. A claim that it does so must provide a new collapse theorem for general circuits, and such a theorem cannot preserve a low-degree terminal profile for parity.

4 Explicit general-circuit stress tests

4.1 Fourier geometry of explicit small circuits

The parity stress test is not isolated. Several elementary circuit families occupy very different Fourier regions.

For a dictator $D(x) = x_1$, one has $\widehat{D}(\{1\}) = 1$. For parity, the unique nonzero coefficient occurs at $[N]$. For the conjunction written in $\{-1,+1\}$ form,

$$A_N(x) = 2\prod_{i=1}^{N} \frac{1+x_i}{2} - 1, \tag{11.1}$$

one computes

$$\widehat{A}_N(\varnothing) = -1 + 2^{1-N}, \qquad \widehat{A}_N(S) = 2^{1-N} \quad (S \ne \varnothing). \tag{11.2}$$

Thus a one-gate unbounded fan-in AND has nonzero coefficients at every level. Its high-tail energy is

$$\sum_{|S|>t} \widehat{A}_N(S)^2 = 4^{1-N} \sum_{j>t} \binom{N}{j}, \tag{11.3}$$

which can be substantial or negligible depending on the threshold. There is no monotone relation between the existence of high Fourier levels and unrestricted circuit size.

Proposition 4.1 (No level-profile separator for all easy families)**.** *A functional depending only monotonically on the quantities*

$$\Big(\sum_{|S|\geq t} \widehat{f}(S)^2\Big)_{t=0}^{N} \tag{11.4}$$

cannot use "more high-degree mass" as a universal witness of large general-circuit complexity.

Proof. Parity maximizes every high-tail sum below level N while having linear circuit size. Therefore increasing the high-degree profile cannot force large general-circuit complexity. □

The proposition does not exclude all Fourier approaches. An invariant could use signed phase relations, algebraic structure among coefficients, robustness under restrictions, or incompatibility with a carefully defined class of circuit-generated spectra. But each such enhancement needs its own universal theorem. A graph built from spectral magnitudes and positive curvature cannot bypass the example merely by changing terminology.

4.2 The failed high-tail route and its precise correction

The claimed general-circuit ceiling in the fixed abstract would require a statement that contradicts Corollary 3.3. We record this contradiction in theorem form so that no subsequent argument can depend on the invalid estimate.

Theorem 4.2 (High-tail no-go for general circuits)**.** *Let $t_N < N$ for infinitely many N, and let $\varepsilon_N < 1$ for those N. The property*

$$\sum_{|S|>t_N} \widehat{C}(S)^2 \leq \varepsilon_N \tag{11.1}$$

cannot hold for all polynomial-size Boolean circuits C on N inputs.

Proof. For each such N choose the linear-size parity circuit. Its entire spectral mass lies at degree $N > t_N$, so the left-hand side of (11.1) is $1 > \varepsilon_N$. □

The correct form of a spectral lower-bound strategy must therefore be selective. One must identify a property $\mathcal{A}_N(f)$ satisfying both

$$\mathcal{A}_N(F_N) \geq a_N, \qquad \sup_{C\in\mathcal{C}_{N,N^k}} \mathcal{A}_N(C) < a_N \tag{11.2}$$

for infinitely many N and for every fixed k, while explicitly allowing easy high-degree functions such as parity. The distance energy $\mathfrak{D}_{N^k}$ has exactly this soundness property, but obtaining a nonzero lower bound for F_N is the circuit lower-bound problem itself.

A curvature construction can be inserted only after it has been linked to (11.2). In particular, suppose a weighted Fourier graph yields a functional $\mathcal{S}_N$. The valid bridge theorem would have to be of the form

$$\mathcal{S}_N(f) > \sup_{C\in\mathcal{C}_{N,N^k}} \mathcal{S}_N(C) \quad \Longrightarrow \quad \mathfrak{D}_{N^k}(f) > 0. \tag{11.3}$$

The implication is elementary if the strict inequality is proved. The mathematical content is the universal circuit upper bound and the target-language lower bound; neither follows from the definition of curvature.

4.3 Fourier truncations and loss of exact computation

Let $P_{\leq d} f = \sum_{|S|\leq d} \widehat{f}(S)\chi_S$ be the low-degree projection. For AC^0 circuits, LMN says that $\|C - P_{\leq d}C\|_2^2$ is small at an appropriate polylogarithmic d. For general circuits this fails, as parity satisfies

$$P_{\leq d}P_N = 0 \qquad (d < N), \qquad \|P_N - P_{\leq d}P_N\|_2^2 = 1. \tag{35.1}$$

Yet P_N has linear circuits.

More generally, let $\mathcal{L}_d$ be any linear subspace generated by characters from a fixed family of degrees. If a small-circuit function is orthogonal to $\mathcal{L}_d$, then distance from $\mathcal{L}_d$ cannot certify hardness. Since parity characters at every level are small circuits, no degree-selected subspace alone separates all small circuits.

Lemma 4.3 (Projection lower bounds require circuit containment)**.** *A conclusion of the form*

$$\|F_N - P_V F_N\|_2^2 > \varepsilon \tag{35.2}$$

excludes size-s circuits only when every size-s circuit lies in, or is within error less than ε of, the subspace V.

Proof. If a size-s circuit g lies outside V by equal or greater error, (35.2) is compatible with $F_N = g$. Thus an approximation theorem for the full circuit class is required. □

This lemma precisely locates the valid content of LMN: for bounded-depth circuits, a low-degree subspace approximates the class. It cannot be used for P/poly in the absence of an approximation theorem that survives parity and other small-circuit functions.

4.4 Multilinear representations and the degree trap

Every real function on $\{-1,+1\}^N$ has a unique multilinear polynomial representation,

$$f(x) = \sum_{S\subseteq[N]} \widehat{f}(S) \prod_{i\in S} x_i. \tag{40.1}$$

The Fourier degree $\deg(f) = \max\{|S| : \widehat{f}(S) \neq 0\}$ is an invariant of the truth table, but it is not a lower bound for general circuit size. The parity circuit already has degree N and linear size. More generally, degree behaves very differently under gates and under arbitrary circuit composition.

If $u, v : \{-1,+1\}^N \to \{-1,+1\}$ represent Boolean predicates in the sign convention, then their conjunction may be expressed by

$$u \wedge v = \frac{-1 + u + v + uv}{2}, \tag{40.2}$$

and negation is $-u$. Thus the exact multilinear polynomial degree of a circuit can at most double at an AND gate and does not increase under NOT. A depth-d fan-in two circuit may have degree as large as 2^d, so depth $O(\log N)$ already permits full degree N.

Lemma 4.4 (Degree growth under De Morgan circuits)**.** *If a fan-in two circuit has depth d, then its multilinear representation has degree at most* 2^d*. This estimate is sharp up to constants for balanced parity or conjunction constructions.*

Proof. Inputs have degree one. Negation preserves degree, while (40.2) shows that conjunction and disjunction of polynomials of degrees a, b have degree at most $a + b$. Induction over depth yields 2^d. A balanced tree whose leaves are distinct literals can produce a monomial involving 2^d variables until the ambient bound N is reached. □

The lemma explains precisely why low-degree Fourier analysis naturally controls very shallow circuits but not polynomial-size circuits with logarithmic or greater depth. Any attempted universal descent from arbitrary circuits to low-degree spectra must eliminate computations such as parity while preserving their truth tables, which is impossible.

Proposition 4.5 (Approximate degree is also insufficient by itself)**.** *A theorem proving that a target function has large exact or approximate polynomial degree does not alone yield a lower bound against general polynomial-size Boolean circuits.*

Proof. Parity has exact degree N in the multilinear Fourier representation and yet has linear-size Boolean circuits. Although approximate degree is powerful for other models, its magnitude without a simulation theorem for the relevant circuit basis does not exclude general circuits. □

This does not make polynomial methods irrelevant. It says that the model must be stated accurately: approximate degree can bound decision trees, formula sizes, quantum query complexity and certain circuit restrictions, but the full P/poly target requires a measure known to be large on SAT and small for every general circuit.

4.5 Formulas, circuits, and fan-out

Formula lower bounds do not immediately become circuit lower bounds. A Boolean formula is a circuit in which each gate has fan-out one. A general circuit may reuse intermediate computations, introducing sharing that can reduce size dramatically. Any descent which unfolds a circuit into a formula may incur exponential growth in depth.

Lemma 4.6 (Unfolding cost)**.** *A fan-in two circuit of size s and depth d can be unfolded into an equivalent formula of size at most $s2^d$. There are circuit families for which straightforward unfolding duplicates subcomputations exponentially many times.*

Proof. Beginning at the output, recursively copy each predecessor subcircuit every time it is used. At depth j there are at most 2^j copied paths, each contributing at most s gate labels. Sharing is precisely what is lost in unfolding. □

Fourier and random-restriction arguments sometimes apply cleanly to formula models, read-once structures or bounded fan-out systems. To infer $P \neq NP$, they must control arbitrary sharing. A curvature graph constructed from the final Fourier spectrum does not record which intermediate subcomputations were shared, and consequently cannot obtain a general-circuit ceiling merely by formula-style descent.

Definition 4.7 (Circuit-sound transformation)**.** *A transformation $\mathcal{T}$ on Boolean functions is circuit-sound with overhead q if every size-s circuit for f yields a size-$q(s, N)$ circuit for $\mathcal{T}(f)$. A reverse obstruction is valid only if hardness of $\mathcal{T}(f)$, together with this soundness relation, forces hardness of f.*

Deleting small Fourier coefficients, replacing negative-curvature components by canonical functions, or flowing coefficients under a Laplacian is not circuit-sound until an explicit gate construction is given. In particular, replacing part of a spectrum by an XOR gadget changes the function unless the equality of truth tables is proved; approximation is insufficient for an exact decision lower bound unless its error is quantified and propagated through a distributional theorem.

4.6 Fourier phases and the failure of magnitude-only geometry

A graph based on weights $|\widehat{f}(S)|^2$ loses the signs of Fourier coefficients. This loss can identify functions with different Boolean behavior and different computational structure. Let $a = (a_S)$ be a Fourier vector and choose signs $\varepsilon_S \in \{-1, +1\}$. The vector $(\varepsilon_S a_S)$ has the same magnitude profile and therefore the same graph under a magnitude-only construction. It need not be the spectrum of a Boolean function, but when it is, the resulting function may have different symmetry and complexity.

For AND and OR the signed spectra in (21.2)–(21.3) have identical magnitudes and distinct truth tables. Both are easy, but the example proves that a magnitude graph cannot distinguish even basic logical direction. If a target lower bound depends only on magnitude curvature, it must prove a ceiling valid for every easy function sharing or dominating the relevant magnitude pattern.

Proposition 4.8 (Phase-blind invariants require phase-blind circuit ceilings)**.** *Let $\mathcal{I}(f)$ depend only on the collection $\{|\widehat{f}(S)| : S \subseteq [N]\}$. If an easy function g has the same Fourier magnitudes as a target f, then $\mathcal{I}$ cannot distinguish f from g. More generally, if $\mathcal{I}(f) > b$ is used to certify hardness, every polynomial-size function with $\mathcal{I}(g) > b$ refutes the claimed circuit ceiling.*

Proof. The first assertion is the definition of phase blindness. The second is Proposition 6.1. □

A viable Fourier invariant may use phases, convolution identities induced by gate composition, restrictions, or higher correlation tensors. Such enriched data may distinguish AND from OR or parity from other functions, but a universal ceiling remains a substantive theorem to be proved.

5 Formula clusters and the encoding boundary

5.1 Why clustered satisfying assignments do not yield a general circuit lower bound

Let $\mathcal{F}_{n,m,k}$ denote a distribution of k-CNF formulas on n variables and m clauses. A cluster theorem may state that a formula $\Phi \sim \mathcal{F}_{n,m,k}$ typically has a solution set that decomposes into components of small diameter separated by linear Hamming distance. This is valuable information about the geometry of f_Φ in (2.1). It does not automatically impose any spectral property on the encoded language function F_N in (2.2).

Definition 5.1 (Internal and external cubes)**.** *The* internal cube *of a formula* Φ *is the assignment space* $\{-1,+1\}^n$ *on which* f_Φ *is defined. The* external cube *at description length* N *is* $\{-1,+1\}^N$ *on which the language slice* F_N *is defined. An encoding-transfer theorem is a quantitative statement converting a property of a set of internal cubes into an invariant of the external slice.*

Theorem 5.2 (No transfer without an encoding theorem)**.** *Suppose* $\mathcal{P}(\Phi)$ *is any property depending only on the subset of satisfying assignments of* Φ *and holding for a positive-measure family of formulas in a distribution on* k*-CNF instances. From* $\mathcal{P}$ *alone one cannot deduce a Fourier tail lower bound for* F_N *or a circuit lower bound for* F_N*.*

Proof. The data $\mathcal{P}(\Phi)$ determine only subsets of the assignment cubes indexed by formulae in the distribution. A coefficient $\widehat{F}_N(T)$ depends on the signs of F_N across all 2^N strings and on their coordinate positions in the external cube. Neither the values on non-formula strings nor the coordinate geometry of the encoding is supplied by $\mathcal{P}$. Thus the value of even a single coefficient of F_N cannot be recovered from $\mathcal{P}$ without additional encoded-instance data. A lower bound on the external Fourier tail is therefore not a logical consequence of the internal property alone. □

An encoding-transfer theorem would need to address three separate losses. It must control the size blow-up from n variables and m clauses to N encoding bits; preserve a spectral or geometric obstruction under the reduction map; and prove that arbitrary polynomial-size circuits cannot exploit the representation of the formula string. These demands are exactly the missing bridge between phase-transition geometry and a nonuniform general-circuit lower bound.

5.2 Encoding transfer as an operator problem

Fix a syntactic encoding of k-CNF formulae by strings of length N. Let $\mathcal{E}_N$ be the subset of valid encodings and let Φ_e be the formula represented by $e \in \mathcal{E}_N$. The map

$$\mathcal{T}_N : e \longmapsto f_{\Phi_e} \tag{13.1}$$

takes a point of the external cube to a Boolean function on an internal cube whose dimension may depend on the instance. A cluster theorem is a statement about $\mathcal{T}_N(e)$ on a distribution of e; a language lower bound concerns the scalar sign $F_N(e)$.

An operator capable of transmitting internal geometry to the language slice would have to assign to every assignment function a scalar statistic $\Gamma(f_{\Phi_e})$ and prove both

$$\Gamma(f_{\Phi_e}) \neq 0 \Longrightarrow F_N(e) = +1, \quad \text{and} \quad \widehat{F_N} \text{ is quantitatively determined by } \Gamma \circ \mathcal{T}_N. \tag{13.2}$$

The first implication alone is simply a witness for satisfiability; the second is a global identity across encoded formula strings. No clustering theorem provides the second identity.

Theorem 5.3 (External-spectrum reconstruction requirement)**.** *Let* $\Gamma_N : \mathcal{E}_N \to \mathbb{R}$ *be any statistic of the satisfying-assignment geometry. To derive a Fourier lower bound on* F_N *from* Γ_N*, one must prove inequalities controlling*

$$2^{-N} \sum_{e\in\{-1,+1\}^N} F_N(e)\chi_T(e) \quad \textit{in terms of} \quad \{\Gamma_N(e) : e \in \mathcal{E}_N\} \tag{13.3}$$

for a family of characters T *sufficiently rich to force a circuit lower bound. Without such inequalities, the internal statistic has no determined external spectral consequence.*

Proof. A Fourier-tail lower bound is a sum of squares of the coefficients in (13.3). Thus every derivation of such a lower bound must supply information about those sums. A statistic that is not connected to them by a proved inequality cannot imply their size. □

This operator formulation is useful because it places the unresolved work in a concrete mathematical location: between the geometry of individual solution spaces and the spectral analysis of a language over formula encodings.

5.3 Random-SAT clusters as a source of candidate tests

A cluster theorem can still inspire a legitimate construction. Let Φ be a formula and let $\mathcal{Z}(\Phi)$ be its satisfying assignments. Suppose a statistic $G(\Phi)$ records, for example, a separation profile of $\mathcal{Z}(\Phi)$ when nonempty. To obtain a lower bound for the encoded SAT language, one could attempt to construct two distributions $\mathcal{D}_N^+$ and $\mathcal{D}_N^-$ on encoded formulas satisfying respectively $F_N = +1$ and $F_N = -1$, such that:

$$\mathbb{E}_{\mathcal{D}_N^+} G - \mathbb{E}_{\mathcal{D}_N^-} G \geq \gamma_N, \tag{36.1}$$

and every size-N^k circuit has advantage less than γ_N in distinguishing the two distributions. The second condition, not the first, is a general-circuit lower bound.

Theorem 5.4 (Distinguishing route)**.** *If distributions* $\mathcal{D}_N^+$ *and* $\mathcal{D}_N^-$ *supported respectively on satisfiable and unsatisfiable encoded instances satisfy*

$$\sup_{C\in\mathcal{C}_{N,N^k}} \left| \Pr_{e\sim\mathcal{D}_N^+}[C(e) = 1] - \Pr_{e\sim\mathcal{D}_N^-}[C(e) = 1] \right| < 1, \tag{36.2}$$

then no circuit in $\mathcal{C}_{N,N^k}$ *decides* F_N *exactly.*

Proof. An exact SAT decider outputs one on all of $\mathcal{D}_N^+$ and zero on all of $\mathcal{D}_N^-$, giving distinguishing advantage one. □

To prove a robust lower bound one would seek advantage $o(1)$ or exponentially small advantage, but any strict loss below one already excludes exact computation. This theorem opens a principled place for geometric statistics: they may be used to define or analyse distributions, provided the circuit indistinguishability estimate is actually proved.

6 Spectral curvature and Boolean realizability

6.1 Spectral invariants and the polynomial-circuit stress test

A spectral invariant is any functional of the Fourier coefficient array. For example, if G_f is a weighted graph on subsets of $[N]$ whose edge weights are determined by $|\widehat{f}(S)|^2$, a curvature functional of G_f is a spectral invariant. Such invariants may encode detailed structure, but they do not by themselves distinguish hard functions from easy ones.

Proposition 6.1 (Stress test for a spectral invariant)**.** *Let* $\mathcal{I}_N$ *be any proposed obstruction satisfying*

$$\mathcal{I}_N(F_N) \geq a_N, \qquad \mathcal{I}_N(C) \leq b_N \quad \textit{for every } C \in \mathcal{C}_{N,k}, \quad a_N > b_N. \tag{6.1}$$

If there is an explicit polynomial-size function $g_N \in \mathcal{C}_{N,k}$ *whose spectral data force* $\mathcal{I}_N(g_N) > b_N$*, then the circuit-ceiling half of (6.1) is false and the invariant does not prove the lower bound.*

Proof. The statement follows immediately from the quantified phrase "for every $C \in \mathcal{C}_{N,k}$". Substituting g_N contradicts the claimed ceiling. □

Parity is the first stress test because it has maximal high-level Fourier mass but linear circuit size. Majority, iterated XORs, addressing functions and shallow algebraic constructions offer further tests depending on the form of the invariant. Any curvature-guided descent that sends every polynomial-size circuit to a low-high-tail canonical form is therefore incompatible with Lemma 3.2 unless parity-like outputs are explicitly permitted with their large high-level mass. Once they are permitted, high-level mass no longer separates the target from small circuits.

Corollary 6.2 (A saturation functional requires an additional invariant)**.** *A functional whose separation theorem uses only a positive lower bound on high-level Fourier mass of* F_N *and an exponentially small upper bound on the corresponding mass for every polynomial-size circuit cannot yield* SAT $\notin P/$poly.

Proof. The upper bound is contradicted by the polynomial-size parity circuit. Therefore the proposed separation is not valid for the target circuit class. □

A possible repair cannot consist of changing constants in the high-level threshold. It must introduce information that parity does not share with the target and prove a universal ceiling for all polynomial-size circuits under that stronger information. Constructing such an invariant is itself tantamount to constructing a new general-circuit lower-bound method.

6.2 Spectral graph constructions and easy-function saturation

Let a spectral graph construction attach a vertex to every set S for which $|\widehat{f}(S)|$ exceeds a cutoff, and join vertices when their symmetric difference is small. If the graph weights depend only on $|\widehat{f}(S)|^2$, parity produces a graph with one active vertex. A conjunction produces a graph with active vertices on many levels; a dictator produces another one-vertex graph at low degree. Any curvature convention applied to these examples must be explicitly calculated before a universal circuit inequality is claimed.

For a one-vertex Markov graph, edge-based Ollivier curvature is either undefined or vacuous unless loops or a background graph are supplied. If loops are supplied, the resulting curvature may be maximal for parity, an easy function. For a graph defined on all subsets with weights derived from coefficients, the zero coefficients of parity force most mass to vanish and the metric convention determines the curvature arbitrarily. These are not objections to defining a graph; they show that a graph definition alone is not a circuit lower bound.

Theorem 6.3 (Curvature-ceiling stress condition)**.** *Suppose* $\mathcal{S}_N(f)$ *is defined for every Boolean function and a theorem claims* $\mathcal{S}_N(C) \le b_N$ *for every polynomial-size circuit. Then a necessary condition for the theorem is*

$$\mathcal{S}_N(P_T) \le b_N \quad \textit{for every parity character } P_T = \chi_T, \tag{25.1}$$

and similarly for conjunctions, disjunctions, dictators and their polynomial-size compositions.

Proof. Each listed function is computed by a polynomial-size circuit; substitute it into the asserted universal bound. □

This theorem is elementary but indispensable. In particular, if the claimed saturation lower bound for the target is obtained merely from its high-level spectral mass, while parity has at least as much high-level mass, then either parity violates the circuit ceiling or the target lower bound does not separate.

6.3 Curvature functionals and exact monotonicity conditions

Suppose a Fourier graph $G_f = (V, E, w)$ is constructed from coefficients of a Boolean function. A Markov kernel $P_f(S, T)$ on V may be defined by normalizing edge weights. Ollivier curvature along an edge compares the Wasserstein distance of neighboring probability measures to the edge metric. Even when this construction is well-defined, it is an invariant of a chosen graph representation, not automatically of circuit size.

Let $\mathcal{S}(f)$ denote a curvature-weighted functional. A descent step in coefficient space may be written formally as

$$\widehat{f_{t+1}} = \widehat{f_t} - \eta_t \mathcal{L}_{f_t} \widehat{f_t}, \tag{14.1}$$

where $\mathcal{L}_{f_t}$ is a graph Laplacian. The quadratic energy

$$\mathcal{E}(f_t) = \frac{1}{2} \sum_{S,T} w_t(S,T)\big(\widehat{f_t}(S) - \widehat{f_t}(T)\big)^2 \tag{14.2}$$

may decrease when the graph is fixed and the step size is controlled. In the proposed application the graph itself depends on f_t, and the evolution of edge weights contributes additional terms.

Lemma 6.4 (Moving-graph derivative). *For a differentiable coefficient path $a_t(S)$ and differentiable symmetric weights $w_t(S,T)$,*

$$\frac{d}{dt}\mathcal{E}_t = \sum_{S,T} w_t(S,T)(a_t(S) - a_t(T))(\dot{a}_t(S) - \dot{a}_t(T)) + \frac{1}{2} \sum_{S,T} \dot{w}_t(S,T)(a_t(S) - a_t(T))^2. \tag{14.3}$$

Proof. Differentiate (14.2) term by term. The first term comes from the coefficient difference and the second from the moving weights. □

Therefore a Bochner inequality for a fixed positive-curvature graph does not prove monotonicity of a functional whose graph is reconstructed after every spectral update. One must control the $\dot{w}_t$ contribution and also prove that the transformed coefficient array is the Fourier spectrum of a Boolean function computed by a circuit of controlled size.

Proposition 6.5 (Boolean-realizability condition). *A coefficient update path in $\mathbb{R}^{2^N}$ proves an upper bound for polynomial-size circuits only if every endpoint to which a circuit spectrum is sent is realized by a Boolean function computed by a circuit with the required size bound, or if the path yields an inequality valid directly for the original circuit. An arbitrary Laplacian update of Fourier coefficients does not ensure either condition.*

Proof. Most points of the Fourier unit sphere are not Fourier transforms of Boolean functions: Boolean realizability imposes the 2^N pointwise conditions that the inverse transform takes values in $\{-1, +1\}$. Even among Boolean spectra, there is no automatic preservation of small circuit size under coefficient deformation. Therefore a continuous descent in coefficient space cannot be substituted for a circuit transformation without an additional construction. □

The missing curvature theorem is thus twofold: a controlled moving-graph monotonicity statement and a circuit-preserving realization theorem. Both must be proved before curvature descent can contribute to (1.2).

6.4 Circuit-preserving transformations and spectral descent

A descent argument beginning with a small circuit C must either transform circuits within a controlled size class or establish an inequality on the untransformed circuit. A coefficient update of the form

$$\widehat{f} \longmapsto \widehat{f} - \eta L \widehat{f} \tag{26.1}$$

usually produces a real function not taking values in $\{-1, +1\}$; even if projected back to a Boolean function, there need not be a small circuit for the projection.

Let $\mathscr{B}_N \subset \mathbb{R}^{2^N}$ denote the set of Fourier transforms of Boolean functions. It is a set of 2^{2^N} isolated points on the sphere. The subset $\mathscr{S}_{N,s}$ of size-s circuit spectra is much smaller. A smooth flow on the ambient coefficient space almost never remains in $\mathscr{B}_N$, still less in $\mathscr{S}_{N,s}$.

Lemma 6.6 (Isolation of Boolean spectra)**.** *If* $f, g : \{-1,+1\}^N \to \{-1,+1\}$ *are distinct, then*

$$\|\widehat{f} - \widehat{g}\|_2 \geq 2^{1-N/2}. \tag{26.2}$$

Consequently every Boolean spectrum is isolated in Euclidean coefficient space.

Proof. By Parseval,

$$\|\widehat{f} - \widehat{g}\|_2^2 = \|f - g\|_2^2.$$

At at least one input $f - g = \pm 2$, so the uniform average of its square is at least $4/2^N$. □

Any continuous spectral flow leaving a Boolean point therefore immediately exits the Boolean truth-table set unless it is stationary up to its first discrete jump. To use a descent flow as a proof about circuits, one must define discrete circuit operations whose Fourier changes satisfy the desired inequalities. Spectral excision by deletion of coefficients has no such automatic realization.

6.5 Higher-order Fourier tensors

The loss of information in level weights suggests retaining products of coefficients. For $r \geq 2$, define an r-fold convolution tensor by

$$\mathcal{T}_r(f; S_1, \ldots, S_{r-1}) = \sum_{T \subseteq [N]} \widehat{f}(T)\widehat{f}(T \triangle S_1) \cdots \widehat{f}(T \triangle S_{r-1}). \tag{43.1}$$

These tensors encode moments of products of translated functions. For instance, by Fourier inversion,

$$\mathcal{T}_2(f; S) = \mathbb{E}_x f(x)^2 \chi_S(x), \tag{43.2}$$

which vanishes for $S \neq \varnothing$ when f is Boolean. Higher tensors carry nontrivial information about Boolean multiplication constraints.

Lemma 6.7 (Boolean convolution constraint)**.** *If* $f : \{-1,+1\}^N \to \{-1,+1\}$, *then its coefficient vector satisfies*

$$\sum_{T \subseteq [N]} \widehat{f}(T)\widehat{f}(S \triangle T) = \mathbf{1}_{S=\varnothing}. \tag{43.3}$$

Proof. The left side is the Fourier coefficient at S of f^2. Since $f^2 = 1$, its only nonzero Fourier coefficient is the constant coefficient equal to one. □

Equation (43.3) is a necessary and sufficient constraint for a real coefficient vector to represent a sign-valued function after Fourier inversion. It makes clear why an arbitrary continuous coefficient descent loses Boolean realizability: the quadratic convolution identities need not be preserved.

One may seek a higher-order separator whose feasible region consists of coefficient tensors satisfying (43.3) together with inequalities valid for size-bounded circuit compositions. Such a construction is mathematically coherent, but its containment and separation theorems are new obligations. It cannot be replaced by a magnitude-only graph with an unproved monotonicity assertion.

6.6 A Boolean-realizable spectral flow condition

Suppose $a_t(S)$ is a differentiable coefficient path intended to remain the Fourier transform of sign functions. Differentiating (43.3) gives, for every S,

$$\sum_T \dot{a}_t(T) a_t(S \triangle T) + \sum_T a_t(T) \dot{a}_t(S \triangle T) = 0. \tag{44.1}$$

Equivalently, by symmetry,

$$\sum_T \dot{a}_t(T) a_t(S \triangle T) = 0. \tag{44.2}$$

At $S = \varnothing$ this says $\langle \dot{a}_t, a_t \rangle = 0$, tangency to the unit sphere; for nonempty S it gives many further constraints.

Theorem 6.8 (No nonconstant differentiable Boolean spectral path)**.** *If $t \mapsto a_t$ is continuous and each a_t is the Fourier vector of a Boolean function on the finite cube $\{-1, +1\}^N$, then a_t is constant on every connected interval.*

Proof. There are only finitely many Boolean functions on $\{-1, +1\}^N$, hence only finitely many Boolean spectral vectors. A continuous image of a connected interval in a finite discrete set is a point. □

Thus a literal curvature flow in spectral coefficient space cannot remain within Boolean functions except by being constant. It may instead be interpreted as a relaxation or an analytic comparison path, but then a separate rounding or inequality theorem is needed to infer a statement about circuits. This exact topological observation removes any ambiguity in a continuous "descent" argument for exact Boolean computation.

7 Exact distance from small circuits

7.1 An exact Fourier distance to small circuits

We now define an invariant that is mathematically exact for general circuits. It does not itself solve the lower-bound problem; rather, it turns the required statement into an unambiguous quantitative target and provides a sound replacement for an invalid universal Fourier-tail ceiling.

For $f : \{-1, +1\}^N \to \{-1, +1\}$ and $s \geq 1$, define

$$\mathfrak{D}_s(f) = \min_{g \in C_{N,s}} \|f - g\|_2^2 = \min_{g \in C_{N,s}} \sum_{S \subseteq [N]} \left(\widehat{f}(S) - \widehat{g}(S)\right)^2, \tag{9.1}$$

where $C_{N,s}$ is the finite set of Boolean functions computable by circuits of at most s gates in a fixed finite basis.

Theorem 7.1 (Exact obstruction energy)**.** *For Boolean f, the following hold.*

(i) $\mathfrak{D}_s(f) = 0$ *if and only if f has a circuit of size at most s.*

(ii) *If* $\mathfrak{D}_s(f) > 0$, *then* $\mathfrak{D}_s(f) \geq 4 \cdot 2^{-N}$.

(iii) $\mathfrak{D}_s(f)$ *is non-increasing in s.*

(iv) *If π is a permutation or sign-flip of input coordinates, then* $\mathfrak{D}_s(f \circ \pi) = \mathfrak{D}_s(f)$ *up to a constant-size change in the circuit convention.*

Proof. Parseval gives the equality of the spatial and Fourier expressions in (9.1). The minimum is zero precisely when some admissible g equals f at every input. If $f \neq g$, their $\{-1, +1\}$ values differ by 2 at at least one input, so $\|f - g\|_2^2 \geq 4 \cdot 2^{-N}$. Monotonicity follows from inclusion of circuit classes. Coordinate symmetries are computable by rewiring and input negation, preserving size up to the stated convention. □

Corollary 7.2 (Exact nonuniform lower-bound formulation)**.** *The statement* SAT $\notin P/$poly *is equivalent to*

$$\forall k \in \mathbb{N}\ \exists^\infty N \quad \mathfrak{D}_{N^k}(F_N) > 0. \tag{9.2}$$

It is sufficient to prove the quantitative strengthening

$$\forall k \in \mathbb{N}\ \exists^\infty N \quad \mathfrak{D}_{N^k}(F_N) \geq 4 \cdot 2^{-N}. \tag{9.3}$$

Proof. Apply Theorem 7.1 to Theorem 1.1. The apparent strengthening is equivalent because any positive distance between two Boolean functions has the discrete lower bound of Theorem 7.1(ii). □

The energy $\mathfrak{D}_s$ is exact but computationally difficult, since evaluating it asks whether a circuit of bounded size computes a given truth table. Its usefulness is conceptual: any proposed smooth or geometric relaxation must lower-bound $\mathfrak{D}_s$ rather than a Fourier tail that is already large on easy functions.

7.2 The circuit-generated spectral set

Fix a finite gate basis and a size bound s. The set

$$\mathcal{S}_{N,s} = \left\{ (\widehat{g}(S))_{S\subseteq[N]} : g \in C_{N,s} \right\} \subset \mathbb{R}^{2^N} \tag{12.1}$$

is finite. Under Parseval it lies on the unit sphere when functions are Boolean-valued. The distance obstruction $\mathfrak{D}_s(f)$ of (9.1) is the squared Euclidean distance from $\widehat{f}$ to $\mathcal{S}_{N,s}$.

Lemma 7.3 (Compact separation). *If $f \notin C_{N,s}$, then there exists a linear functional L on $\mathbb{R}^{2^N}$ and a number $\gamma > 0$ such that*

$$L(\widehat{f}) \geq \sup_{u\in\mathcal{S}_{N,s}} L(u) + \gamma. \tag{12.2}$$

Proof. The finite set $\mathcal{S}_{N,s}$ is compact. Let u_0 be a closest point to $\widehat{f}$. If $\widehat{f}$ is outside the convex hull, the separating hyperplane theorem yields (12.2). If it lies in the convex hull while not in the set, a linear separator may not exist; in that case the nonlinear distance $\mathfrak{D}_s$ remains positive. Thus linear Fourier certificates detect only those lower bounds visible outside the convex hull of small-circuit spectra. □

The proof identifies another obstruction to a purely linear spectral method: taking convex combinations can erase the discreteness of circuit computation. The set of truth tables computed by small circuits is nonlinear; its convex hull may contain spectra of functions not computed by any individual small circuit.

Definition 7.4 (Convex relaxation defect). *Define*

$$\mathfrak{D}_s^{\mathrm{cvx}}(f) = \operatorname{dist}\big(\widehat{f}, \operatorname{conv}(\mathcal{S}_{N,s})\big)^2. \tag{12.3}$$

Then $0 \leq \mathfrak{D}_s^{\mathrm{cvx}}(f) \leq \mathfrak{D}_s(f)$.

A positive lower bound for $\mathfrak{D}_s^{\mathrm{cvx}}(F_N)$ would prove a circuit lower bound, and might be susceptible to duality. But proving such positivity for SAT is not made easier by high Fourier mass; the convex hull already contains spectra of parity, conjunctions, majorities and their small-circuit combinations.

7.3 A spectral norm that exactly detects disagreement

The circuit-distance functional admits useful equivalent forms. For Boolean f, g,

$$\langle f, g\rangle = 1 - 2\Pr[f \neq g], \qquad \|f - g\|_2^2 = 2 - 2\langle f, g\rangle = 4\Pr[f \neq g]. \tag{34.1}$$

Parseval therefore yields

$$\Pr[f \neq g] = \frac{1}{4} \sum_{S\subseteq[N]} \big(\widehat{f}(S) - \widehat{g}(S)\big)^2. \tag{34.2}$$

The equality is exact, not an approximation theorem. It states that a lower bound on Euclidean distance of full spectra is exactly a lower bound on uniform disagreement.

Definition 7.5 (Circuit correlation radius). *For $s \geq 1$, define*

$$\mathfrak{C}_s(f) = \max_{g\in C_{N,s}} \langle f, g\rangle. \tag{34.3}$$

Then

$$\mathfrak{D}_s(f) = 2 - 2\mathfrak{C}_s(f). \tag{34.4}$$

Proposition 7.6 (Exact computing criterion). *A Boolean function f is computed by a size-s circuit if and only if $\mathfrak{C}_s(f) = 1$. If every size-s circuit disagrees with f on at least an ε fraction of inputs, then $\mathfrak{C}_s(f) \leq 1 - 2\varepsilon$ and $\mathfrak{D}_s(f) \geq 4\varepsilon$.*

Proof. Apply (34.1) and take the maximum or minimum over $g \in C_{N,s}$. □

This form exposes the possible role of harmonic analysis: one may seek a bound on correlations $\langle F_N, g \rangle$ valid for every small circuit g. Such a bound may be proved by pseudorandomness, communication lower bounds, algebraic rigidity, geometric separators, or a new method. But a bound only for restricted circuits remains a restricted conclusion.

7.4 Tensor products, block composition, and direct-sum tests

A separator for circuit spectra must be stable not only under individual gates but also under independent block constructions. If $f : \{-1,+1\}^m \to \{-1,+1\}$ and $g : \{-1,+1\}^n \to \{-1,+1\}$ act on disjoint variables, define their product

$$(f \otimes g)(x, y) = f(x)g(y). \tag{48.1}$$

Its Fourier transform factorizes:

$$\widehat{f \otimes g}(S \sqcup T) = \widehat{f}(S)\widehat{g}(T), \tag{48.2}$$

and therefore

$$W_r(f \otimes g) = \sum_{i+j=r} W_i(f)W_j(g). \tag{48.3}$$

If f and g have circuits of sizes s and t, their product is computed with size $s + t + O(1)$ because multiplication of signs is XOR/equality up to a fixed Boolean gadget.

Proposition 7.7 (Tensor stress test)**.** *Any level-based circuit ceiling stable under polynomial-size composition must contain the convolution powers of the level distributions of every elementary small-circuit function. In particular, a ceiling excluding broad high-level concentration fails on tensor powers of parity characters.*

Proof. Equation (48.3) gives level distribution under tensor product, and the circuit-size statement keeps each finite tensor power within polynomial size when block dimensions sum to N. Taking parity on each block yields parity on their union, with mass at the total degree. □

Circuit-distance energies behave compatibly with products in one direction. If f_1, f_2 are approximated by g_1, g_2 respectively, then

$$\Pr[f_1 \otimes f_2 \neq g_1 \otimes g_2] \leq \Pr[f_1 \neq g_1] + \Pr[f_2 \neq g_2]. \tag{48.4}$$

Thus upper bounds on distance tensorize subadditively. A lower bound for a tensor product does not automatically lower-bound either factor because errors can be correlated; a separator proof must control the chosen composition explicitly.

Block composition also illustrates the encoding issue. A formula encoding can concatenate independent subformula descriptions, and satisfiability of the conjunction becomes a gate combination of the subformula language values. Any claimed obstruction for SAT must be compatible with this operation on encodings. A fibrewise solution-cluster statistic need not be stable under syntactic concatenation unless its transfer theorem is proved.

7.5 The dual formulation of a circuit separator

The finite set $\mathcal{S}_{N,s}$ allows an exact dual description of convex separation. Let $K_{N,s} = \mathrm{conv}(\mathcal{S}_{N,s})$. By finite-dimensional separation,

$$\widehat{F}_N \notin K_{N,s}$$

if and only if there exists a coefficient vector $u = (u_S)$ and a real number b such that

$$\sum_S u_S \widehat{F}_N(S) > b \quad \text{while} \quad \sum_S u_S \widehat{g}(S) \leq b \quad (g \in C_{N,s}). \tag{49.1}$$

Writing $\psi_u(x) = \sum_S u_S \chi_S(x)$, the inequality becomes

$$\langle F_N, \psi_u \rangle > b \quad \text{and} \quad \langle g, \psi_u \rangle \le b \quad (g \in C_{N,s}). \tag{49.2}$$

Thus every linear spectral separator is exactly a correlation test against a witness function ψ_u.

Theorem 7.8 (Dual witness theorem)**.** *If, for every k, infinitely many lengths admit a function $\psi_{N,k}$ and threshold $b_{N,k}$ satisfying (49.2) with $s = N^k$, then $P \neq NP$.*

Proof. The strict correlation separation implies F_N is not the spectrum of a size-N^k circuit at each chosen length. Apply Theorem 1.1. □

This theorem is a precise version of a witness-based proof programme. Entropy, curvature or geometric features may be used to construct $\psi_{N,k}$, but two estimates must be verified: a universal correlation ceiling against all size-bounded circuits and a strict target correlation. The ceiling cannot follow from the LMN theorem unless the target circuit class is restricted to AC^0.

The nonlinear distance $\mathfrak{D}_s$ can remain positive even if $\widehat{F}_N \in K_{N,s}$, in which case no linear witness separates it from the convex hull. A hierarchy of polynomial or semidefinite witnesses may be needed. This is another reason that a single scalar Fourier-tail statistic is too coarse for unrestricted circuits.

8 Distributional and reduction-stable formulations

8.1 Reduction stability and the SAT target

Let A be an NP language and let r be a polynomial-time many-one reduction from A to SAT. On an input of length n, the output length is at most $p(n)$ for a polynomial p. If SAT $\in P/\mathrm{poly}$, composition of r with the circuit family for SAT gives polynomial-size circuits for A. This is the standard closure under reductions.

For obstruction energies, the inverse direction is subtle. An invariant lower bound for a distribution of formulas cannot simply be pulled back to all NP languages unless the invariant is stable under the reduction map. If $R_n : \{-1,+1\}^n \to \{-1,+1\}^{p(n)}$ is the circuit implementing the reduction, then

$$\mathbf{1}_A(x) = F_{p(n)}(R_n(x)). \tag{12.1}$$

A Fourier coefficient of a composition with R_n generally expands nonlinearly in the Fourier coefficients of $F_{p(n)}$; it is not preserved level by level.

Proposition 8.1 (Circuit-distance stability under reductions)**.** *If $F_{p(n)}$ has a circuit of size s and R_n has a circuit of size $r(n)$, then $\mathbf{1}_A$ has a circuit of size $s + r(n) + O(1)$. Equivalently,*

$$\mathfrak{D}_{s+r(n)+O(1)}(\mathbf{1}_A) > 0 \quad \Longrightarrow \quad \mathfrak{D}_s(F_{p(n)}) > 0. \tag{12.2}$$

Proof. Compose the circuits. The contrapositive gives (12.2). □

Unlike a raw Fourier level profile, the distance obstruction is stable under arbitrary polynomial-time reductions because it is defined directly against circuit computation. This is a reason to treat it as the correct terminal quantity. The difficulty remains to lower-bound it by a tractable analytic or combinatorial functional.

8.2 Restriction-averaged obstruction energies

An invariant capable of supporting a new method should be compatible with restrictions, because restrictions are central both to circuit simplification and to local analysis of satisfiability. Let ν_θ denote the product distribution on restrictions in which a coordinate is live with probability θ and fixed to either sign otherwise. Define

$$\mathfrak{R}_{s,\theta}(f) = \mathbb{E}_{\rho\sim\nu_\theta} \mathfrak{D}_{s(\rho)}(f \restriction \rho), \tag{10.1}$$

where $s(\rho)$ is a size allowance for the restricted circuit, for example $s(\rho) = s + O(N)$ under a fixed basis convention.

Proposition 8.2 (Soundness of restriction-averaged distance)**.** *If f is computed by a circuit of size at most s, then $\mathfrak{R}_{s',\theta}(f) = 0$ for every θ and every restriction-size allowance s' large enough to contain restrictions of size-s circuits. Therefore any lower bound $\mathfrak{R}_{s',\theta}(F_N) > 0$ proves that F_N has no size-s circuit.*

Proof. Restricting a fixed circuit yields a circuit for the restricted function, with no increase beyond the chosen allowance. Every distance in (10.1) is then zero. The contrapositive gives the second statement. □

A Fourier relaxation of (10.1) may be constructed by projecting away spectra of known easy restricted functions. Let $\mathcal{V}_{s,m}$ be the linear span in $L^2(\{-1,+1\}^m)$ of all functions computed by circuits of size at most s. Define

$$\mathfrak{Q}_{s,\theta}(f) = \mathbb{E}_{\rho\sim\nu_\theta}\|P_{\mathcal{V}^{\perp}_{s,|L(\rho)|}}(f \restriction \rho)\|_2^2. \tag{10.2}$$

This is computable from Fourier coefficients once an orthogonal basis for $\mathcal{V}_{s,m}$ is specified, and it satisfies

$$\mathfrak{Q}_{s,\theta}(f) \le \mathfrak{R}_{s,\theta}(f) \tag{10.3}$$

because projection to the orthogonal complement of the span is no larger than distance to the nonlinear set that spans it.

Remark 8.3. *The span $\mathcal{V}_{s,m}$ can be the entire function space long before the nonlinear circuit class contains every Boolean function. Thus $\mathfrak{Q}$ is sound when positive but may be identically zero and hence too weak. This illustrates the central difficulty: linear spectral relaxations easily lose the nonlinear information required for general circuit lower bounds.*

8.3 Robust obstruction energies under approximation

One might weaken exact computation to approximation under the uniform distribution. Let

$$\mathfrak{D}_s^{(\varepsilon)}(f) = \min_{g\in C_{N,s}} \Pr_x[f(x) \ne g(x)]. \tag{15.1}$$

For $\{-1,+1\}$ functions,

$$\mathfrak{D}_s(f) = 4\mathfrak{D}_s^{(\varepsilon)}(f), \tag{15.2}$$

where the superscript in the left definition is merely a reminder that the quantity is an error probability. The Fourier representation remains

$$\Pr[f \ne g] = \frac{1}{4}\sum_S(\widehat{f}(S) - \widehat{g}(S))^2. \tag{15.3}$$

Theorem 8.4 (Approximate lower-bound gate)**.** *If for every fixed k there are infinitely many N such that*

$$\min_{g\in C_{N,N^k}} \Pr_x[F_N(x) \ne g(x)] \ge \varepsilon_N > 0, \tag{15.4}$$

then SAT $\notin P/\mathrm{poly}$*. If ε_N is inverse polynomial or constant, the conclusion is a correspondingly robust average-case lower bound on those slices.*

Proof. An exact computing circuit would have error zero, contradicting (15.4). The robust interpretation follows directly from the definition. □

A successful geometric programme might aim at (15.4) rather than exact distance, because averages and transport inequalities naturally produce nonzero error. But it must do so on encoded SAT slices and against general circuits. Random formula clustering may suggest a distributional route only after a reduction and a hardness amplification theorem are proved.

8.4 Distributional formulations and worst-case recovery

Let $\mathcal{D}_N$ be a distribution on valid formula encodings of length N. Define distributional circuit error

$$\mathrm{err}_{\mathcal{D}_N}(C) = \Pr_{e\sim\mathcal{D}_N}[C(e) \neq F_N(e)]. \tag{16.1}$$

A theorem proving $\mathrm{err}_{\mathcal{D}_N}(C) > 0$ for every size-N^k circuit rules out exact size-N^k computation at that length, because an exact circuit has zero error under every distribution. Thus one may use a hard distribution, but the circuit acts on encoded formulae, not on assignments.

Proposition 8.5 (Valid distributional route)**.** *Suppose that for every k there exist infinitely many N, efficiently samplable distributions $\mathcal{D}_N$ supported on encoded* SAT *instances and numbers $\varepsilon_N > 0$ such that*

$$\inf_{C\in\mathcal{C}_{N,N^k}} \mathrm{err}_{\mathcal{D}_N}(C) \geq \varepsilon_N. \tag{16.2}$$

Then SAT $\notin P/\mathrm{poly}$.

Proof. A circuit computing F_N on every input has error zero on $\mathcal{D}_N$, contradicting (16.2) for each selected length. □

The proposition isolates how random-SAT phenomena could legitimately enter a lower-bound proof. One would need to define $\mathcal{D}_N$ on formula descriptions, use cluster geometry to construct an observable label or test on those descriptions, and prove that all polynomial-size circuits incur nonzero error on that distribution. The geometry of satisfying assignments may be an ingredient, but it is not the result.

8.5 Hard distributions and minimax structure

A distributional lower bound can be related to randomized circuits by a finite minimax principle. Let $\mathcal{C}_{N,s}$ be the finite class of deterministic size-s circuits, and let F_N be a Boolean target. Consider the payoff

$$\ell(C,x) = \mathbf{1}_{C(x)\neq F_N(x)}. \tag{27.1}$$

If every distribution over circuits has expected error at least ε on some input distribution, then every randomized mixture of size-s circuits also fails with error at least ε on that distribution. For exact nonuniform computation, it suffices to prove positive error against deterministic circuits under one distribution.

Proposition 8.6 (Distributional exact-lower-bound criterion)**.** *If there exists a distribution $\mathcal{D}_N$ on encoded inputs such that*

$$\min_{C\in\mathcal{C}_{N,s}} \mathbb{E}_{x\sim\mathcal{D}_N}\ell(C,x) > 0, \tag{27.2}$$

then $F_N \notin \mathcal{C}_{N,s}$.

Proof. If an exact circuit existed, its error would be zero at every input and hence under $\mathcal{D}_N$. □

This provides a legitimate location for probability and phase-transition inputs: one may seek a distribution on formula encodings for which cluster-related observables make decision impossible for every small general circuit. However, the proof must compare circuits on the encoding strings. The assignment-space geometry of sampled formulae can be used only after it has been converted into an encoded-input indistinguishability or error theorem.

8.6 Hardness amplification is not automatic

Suppose one has shown that for a particular distribution on encoded formulas, every small circuit errs with probability at least ε_N, where ε_N may be tiny. This already excludes exact computation at that length when $\varepsilon_N > 0$. To obtain robust average-case separation, one may attempt amplification by taking products or direct-product encodings of independent instances. Yet the language operation and the circuit size must be tracked.

For example, define an AND-composition language slice on r instances by

$$F_N^{\wedge r}(e_1,\dots,e_r) = +1 \quad \Longleftrightarrow \quad F_N(e_i) = +1 \text{ for every } i. \tag{51.1}$$

This is an encoded conjunction of satisfiability instances. A circuit deciding F_N exactly gives a circuit deciding $F_N^{\wedge r}$ with only linear overhead in r, but hardness of the conjunction need not return hardness of a single instance without a reduction in the correct direction.

Proposition 8.7 (Safe amplification implication)**.** *If F_N is exactly computable by size s circuits, then $F_N^{\wedge r}$ is exactly computable by size $rs + O(r)$ circuits. Hence a lower bound against circuits of size $rs + O(r)$ for $F_N^{\wedge r}$ implies a lower bound against size-s circuits for F_N.*

Proof. Compute each instance by a copy of the F_N circuit and combine the outputs with an AND tree. The contrapositive proves the final assertion. □

An amplified distributional route remains legitimate if the composite target is encoded and the lower bound applies to general circuits at the enlarged input length. Cluster decomposition of witnesses may guide selection of composite instances, but it still must be accompanied by a circuit indistinguishability theorem on their encodings.

9 Gate-recursive separator machinery

9.1 Gate polynomials and convolution closure

The Boolean gate operations induce precise Fourier algebra. In sign convention, NOT sends f to $-f$. If truth is $+1$, AND is represented by (40.2), and OR by

$$u \vee v = \frac{1 + u + v - uv}{2}. \tag{45.1}$$

If a, b are coefficient vectors of u, v, write their symmetric-difference convolution as

$$(a * b)(S) = \sum_T a(T) b(S \triangle T). \tag{45.2}$$

Then the coefficient vector of AND is

$$\widehat{u \wedge v} = \frac{-e_\varnothing + a + b + a * b}{2}, \tag{45.3}$$

and similarly OR replaces $+a * b$ by $-a * b$ and the constant sign accordingly.

Proposition 9.1 (Exact spectral circuit recursion)**.** *The set of spectra of size-s De Morgan circuits is generated from input-character vectors and constants by at most s applications of sign change and the polynomial operations (45.3) and its OR analogue.*

Proof. Induct over a topological ordering of gates. Every input and constant has the given initial spectrum; each gate transforms the spectra of predecessors by the displayed operation. □

This recursion is a genuine algebraic encoding of circuit computation. It suggests an alternative construction: rather than a curvature graph based only on Fourier mass, define an outer relaxation of the recursively generated spectral set under convolution-polynomial gate operations. A proof that F_N is outside all polynomial-size stages of this recursion would establish the required lower bound. However, proving such exclusion is again exactly the lower-bound problem, now written in a faithful Fourier algebra.

9.2 An outer-cone relaxation of gate recursion

Let $\mathcal{K}_{N,0}$ contain the spectra of constants and input literals. For $s \geq 1$, let $\mathcal{K}_{N,s}$ be a closed outer set containing $\mathcal{K}_{N,s-1}$ and every vector produced from two vectors in $\mathcal{K}_{N,s-1}$ by the gate operations in (45.3), allowing convexification or semidefinite relaxation if desired. By induction,

$$\mathcal{S}_{N,s} \subseteq \mathcal{K}_{N,s}. \tag{46.1}$$

Theorem 9.2 (Sound gate-recursion separator)**.** *If, for each fixed k, there are infinitely many N such that*

$$\widehat{F}_N \notin \mathcal{K}_{N,N^k}, \tag{46.2}$$

then $P \neq NP$.

Proof. Containment (46.1) implies that no size-N^k circuit computes F_N on the selected lengths. Apply Theorem 1.1. □

The merit of this framework is soundness: it explicitly includes parity, conjunction and every circuit generated by the basis. Its difficulty is strength: a relaxation broad enough to be tractable may quickly include $\widehat{F}_N$. Establishing (46.2) with a tractable $\mathcal{K}$ would be an authentic new proof method.

A curvature or entropy functional could arise as a separating functional for $\mathcal{K}_{N,s}$ if its gate-closure inequalities are proved. For example, one would need inequalities showing that if a, b satisfy the separator ceiling at budgets s_1, s_2, then each gate-combination spectrum satisfies the ceiling at budget $s_1 + s_2 + 1$. No such gate-closure theorem follows from high-degree mass, because convolution creates and destroys levels in ways exhibited by simple gates.

9.3 Proof obligations for a gate-stable functional

Let $\Lambda_{N,s}$ be a candidate functional intended to vanish or remain bounded on all size-s circuit spectra. To certify it inductively through gate recursion, it is sufficient to prove:

$$\Lambda_{N,0}(e_\varnothing) \leq b_0, \quad \Lambda_{N,0}(e_{\{i\}}) \leq b_0; \tag{47.1}$$

$$\Lambda_{N,s+1}(-a) \leq b_{s+1} \quad \text{whenever } \Lambda_{N,s}(a) \leq b_s; \tag{47.2}$$

$$\Lambda_{N,s_1+s_2+1}\left(\frac{-e_\varnothing + a + b + a * b}{2}\right) \leq b_{s_1+s_2+1} \tag{47.3}$$

whenever a, b arise from spectra satisfying the preceding budgets, and the analogous OR inequality. If one then proves

$$\Lambda_{N,N^k}(\widehat{F}_N) > b_{N^k} \tag{47.4}$$

for infinitely many N, separation follows.

Proposition 9.3 (Inductive soundness)**.** *Under (47.1)–(47.3), every spectrum of a size-s circuit satisfies $\Lambda_{N,s}(\widehat{C}) \leq b_s$.*

Proof. Induct on circuit size using the exact gate-recursion proposition. Inputs and constants are handled by (47.1), NOT by (47.2), and binary gates by (47.3). □

This is the correct mathematical form of an invented spectral-curvature mechanism: one must prove gate-stable inequalities, not merely a flow on a graph reconstructed from the final spectrum. The obstacle is now transparent and local. A successful invariant must remain bounded under every circuit-building gate while taking a provably larger value on the SAT language slices.

9.4 A moment hierarchy for small-circuit spectra

A general-circuit separator may be formulated through moment inequalities satisfied by all small-circuit functions. For a finite list of test functions $\psi_1, \dots, \psi_m$ on $\{-1, +1\}^N$, set

$$\mathbf{m}(f) = \big(\langle f, \psi_1\rangle, \dots, \langle f, \psi_m\rangle\big). \tag{28.1}$$

Let

$$K^{(m)}_{N,s} = \operatorname{conv}\{\mathbf{m}(g) : g \in C_{N,s}\} \subset \mathbb{R}^m. \tag{28.2}$$

If $\mathbf{m}(F_N) \notin K^{(m)}_{N,s}$, a separating hyperplane gives an explicit linear test violated by the target but satisfied by every size-s circuit.

Theorem 9.4 (Finite moment separator). *If there exist test functions $\psi_1, \dots, \psi_m$ and coefficients $a_1, \dots, a_m$ such that*

$$\sum_{j=1}^{m} a_j \langle F_N, \psi_j\rangle > \sup_{g \in C_{N,s}} \sum_{j=1}^{m} a_j \langle g, \psi_j\rangle, \tag{28.3}$$

then F_N has no size-s circuit.

Proof. If a size-s circuit computed F_N, substituting it in the supremum would contradict the strict inequality. □

Fourier coefficients arise by choosing $\psi_j = \chi_{S_j}$. Curvature or entropy functionals may guide the selection of nonlinear tests, but the universal right-hand bound in (28.3) must still be proved. The parity no-go theorem says that tests measuring only total high-degree mass cannot supply such a bound for all circuits.

9.5 A separator hierarchy for prospective new machinery

The exact distance energy can be relaxed through a hierarchy of functionals. Let $\mathcal{K}^{(r)}_{N,s}$ be a tractable outer approximation to the convex hull of size-s circuit spectra, indexed by a complexity parameter r, with

$$\operatorname{conv}(\mathcal{S}_{N,s}) \subseteq \mathcal{K}^{(r+1)}_{N,s} \subseteq \mathcal{K}^{(r)}_{N,s}. \tag{17.1}$$

Define

$$\Delta^{(r)}_{N,s}(f) = \operatorname{dist}(\widehat{f}, \mathcal{K}^{(r)}_{N,s})^2. \tag{17.2}$$

Then

$$0 \le \Delta^{(r)}_{N,s}(f) \le \mathfrak{D}^{\mathrm{cvx}}_{s}(f) \le \mathfrak{D}_s(f). \tag{17.3}$$

Theorem 9.5 (Sound relaxation principle). *For any r, positivity of $\Delta^{(r)}_{N,s}(F_N)$ implies that F_N has no circuit of size s. Consequently, a hierarchy satisfying*

$$\forall k\ \exists r(k)\ \exists^{\infty} N : \Delta^{(r(k))}_{N,N^k}(F_N) > 0 \tag{17.4}$$

would prove $P \neq NP$.

Proof. Every size-s circuit spectrum belongs to $\mathcal{S}_{N,s}$ and hence to each outer approximation. If the distance to an outer approximation is positive, the target cannot equal the spectrum of such a circuit. Apply Theorem 1.1. □

This formulation permits genuinely new mathematics: the sets $\mathcal{K}^{(r)}_{N,s}$ could incorporate restriction identities, low-degree polynomial simulations, communication constraints, pseudorandomness tests, or geometric semidefinite inequalities. Yet the hierarchy remains sound only when the containment in (17.1) is proved for every polynomial-size circuit. A curvature or entropy inequality may serve as a separating functional for one level of the hierarchy after this containment has been verified.

9.6 A circuit-composition consistency law

Circuit classes are closed under composition. If $g : \{-1, +1\}^m \to \{-1, +1\}$ has a size-s_0 circuit and each $h_i : \{-1, +1\}^N \to \{-1, +1\}$ has a size-s_i circuit, then

$$x \longmapsto g(h_1(x), \dots, h_m(x)) \tag{38.1}$$

has a circuit of size at most $s_0 + \sum_i s_i$. A proposed outer approximation $K_{N,s}$ to small-circuit spectra must reflect this closure, or at least it must include all spectra produced by such compositions.

On the Fourier side, composition is nonlinear. If g is expanded as a multilinear polynomial on m signs and h_i are Boolean functions, then

$$g(h_1(x), \dots, h_m(x)) = \sum_{A\subseteq[m]} \widehat{g}(A) \prod_{i\in A} h_i(x). \tag{38.2}$$

The Fourier coefficients of a product are convolution sums:

$$\widehat{uv}(S) = \sum_{T\subseteq[N]} \widehat{u}(T)\widehat{v}(S\triangle T). \tag{38.3}$$

Thus a circuit-generated spectral set must be stable under repeated symmetric-difference convolution. Level-mass bounds and simple graph curvature generally do not record this algebraic closure.

Definition 9.6 (Composition-closed spectral outer family)**.** *An outer family $K_{N,s}$ is composition-closed if it contains all character vectors of literals and constants and, whenever it contains spectra of $h_1, \dots, h_m$ and a spectrum of a gate function g, it contains the convolution-composition spectrum in (38.2)–(38.3) with the associated size budget.*

A separator theorem based on a composition-closed outer family would be structurally suited to general circuits. Proving that encoded SAT spectra lie outside every polynomial-size level of such a family is a concrete but unresolved form of the desired new mathematics.

9.7 Semidefinite separators and curvature candidates

A quadratic spectral invariant has the form

$$\mathcal{A}_M(f) = \langle \widehat{f}, M\widehat{f}\rangle, \tag{37.1}$$

where M is a symmetric positive semidefinite matrix indexed by subsets of $[N]$. Noise energies and graph Laplacian energies are examples. If one could construct $M = M_{N,k}$ satisfying

$$\mathcal{A}_M(F_N) > b_{N,k} \quad \text{and} \quad \mathcal{A}_M(g) \le b_{N,k} ext for all g \in C_{N,N^k}, \tag{37.2}$$

then the separation follows. A curvature-weighted Laplacian may be viewed as an attempt to choose such an M adaptively.

Proposition 9.7 (Quadratic stress-test requirement)**.** *For any matrix M used in (37.2), one must verify*

$$\langle e_T, Me_T\rangle \le b_{N,k} \quad (T \subseteq [N]), \tag{37.3}$$

where e_T is the Fourier coefficient vector of the parity character χ_T. Further inequalities must hold for the coefficient vectors of AND, OR, majority, addressing, and all their size-N^k compositions.

Proof. Each coefficient vector in (37.3) is realized by a polynomial-size circuit. The universal ceiling in (37.2) applies to every one of them. □

A semidefinite outer approximation to small-circuit spectra could encode finitely many such constraints and perhaps more sophisticated composition constraints. Its separation from F_N would be a genuine proof object. The construction of an adequate outer approximation is precisely the nontrivial machinery required; choosing a graph energy without proving containment does not accomplish it.

9.8 A quadratic moment relaxation for gate recursion

The gate-recursive formulation can be made more explicit by lifting coefficient vectors to moment matrices. For a Fourier vector $a = (a_S)_{S \subseteq [N]}$, define the rank-one matrix

$$X(a) = \begin{pmatrix} 1 \\ a \end{pmatrix} \begin{pmatrix} 1 & a^{\mathsf{T}} \end{pmatrix} \succeq 0. \tag{54.1}$$

The Boolean constraint (43.3) becomes a family of linear identities in the quadratic entries of $X(a)$:

$$\sum_T X(a)_{T, S \triangle T} = \mathbf{1}_{S=\varnothing} \qquad (S \subseteq [N]). \tag{54.2}$$

Rank one is nonconvex; dropping rank gives a convex outer set containing every Boolean spectrum. Gate recursions can be further relaxed by imposing linear identities between parent and child moment blocks derived from (45.3).

For a circuit parse tree of size s, associate a Fourier vector a_v and a moment matrix X_v to each gate v. At an input gate, a_v is a character vector. At a NOT gate, $a_v = -a_u$. At an AND gate with predecessors u, w, equation (45.3) is quadratic in the coefficient vectors because of convolution. Introduce cross-moment variables

$$Y_{u,w}(T, U) = a_u(T) a_w(U). \tag{54.3}$$

Then every output coefficient is a linear expression in $a_u, a_w, Y_{u,w}$:

$$a_v(S) = -\mathbf{1}_{S=\varnothing} + a_u(S) + a_w(S) + \sum_T Y_{u,w}(T, S \triangle T)2. \tag{54.4}$$

A block positive-semidefinite constraint on $(1, a_u, a_w)$ provides a convex outer relaxation of the cross moments.

Proposition 9.8 (Soundness of the moment relaxation)**.** *Every Fourier vector produced by a size-s Boolean circuit admits moment and cross-moment variables satisfying the Boolean linear identities, the gate linear identities, and the associated positive-semidefinite block constraints.*

Proof. Given an actual circuit, set every vector and every cross moment equal to its genuine product values. Rank-one moment matrices are positive semidefinite, Boolean identities follow from $f^2 = 1$, and gate identities follow from the exact sign-polynomial formulas. Dropping the rank conditions enlarges the feasible set but cannot remove any actual circuit spectrum. □

This construction is a mathematically sound outer hierarchy. A separating dual certificate showing that $\widehat{F}_N$ violates it at all polynomial size budgets would prove the target lower bound. Its weakness is computational scale: the full Fourier vector has dimension 2^N, and a tractable truncation must retain enough gate information to exclude SAT while still containing all small circuits. Finding such a truncation is a concrete version of the required new mathematics.

9.9 Dual certificates for the moment relaxation

A semidefinite outer relaxation has a dual: a linear combination of Boolean and gate identities together with positive-semidefinite multiplier matrices. If the dual evaluates to a negative number on the target vector while remaining nonnegative on every feasible circuit construction, it is an explicit lower-bound certificate. In symbolic form, a certificate would assert

$$\mathfrak{L}_{N,k}(a, \{X_v, Y_{u,w}\}) \geq 0 \quad \text{for every feasible size-}N^k \text{ gate construction,} \tag{54.5}$$

while

$$\mathfrak{L}_{N,k}(\widehat{F}_N, \cdots) < 0. \tag{54.6}$$

Soundness follows from weak semidefinite duality; the difficulty is constructing the target-side evaluation without writing the entire SAT truth table.

The formulation also explains why a curvature statistic may be informative only after a gate-containment proof. Curvature can contribute terms to a dual functional, but feasibility must be established recursively for all gate-generated spectra. A terminal graph statistic unaccompanied by gate constraints has no reason to bound the full circuit class.

9.10 Uniform definability and the explicit target

For a nonuniform lower bound, the certificate may itself be nonuniform in N; to constitute a mathematical proof it must nevertheless be describable and verified for infinitely many lengths by a finite argument or a uniform theorem. An assertion that a certificate exists for each N without construction is simply a restatement of the lower bound.

Let $\mathrm{Cert}_{N,k}$ be a proposed family of separator certificates. The useful theorem has the form

$$\forall k\ \exists N_0\ \forall N \geq N_0 : \mathrm{Cert}_{N,k} \text{ is defined and satisfies (54.5)–(54.6).} \tag{54.7}$$

A weaker infinite-subsequence theorem is sufficient for exclusion from P/poly:

$$\forall k\ \exists^{\infty} N : \mathrm{Cert}_{N,k} \text{ is defined and satisfies (54.5)–(54.6).} \tag{54.8}$$

Either conclusion would be a decisive breakthrough. The spectral calculations in the present manuscript specify admissible certificate types and rule out the magnitude-tail form; they do not construct certificates satisfying (54.7) or (54.8).

9.11 Algebraic gate ideals as exact finite certificates

For each fixed input length N and size budget s, circuit nonexistence can be represented by a finite polynomial feasibility problem. This observation does not solve the asymptotic lower bound, but it supplies a completely faithful algebraic certificate space in which a future separator may be constructed.

Choose a finite complete gate basis. A labelled circuit of size s is specified by finitely many Boolean selector variables indicating, at each gate, its type and the predecessor wires it reads. For every input $x \in \{0,1\}^N$ and every gate position v, introduce a Boolean value variable $y_{v,x}$. Booleanity is imposed by

$$y_{v,x}^2 - y_{v,x} = 0. \tag{55.1}$$

For each choice of gate type and predecessors, polynomial selector constraints express the gate truth table. At the output gate o, impose

$$y_{o,x} = \mathbf{1}_{\mathrm{SAT}\cap\Sigma^N}(x) \qquad (x \in \{0,1\}^N). \tag{55.2}$$

Together with the selector constraints that choose one permitted wiring per gate, these equations have a common zero over $\{0,1\}$ if and only if a size-s circuit computes the length-N SAT slice.

Theorem 9.9 (Finite algebraic certificate equivalence)**.** *For fixed N and s, let $I_{N,s}^{\mathrm{SAT}}$ be the ideal generated by the Boolean, wiring, gate-consistency and output equations above. Then*

$$F_N \notin C_{N,s} \iff 1 \in I_{N,s}^{\mathrm{SAT}} \tag{55.3}$$

after adjoining Boolean equations for all selector variables. Equivalently, nonexistence of a size-s circuit admits a finite polynomial identity certificate.

Proof. The generators describe a finite system over the Boolean cube of selector and gate-value variables. A common Boolean zero is exactly a valid circuit agreeing with F_N on every input. If there is no common zero, the finite Boolean coordinate algebra is a product of copies of the ground field over all Boolean assignments; an ideal with empty zero set is the unit ideal. Conversely, if 1 belongs to the ideal, no common zero exists because every generator vanishes at any proposed circuit assignment while the constant polynomial does not. □

This theorem is an exact algebraic analogue of the circuit-distance energy. It incorporates gate composition rather than replacing it by level masses, and it automatically passes the parity test because parity circuits give feasible points for their own target truth tables. Its cost is enormous: it contains output equations for every input string and has no known uniform low-degree certificate for the SAT slices at polynomial size budgets.

A prospective "curvature-guided" or Fourier–entropy certificate could be made rigorous by mapping it into this ideal framework. It would be sufficient to construct, for infinitely many N and every k, polynomial identities witnessing

$$1 \in I^{\mathrm{SAT}}_{N,N^k} \tag{55.4}$$

whose existence is proved uniformly from a tractable structural theorem. Such identities would be unambiguous nonuniform lower-bound certificates. No high-tail estimate or assignment-cluster theorem presently supplies them, because the former fails on easy circuits and the latter is not an output-equation statement on encoded instances.

9.12 Relationship between spectral and algebraic certificates

The Fourier representation is an invertible linear transformation of a truth table. Therefore a polynomial identity in gate-value variables may be transported to Fourier coordinates, and a complete Fourier separator may be interpreted as a dual witness to circuit infeasibility. Conversely, a spectral inequality that is valid only after discarding gate-consistency constraints cannot certify membership of the unit ideal.

Let $\mathcal{P}_{N,s}$ be the finite Boolean feasibility polytope obtained by taking convex combinations of valid size-s circuit truth tables, and let $\widehat{\mathcal{P}}_{N,s}$ be its Fourier image. A linear spectral witness separating $\widehat{F}_N$ from $\widehat{\mathcal{P}}_{N,s}$ is a valid certificate against randomized mixtures of size-s circuits and hence against exact deterministic circuits. The moment relaxations of Section 9 provide outer approximations to this polytope. The gate ideal provides a nonlinear exact description. A new proof can use either language, but must establish target exclusion from an outer set that contains all valid small circuits.

This unifies the constructive alternatives. A gate-stable Fourier functional, a semidefinite moment certificate, a distributional correlation witness, or a low-degree algebraic refutation all become sound once their containment of every size-bounded circuit and their exclusion of F_N are proved. The unconditional separation claimed by the fixed front matter would require such certificates for unbounded input lengths and for every polynomial size exponent.

10 Barriers, explicitness, and proof discipline

10.1 Barrier statements and what an unconditional proof must show

A proof of $P \neq NP$ need not resemble earlier lower-bound arguments, but it cannot be certified by declaring that known barriers have been bypassed. Three distinctions are relevant.

A proof *relativizes* when it remains valid after both classes are equipped with the same oracle. Baker, Gill and Solovay produced oracles giving opposite answers to relativized versions of the question. Thus a relativizing argument cannot settle the unrelativized problem.

A proof property is *natural* in the Razborov–Rudich sense when it is constructive and large while separating a hard function from small circuits. Subject to standard pseudorandom-function assumptions, this class of properties cannot prove superpolynomial lower bounds for general circuits. A functional is not shown to avoid naturalness merely by being described geometrically; its largeness and constructivity properties must be analysed.

Algebrization extends the relativization barrier to arguments that admit low-degree algebraic access to the oracle. Again, a proposed invariant does not automatically avoid the barrier because it has no advertised finite-field analogue; one needs a proof that the entire argument is outside the algebrizing framework.

Proposition 10.1 (Barrier declarations are not lower bounds)**.** *Neither non-relativizing vocabulary, nor the absence of an obvious algebraic extension, nor an asserted sparsity of an invariant implies (1.2). A separation theorem requires an independently proved upper bound for every polynomial-size circuit family and a lower bound for the target language slices.*

Proof. Each barrier theorem constrains classes of proof methods; it does not furnish a lower bound for a new method. The implication (1.2) follows only from the two numerical inequalities displayed there. Descriptive properties of the proposed method do not logically imply either inequality. □

This observation does not obstruct invention of new mathematics. It specifies what new mathematics must deliver: a functional or structural theorem that passes the parity stress test, is attached to encoded language slices, is stable under NP reductions at polynomial cost, and is quantitatively bounded for all general circuits.

10.2 Relation to natural-proof barriers

A gate-stable separator might still confront the natural-proofs barrier. In a truth-table formulation, a property is useful against size-s circuits if it holds for a target hard function but fails for every size-s function; it is large if it holds for a substantial fraction of all truth tables and constructive if membership can be tested efficiently from the truth table. Under pseudorandom-function assumptions, large constructive useful properties cannot prove strong general-circuit lower bounds.

The exact distance $\mathfrak{D}_s$ is useful but not known to be constructive at the required scale; its evaluation directly searches a vast circuit class. An outer-cone separator may become constructive if described by a polynomial-size semidefinite program, but then its largeness must be analysed. Conversely, declaring a curvature statistic sparse does not establish non-naturalness: one must prove the density and constructivity properties in the Razborov–Rudich model.

Proposition 10.2 (Barrier-compatible statement)**.** *A proof of (46.2) or (47.4) would establish a circuit lower bound regardless of whether its method is subsequently classified as natural, non-natural, algebrizing or non-algebrizing. Conversely, asserting such a classification without proving (46.2) or (47.4) yields no lower bound.*

Proof. The first implication follows from sound gate containment and separation. The converse follows because methodological classification is not one of the hypotheses in Theorem 1.1. □

This places the emphasis correctly: the primary theorem is a separator with proven gate containment and SAT exclusion. Barrier discussion is secondary validation of the method, not a replacement for the separator.

10.3 Counting lower bounds and the explicitness barrier

There is a simple nonconstructive lower bound for general circuits. It is useful because it reveals how far an explicit language such as SAT lies from a counting argument. Fix the fan-in two De Morgan basis. A gate in a size-s circuit chooses its type from a constant set and chooses at most two predecessors from the N inputs and earlier gates. Thus the number of labelled size-s circuits is at most

$$\left(c(N+s)^2\right)^s = 2^{O(s\log(N+s))}. \tag{30.1}$$

Every such circuit defines one Boolean function on N inputs, whereas the number of Boolean functions is 2^{2^N}.

Theorem 10.3 (Shannon counting bound)**.** *If $s\log(N+s) = o(2^N)$, then the fraction of Boolean functions on N bits computable by size-s circuits tends to zero. In particular there exist functions requiring circuit size $\Omega(2^N/N)$.*

Proof. By (30.1), the number of size-s computable functions is at most $2^{O(s\log(N+s))}$. Divide by 2^{2^N}. When the exponent in the numerator is $o(2^N)$, this ratio tends to zero. Choosing $s = c2^N/N$ with sufficiently small absolute c yields the stated existence bound. □

This theorem does not identify an NP language with large circuit complexity. Indeed, almost all truth tables are not even succinctly describable, whereas a language in NP has a polynomial-time verifier and therefore an extremely compact description. The explicitness requirement is the heart of the lower-bound problem.

Proposition 10.4 (Why counting does not close the target)**.** *The existence of functions outside* P/poly *does not imply* SAT $\notin P/\text{poly}$*, nor* $P \neq NP$*.*

Proof. The class of all Boolean function families contains undecidable and arbitrarily specified functions. Counting shows that most such families have large circuits; it does not locate any NP-complete language among them. A proof concerning SAT must establish a lower bound for the explicit slices F_N. □

A spectral separator offers one possible route to explicitness: it can be computed or estimated from structural data of F_N rather than its full truth table. But this advantage disappears if the separator does not distinguish F_N from explicit small-circuit stress tests. The parity and encoding-domain theorems therefore sit between nonconstructive counting and an actual explicit lower bound.

10.4 Karp–Lipton consequences and their limitation

The target SAT $\notin P/\text{poly}$ is stronger than $P \neq NP$. The classical Karp–Lipton theorem states that if $NP \subseteq P/\text{poly}$, then the polynomial hierarchy collapses to its second level. This provides evidence for the strength of nonuniform lower bounds, but it is not a proof of them.

The logical direction needed here is simpler:

$$\text{SAT} \notin P/\text{poly} \quad \Longrightarrow \quad NP \not\subseteq P/\text{poly} \quad \Longrightarrow \quad P \neq NP. \tag{33.1}$$

The first implication holds since SAT $\in NP$; the second since $P \subseteq P/\text{poly}$. A Fourier or geometric programme targeting P/poly is therefore entitled to use (33.1) after it has actually proved a circuit lower bound for the language slice.

Proposition 10.5 (No converse from collapse evidence)**.** *The known consequence* $NP \subseteq P/\text{poly} \Rightarrow PH = \Sigma_2^P$ *does not supply* $NP \not\subseteq P/\text{poly}$ *unless one separately proves that the polynomial hierarchy does not collapse.*

Proof. It is a conditional implication. Its contrapositive requires the premise $PH \neq \Sigma_2^P$, which is itself unproved. □

This observation parallels the barrier analysis: consequences of a hypothetical small-circuit family may indicate why the hypothesis is implausible, but only a direct lower bound establishes the separation.

10.5 Adversarial checks required of a proposed separator

Any proposed Δ or $\mathcal{J}$ intended for general circuits should be tested against explicit classes. The following list is mathematical rather than procedural: each item represents a necessary inequality.

First, parity circuits require that a separator not infer hardness from high Fourier degree alone. Second, majority and threshold circuits require control of functions with strong low-degree but nontrivial long tails. Third, addressing and multiplexer functions test whether sparse structured supports are mistaken for hardness. Fourth, compositions of small circuits test whether the proposed invariant is stable under circuit substitution. Fifth, syntactic recodings of formulas test whether the target quantity is defined on a language rather than on an arbitrary presentation.

Definition 10.6 (Admissible general-circuit separator)**.** *A sequence of functionals* $\mathcal{J}_{N,k}$ *is admissible if it* *satisfies:*

$$\mathcal{J}_{N,k}(C) = 0 \qquad \textit{for every size-}N^k \textit{ circuit } C, \tag{18.1}$$

$$\mathcal{J}_{N,k}(f \circ \pi) = \mathcal{J}_{N,k}(f) \qquad \textit{for input rewiring symmetries } \pi, \tag{18.2}$$

$$\mathcal{J}_{N,k}(F_N) > 0 \qquad \textit{for infinitely many } N \textit{ for every } k. \tag{18.3}$$

Conditions (18.1)–(18.3) are deliberately strong: together they prove the lower bound. The original Fourier-tail and cluster-curvature proposals fail before reaching (18.1), because easy explicit circuits already occupy their purported forbidden spectral region.

11 The complete separation gate

11.1 A valid descent programme and its unresolved theorem

We may now formulate a mathematically coherent replacement for a curvature-guided descent claim. Fix k. A *sound descent functional* is a nonnegative functional $\mathcal{J}_{N,k}$ on Boolean functions such that

$$\mathcal{J}_{N,k}(g) = 0 \quad (g \in C_{N,N^k}), \qquad \mathcal{J}_{N,k}(f) > 0 \Longrightarrow \mathfrak{D}_{N^k}(f) > 0. \tag{13.1}$$

A spectral, geometric or curvature formula may be used for $\mathcal{J}$ only after (13.1) is proved. For example, the exact choice

$$\mathcal{J}_{N,k}(f) = \mathfrak{D}_{N^k}(f) \tag{13.2}$$

is sound. A useful new theory would replace (13.2) by a lower-bounding functional whose positivity on F_N is accessible without enumerating all circuits.

Proof obligation 11.1 (General-circuit obstruction theorem). *Construct functionals* $\mathcal{J}_{N,k}$ *and an infinite set* $\mathcal{N}_k$ *for each* k *such that*

$$\mathcal{J}_{N,k}(g) = 0 \quad \textit{for every } g \in C_{N,N^k}, \qquad \mathcal{J}_{N,k}(F_N) > 0 \quad (N \in \mathcal{N}_k). \tag{13.3}$$

The functionals must be defined on encoded language slices, remain sound in the presence of parity-like polynomial circuits, and admit a proof of positivity independent of the assumption $P \neq NP$.

Theorem 11.2 (Conditional completion through a sound functional). *If Obligation 11.1 is fulfilled, then* SAT $\notin P/\text{poly}$*, and hence* $P \neq NP$.

Proof. For each k and each $N \in \mathcal{N}_k$, positivity of $\mathcal{J}_{N,k}(F_N)$ implies $\mathfrak{D}_{N^k}(F_N) > 0$, so no size-$N^k$ circuit computes that slice. Theorem 1.1 gives SAT $\notin P/\text{poly}$. Since $P \subseteq P/\text{poly}$ and SAT is NP-complete, $P \neq NP$ follows. □

The theorem is exact. It does not convert a restricted lower bound into a general one, and it does not confuse random instance geometry with the encoded language function. It also explains why the unresolved step is hard: a proof of (13.3) is precisely a new nonuniform general-circuit lower bound for an NP-complete language.

11.2 Necessary properties of a new lower-bound machine

The preceding analysis yields an exact specification for any new mathematical machine aimed at the fixed conclusion.

First, it must be a theorem about a standard NP-complete language on encoded inputs. Secondly, it must control unrestricted polynomial-size circuits, not merely bounded-depth, monotone or algebraic subclasses. Thirdly, it must admit all explicit easy high-degree functions, particularly parity. Fourthly, any use of random formula geometry must be converted into a hard distribution or a separator on formula encodings. Fifthly, any continuous or geometric descent must either preserve Boolean circuit realizability or operate only as an inequality on original circuits. Finally, barrier analysis must accompany, not replace, the lower bound.

Proof obligation 11.3 (Separator construction in a tractable hierarchy)**.** *For each k, construct a tractable outer set K_{N,N^k} containing the spectra or tested moments of every size-N^k circuit and prove, on infinitely many lengths,*

$$\mathbf{m}(F_N) \notin K_{N,N^k}. \tag{29.1}$$

The containment theorem must pass the parity, conjunction, majority, addressing and composition stress tests; the exclusion theorem must apply to encoded SAT slices.

If this obligation is achieved, Theorem 28.1 and Theorem 1.1 complete the separation. If it is not achieved, no terminology of spectral saturation or curvature descent converts restricted or fibrewise facts into $P \neq NP$.

11.3 Final precise form of the missing separation theorem

Combining the rigorous statements above yields a single theorem whose proof would establish the conclusion named in the fixed front matter. Let F_N be the standard encoded SAT slice and let $K_{N,k}$ be a family of spectral or moment outer sets.

Proof obligation 11.4 (Complete spectral separator theorem)**.** *For every $k \in \mathbb{N}$, construct infinitely many lengths N and a closed set $K_{N,k} \subseteq \mathbb{R}^{2^N}$ satisfying:*

(a) $\widehat{g} \in K_{N,k}$ for every Boolean function g computable by a size-N^k circuit;

(b) membership containment in (a) is stable under gate composition and input recoding;

(c) $\widehat{F}_N \notin K_{N,k}$, certified by an explicit separating inequality or positive distance;

(d) the proof of (c) uses the encoded language slices and not merely the assignment geometry of individual random formulae.

Theorem 11.5 (Closure from the complete spectral separator theorem)**.** *If Obligation 39.1 is proved, then* SAT $\notin P/$poly *and $P \neq NP$.*

Proof. For each k, choose infinitely many N furnished by the obligation. If F_N had a size-N^k circuit, property (a) would place $\widehat{F}_N$ in $K_{N,k}$, contradicting (c). Thus Theorem 1.1 gives SAT $\notin P/$poly; equation (33.1) gives $P \neq NP$. □

The obligation expressly incorporates the corrections forced by the mathematics. It admits parity and other small circuits through containment, demands external encoded-language analysis rather than fibrewise cluster rhetoric, and provides the correct location for any future curvature, entropy, semidefinite or geometric invariant. It is a genuine path to the desired separation, but its exclusion clause is not proved by the existing high-tail argument.

11.4 A defect calculus for the spectral-separation route

It is helpful to express the gap of a proposed argument by nonnegative defect quantities rather than by informal closure language. Let F_N be the encoded SAT slice and fix a size budget $s = N^k$. Suppose a proposed proof uses a functional $\mathcal{S}_N$. Define three defects:

$$\delta_{\mathrm{circ}}(N,k) = \max\left\{0,\ \sup_{g \in C_{N,s}} \mathcal{S}_N(g) - B_{N,k}\right\}, \tag{50.1}$$

$$\delta_{\mathrm{tar}}(N,k) = \max\left\{0, A_{N,k} - \mathcal{S}_N(F_N)\right\}, \tag{50.2}$$

and

$$\delta_{\mathrm{sep}}(N,k) = \max\{0, B_{N,k} - A_{N,k}\}. \tag{50.3}$$

If all three defects vanish and $A_{N,k} > B_{N,k}$ has been arranged strictly, the functional separates the target from all size-s circuits. Conversely, any positive defect identifies a precise missing theorem or counterexample.

For the high-tail functional

$$\mathcal{S}_N(f) = \sum_{|S|>N^{\alpha}} \widehat{f}(S)^2 \tag{50.4}$$

and a purported exponentially small circuit ceiling $B_{N,k} < 1$, the parity theorem gives

$$\delta_{\mathrm{circ}}(N,k) \geq 1 - B_{N,k} > 0 \tag{50.5}$$

for all sufficiently large N. Thus the defect is not merely unproved: it is provably nonzero. No strengthening of the target lower bound can repair an already failed circuit ceiling.

For a cluster-derived target estimate the domain defect is separate. Let $\mathcal{S}_N^{\mathrm{int}}(\Phi)$ be a statistic on assignment-space functions and $\mathcal{S}_N^{\mathrm{ext}}(F_N)$ a statistic on the language slice. Define

$$\delta_{\mathrm{enc}}(N) = \inf\left\{D \geq 0 : \mathcal{S}_N^{\mathrm{ext}}(F_N) \geq \mathbb{E}_{\Phi}\mathcal{S}_N^{\mathrm{int}}(\Phi) - D \text{ is proved}\right\}. \tag{50.6}$$

Without an encoding-transfer theorem, the right inequality is unavailable, so the intended target estimate has an unresolved domain defect. The two defects are independent: even a successful transfer would not overcome parity's violation of the circuit ceiling.

Theorem 11.6 (Zero-defect separation criterion)**.** *If for every k there are infinitely many N and functionals $\mathcal{S}_{N,k}$ with strict thresholds $A_{N,k} > B_{N,k}$ for which the circuit, target and encoding defects vanish, then $P \neq NP$.*

Proof. Vanishing circuit and target defects gives

$$\mathcal{S}_{N,k}(F_N) \geq A_{N,k} > B_{N,k} \geq \sup_{g \in C_{N,N^k}} \mathcal{S}_{N,k}(g),$$

so no size-N^k circuit computes F_N. Apply Theorem 1.1. □

This defect calculus is a rigorous bookkeeping mechanism: it does not itself set any defect to zero. It makes clear which theorem has to be constructed and prevents a false inference from being hidden in a descriptive phrase.

11.5 Closing theorem dependency chain

The logically valid dependency chain for a prospective general-circuit proof can now be stated without ambiguity:

$$\begin{gathered}\text{encoded language slice } F_N \longrightarrow \text{admissible separator } \mathcal{J}_{N,k} \longrightarrow \mathcal{J}_{N,k}(F_N) > 0, \\ \mathcal{J}_{N,k}(C) = 0\ (C \in C_{N,N^k}) \longrightarrow F_N \notin C_{N,N^k} \longrightarrow \mathrm{SAT} \notin P/\mathrm{poly} \longrightarrow P \neq NP.\end{gathered} \tag{19.1}$$

The Fourier–Walsh calculus, restriction identities and spectral-distance formalism proved above justify the analytic language in which a separator may be sought. The parity theorem and encoding-domain theorem prove that the former candidate separator does not meet its universal circuit duties. Thus the precise remaining mathematical theorem is not hidden behind a ledger: it is the construction of an admissible separator with positive value on infinitely many encoded SAT slices for every polynomial circuit exponent.

11.6 From a language separator to the final complexity consequence

Assume, hypothetically, that a gate-stable functional Λ satisfying (47.1)–(47.4) has been obtained. Then for each k, infinitely many slices F_N are not computed by circuits of size N^k. The conclusion $\mathrm{SAT} \notin P/\mathrm{poly}$ follows. Since every language in P has a uniform polynomial-time machine and hence a polynomial-size circuit family obtained by unrolling its computation at each input length, one has $P \subseteq P/\mathrm{poly}$. Since $\mathrm{SAT} \in NP$, the equality $P = NP$ would imply $\mathrm{SAT} \in P \subseteq P/\mathrm{poly}$, a contradiction.

Theorem 11.7 (Final implication from a gate-stable separator)**.** *If a functional family* $\Lambda_{N,s}$ *satisfies the circuit-generation ceiling (47.1)–(47.3) and the SAT separation (47.4) for every fixed polynomial exponent on infinitely many input lengths, then* $P \neq NP$.

Proof. Inductive soundness excludes size-N^k circuits for each selected SAT slice. The exact general-circuit proof obligation in Theorem 1.1 yields SAT $\notin P/\text{poly}$. The standard containment $P \subseteq P/\text{poly}$ completes the implication. □

The theorem is a complete and rigorous terminal implication. Its hypotheses are exactly where new mathematics is needed. Fourier entropy, geometric curvature, restriction averages or phase-transition observables may enter the construction of Λ, but only after their gate stability and language-slice separation have been proved. The fixed published abstract asserts that these hypotheses have been supplied; the mathematical analysis in the present body shows that the former high-tail and cluster-transfer arguments do not do so.

11.7 Why a fixed abstract cannot replace the missing theorem

The mathematical implication needed for the published conclusion is the separator theorem: for every polynomial circuit exponent, infinitely many SAT slices must be excluded from that size class. The Fourier calculus and the exact defect framework above provide a clean formal language for this objective. They also prove two negative facts about the prior route: the universal high-tail ceiling is false for general circuits, and random formula clustering is not, by itself, a Fourier statement about encoded SAT.

A genuine completion would require one of the following forms of new mathematics:

(i) a gate-stable functional obeying inductive inequalities through all polynomial-size circuit constructions and separating encoded SAT slices;

(ii) a tractable outer hierarchy containing all size-bounded circuit spectra but excluding encoded SAT spectra at infinitely many lengths;

(iii) a hard encoded-instance distribution together with a proved positive error bound against every polynomial-size circuit;

(iv) another method yielding the exact nonuniform lower bound without relying on the falsified spectral ceiling.

Every route ends at the same quantified conclusion (1.1). The present analysis advances the manuscript by removing invalid implications and by expressing the necessary new theorem exactly. It cannot legitimately write the conclusion as proved until one of these constructions is completed and verified.

12 Conclusions

Fourier analysis on the Boolean cube is exact and powerful. It supplies orthogonal decompositions, noise energies, influences, and sharp lower bounds for restricted circuit classes. Random restrictions and clustering theorems are likewise substantive mathematics in their appropriate settings. The question is whether these tools, as assembled, prove a lower bound against unrestricted polynomial-size circuits deciding the encoded satisfiability language.

The stress tests above give a definitive answer for the high-tail route. A universal Fourier-tail ceiling for polynomial-size circuits is false because parity is computed by a linear-size circuit and has all of its Fourier mass at the highest degree. Cluster geometry for the satisfying assignments of a random instance does not determine the Fourier spectrum of the language slice on formula encodings. Declaring the resulting functional non-relativizing, non-algebrizing or non-natural supplies no missing numerical separation.

The sound terminal object is the circuit-distance energy $\mathfrak{D}_s$, together with any rigorously proved lower-bounding relaxation that remains positive on infinitely many satisfiability slices and vanishes for every size-s

circuit. This formulation gives a precise research route for genuinely new machinery: it must detect a feature of SAT that is absent from parity-like and all other polynomial-size circuits, it must be stable under the encoding and reduction maps, and its positivity must be proved without already assuming the desired lower bound.

Accordingly, the Fourier–entropy programme has a rigorous analytic foundation and a rigorous corrected endpoint. The unconditional general separation asserted by the retained published front matter would require the new obstruction theorem in Obligation 11.1; that theorem is not obtained by the high-tail, cluster-transfer, or curvature-descent assertions considered here.

12.1 Terminal verification of the separator obligation

The final lower-bound obligation may be stated without any reference to presentation choices. For each input length N and size exponent k, let F_N be the sign-valued characteristic function of encoded satisfiable formulae and let $\mathcal{S}_{N,N^k}$ denote the set of Fourier vectors of all Boolean circuits of size at most N^k. The exact obstruction is

$$\delta_{N,k} = \operatorname{dist}\big(\widehat{F}_N, \mathcal{S}_{N,N^k}\big)^2. \tag{56.1}$$

By Parseval and discreteness of Boolean truth tables,

$$\delta_{N,k} = 0 \quad\Longleftrightarrow\quad F_N \text{ has a size-}N^k \text{ circuit}, \qquad \delta_{N,k} > 0 \quad\Longrightarrow\quad \delta_{N,k} \geq 4\,2^{-N}. \tag{56.2}$$

Thus a complete separation needs an explicit method proving $\delta_{N,k} > 0$ for infinitely many N, for each fixed k. Every spectral, entropic, geometric or curvature construction is valid only insofar as it gives a lower bound for this distance while retaining the complete gate-generated feasible set.

The exact calculations of this article enforce three non-negotiable checks on any such lower bound. The first is the gate check: constants, literals, parity trees, conjunctions, disjunctions, threshold circuits and their compositions must remain inside the feasible circuit region. The second is the domain check: the object to be separated is F_N, a Boolean function of encoded formula descriptions, not the assignment indicator attached to one sampled formula. The third is the soundness check: a flow or relaxation may be used only when every circuit spectrum remains feasible under the bounding inequalities, or when a dual certificate directly separates the original target from all original circuits.

These checks turn the desired construction into a finite but scale-dependent inequality problem. In a gate-recursive outer model $K_{N,k}$, one must prove

$$\mathcal{S}_{N,N^k} \subseteq K_{N,k}, \qquad \widehat{F}_N \notin K_{N,k} \quad \text{for infinitely many } N \text{ and each } k. \tag{56.3}$$

The first inclusion is the circuit-ceiling theorem; the second exclusion is the SAT lower-bound theorem. A separator functional $L_{N,k}$ would verify the second statement by

$$L_{N,k}(\widehat{F}_N) > \sup_{u \in K_{N,k}} L_{N,k}(u). \tag{56.4}$$

Equations (56.3)–(56.4) are compatible with every elementary circuit stress test and with composition. They also show why the previously asserted high-degree Fourier ceiling cannot be repaired by a constant adjustment: parity lies in $\mathcal{S}_{N,N^k}$ and has unit mass at the highest Fourier level.

The contribution of the corrected machinery is therefore exact. It proves the Fourier identities on which any spectral separator must rest, gives a faithful gate-recursive and algebraic-certificate formulation of the general-circuit feasible set, proves the failure of the earlier magnitude-tail shortcut, and identifies a sound dual certificate that would imply the desired nonuniform lower bound. An unconditional conclusion requires an actual family of certificates satisfying (56.4) at unbounded lengths. Without that final construction, the fixed conclusion in the published front matter is not established by the internal argument.

12.2 Exact finite-length evaluation and the unbounded duty

For completeness, every finite slice of the obstruction can be decided by exact arithmetic. For a Boolean truth table $f : \{-1,+1\}^N \to \{-1,+1\}$ write the integer Walsh numerator

$$A_f(S) = \sum_{x\in\{-1,+1\}^N} f(x)\chi_S(x), \qquad \widehat{f}(S) = 2^{-N} A_f(S). \tag{57.1}$$

Given a fixed size budget s, enumerate the finitely many labelled circuits of at most s gates, evaluate each truth table, compute its integer Walsh numerators, and test equality with A_f. Equivalently, compute

$$2^{2N}\mathfrak{D}_s(f) = \min_{g\in C_{N,s}} \sum_{S\subseteq[N]} \left(A_f(S) - A_g(S)\right)^2, \tag{57.2}$$

an integer which is zero exactly when a size-s circuit computes f.

Proposition 12.1 (Finite verification is exact but not asymptotic separation)**.** *For every fixed N and s, the quantity $\mathfrak{D}_s(F_N)$ is decidable by a finite exact calculation. A list of successful calculations at finitely many input lengths cannot by itself imply* SAT $\notin P/$poly.

Proof. Finiteness follows from enumeration of circuits and truth tables, and equation (57.2) uses only integers. The lower-bound condition for $P/$poly quantifies over every polynomial exponent and infinitely many input lengths. Any finite calculation leaves all sufficiently large lengths undecided and therefore does not imply Theorem 1.1. □

This proposition marks the precise difference between verification and proof in the nonuniform setting. Computation can test a candidate separator or a proposed inductive gate inequality at selected sizes; a theorem must establish a uniform law producing separators at unbounded lengths. The gate ideal formulation, the moment outer hierarchy, and the dual witness inequalities are acceptable frameworks because each specifies a finite object at every length together with a soundness theorem. They become a solution only after a construction proves target exclusion for the quantified family of lengths and size budgets.

Accordingly, the final mathematical duty is neither an unexplained barrier nor a matter of typography. It is the production of a uniformly justified family of exact certificates, in any of the equivalent sound languages developed above, that separates the encoded SAT slices from every polynomial-size gate-recursive feasible region. Until that duty is discharged, a finite spectral calculation or a restricted-circuit estimate cannot support the universal separation.

12.3 Restriction stability of exact circuit distance

Random restrictions are useful only after their relation to the original encoded language has been stated on the correct domain. Let $\mathcal{R}_{N,m}$ be the uniform family of restrictions fixing $N-m$ encoding variables and leaving m variables free. For a Boolean function f write f_ρ for its restricted function, and define $\mathfrak{D}_s^{(m)}(f_\rho)$ by the same Fourier-distance minimum on m variables. Restricting a circuit never increases its number of gates, so every size-s circuit for f yields a size-s circuit candidate for every f_ρ.

Lemma 12.2 (Averaged restricted distance is a sound lower bound)**.** *For every $f : \{-1,+1\}^N \to \{-1,+1\}$, every $0 \le m \le N$, and every size bound s,*

$$\mathbb{E}_{\rho\in\mathcal{R}_{N,m}} \mathfrak{D}_s^{(m)}(f_\rho) \le \mathfrak{D}_s(f). \tag{57.3}$$

Consequently, any positive lower bound for the left side is a valid lower bound against size-s circuits for f itself.

Proof. Choose a circuit function g attaining the minimum in $\mathfrak{D}_s(f)$. For every ρ, the restriction g_ρ is computed by a circuit of size at most s, and hence

$$\mathfrak{D}_s^{(m)}(f_\rho) \le \|f_\rho - g_\rho\|_2^2.$$

Averaging over uniformly chosen restrictions and their uniformly chosen free inputs samples every $x \in \{-1,+1\}^N$ uniformly. Fubini's theorem therefore gives

$$\mathbb{E}_\rho \|f_\rho - g_\rho\|_2^2 = \|f - g\|_2^2 = \mathfrak{D}_s(f),$$

which proves the inequality. □

This lemma replaces an informal transfer from random formula geometry by a sound operation on encoded language slices. It does not assert that the restricted distances are positive; rather, it identifies a legitimate place where a newly constructed restriction theorem could enter a complete lower-bound proof without changing the domain of the Boolean function or losing the general-circuit quantifier.

12.4 Dual certificates in the gate-ideal model

The preceding distance formulation can be expressed by a certificate that is finite at each input length and remains faithful to circuit composition. Fix N and s, introduce variables for the value of each allowed gate on each input $x \in \{-1,+1\}^N$, and let $I_{N,s}$ be the ideal generated by the Boolean relations and the local gate equations. A labelled circuit of size at most s determines a real point of the variety of $I_{N,s}$; conversely, after fixing a gate wiring and output gate, every Boolean point of the corresponding component is its truth table. The union over the finitely many wirings is therefore exactly the size-s circuit family.

Proposition 12.3 (Sound dual separation criterion)**.** *Let F_N be the encoded satisfiability slice and suppose that, for a fixed size bound s, there are real numbers c_x and, for every labelled wiring ω with at most s gates, polynomials $q_{\omega,j}$ and sums of squares $\sigma_{\omega,r}$ such that*

$$\sum_{x\in\{-1,+1\}^N} c_x\big(F_N(x) - z_{\omega,\mathrm{out}}(x)\big) - 1 = \sum_j q_{\omega,j} G_{\omega,j} + \sum_r \sigma_{\omega,r}^2 \tag{58.1}$$

on the Boolean gate variables, where $G_{\omega,j} = 0$ are the gate and Boolean equations for ω. Then no circuit of size at most s computes F_N.

Proof. If a circuit with wiring ω computed F_N, substitution of its Boolean gate values would set every $G_{\omega,j}$ equal to zero and would make the linear expression on the left of (58.1) equal to -1. The right side would be a sum of real squares and hence nonnegative, a contradiction. □

The proposition is elementary, but it fixes the standard that a general-circuit lower-bound construction must meet. A separator is not merely a numerical statistic of a chosen Fourier profile: it must separate the target truth table from every gate-realizable component and it must do so by an identity that can be checked without presupposing the lower bound. If for every fixed k such certificates were constructed for $s = N^k$ at infinitely many N, then SAT $\notin P/\mathrm{poly}$ would follow immediately. This exact algebraic duty is compatible with Fourier coordinates, because Walsh transformation is an invertible linear change of the output variables. It is also compatible with semidefinite relaxations, provided the resulting dual identity has the form (58.1). The construction of an unbounded certificate family, rather than a restriction to a circuit subclass, is the unresolved theorem required by the fixed published claim.

12.5 Family-level consequence of exact dual separation

For $k \geq 1$, let $\mathcal{W}_{N,k}$ be the finite set of labelled Boolean wirings with at most N^k gates, and let F_N denote the satisfiability language slice in the fixed encoding of Section 1. The gate-ideal formulation acquires its full nonuniform meaning through the following direct consequence.

Proposition 12.4 (Unbounded dual-family criterion)**.** *If, for each $k \geq 1$, there are infinitely many lengths N such that every wiring $\omega \in \mathcal{W}_{N,k}$ admits a dual identity of the form* (58.1) *separating its output from F_N, then* SAT $\notin P/\mathrm{poly}$ *and hence $P \neq NP$.*

Proof. Fix k. At each certified length, Proposition 12.3 excludes every output computed by a circuit of size at most N^k, and hence

$$F_N \notin C_{N,k} \qquad \text{for infinitely many } N. \tag{58.2}$$

Since this is true for every k, Theorem 1.1 gives SAT $\notin P/\text{poly}$. The inclusion $P \subseteq P/\text{poly}$ and the NP-completeness of satisfiability then imply $P \neq NP$. □

Thus any Fourier completion must yield exact gate-ideal identities for an unbounded family of general-circuit wirings.

Corollary 12.5 (Required form of a Fourier completion)**.** *A Fourier-coordinate proof implies the stated separation only if, after Walsh inversion, it produces the dual-family criterion of Proposition 12.4.*

Proof. Walsh inversion is an invertible linear change of truth-table coordinates; apply Proposition 12.4 and Theorem 1.1. □

References

[1] S. A. Cook, The complexity of theorem-proving procedures, Proc. 3rd ACM Symp. Theory of Computing (STOC 1971), 151–158.

[2] D. Bhattacharjee, P. Nandi, S. N. Thakur, O. Frederick, and S. Ghosh, Almost Impossible Calabi–Yau Manifolds: Hodge Realization, Full-Measure SYZ Lifting, and Dimensional Saturation, PhilArchive ID: BHAAIC, 2026.

[3] R. M. Karp, Reducibility among combinatorial problems, in R. E. Miller and J. W. Thatcher (eds.), Complexity of Computer Computations, Plenum, 1972, 85–103.

[4] A. Bonami, *Étude des coefficients de Fourier des fonctions de* $L^p(G)$, Ann. Inst. Fourier (Grenoble) 20(2) (1970), 335–402.

[5] D. Bhattacharjee, Higher-Dimensional Calabi–Yau Manifolds and Dimensional Saturation, PhilArchive ID: BHAHCM, 2026.

[6] L. Gross, Logarithmic Sobolev inequalities, Amer. J. Math. 97(4) (1975), 1061–1083.

[7] W. Beckner, Inequalities in Fourier analysis, Ann. of Math. 102(1) (1975), 159–182.

[8] D. Bhattacharjee and O. Frederick, Hopf-Like Fibrations on Calabi–Yau Manifolds, Preprints 2025, DOI: 10.20944/preprints202504.2581.v4.

[9] M. Furst, J. B. Saxe, and M. Sipser, Parity, circuits, and the polynomial-time hierarchy, Math. Systems Theory 17(1) (1984), 13–27.

[10] J. Håstad, Almost optimal lower bounds for small depth circuits, Proc. 18th ACM Symp. Theory of Computing (STOC 1986), 6–20.

[11] D. Bhattacharjee, Homotopy Groups of Spheres, Hopf Fibrations and Villarceau Circles II: Advances in Fiber Bundle Topology and Spherical Homotopy Theory, Preprints 2026, DOI: 10.20944/preprints202602.2038.v1.

[12] A. A. Razborov, Lower bounds for the monotone complexity of some Boolean functions, Soviet Math. Dokl. 31 (1985), 354–357.

[13] A. A. Razborov, Lower bounds for the size of circuits of bounded depth with basis $\{\wedge, \oplus\}$, Math. Notes 41(4) (1987), 333–338.

[14] R. Smolensky, Algebraic methods in the theory of lower bounds for Boolean circuit complexity, Proc. 19th ACM Symp. Theory of Computing (STOC 1987), 77–82.

[15] D. Bhattacharjee and P. Nandi, Constructing Exotic Calabi–Yau 3-Folds via Quantum Inner State Manifolds, Preprints 2025, DOI: 10.20944/preprints202505.0700.v1.

[16] J. Kahn, G. Kalai, and N. Linial, The influence of variables on Boolean functions, Proc. 29th IEEE Symp. Foundations of Computer Science (FOCS 1988), 68–80.

[17] P. Erdős and R. Rado, Intersection theorems for systems of sets, J. London Math. Soc. 35 (1960), 85–90.

[18] N. Linial, Y. Mansour, and N. Nisan, Constant depth circuits, Fourier transform, and learnability, J. ACM 40(3) (1993), 607–620.

[19] D. Bhattacharjee, String Vibrations and Particle Families: A Resonance Classification Framework in String Phenomenology, Preprints 2026, 2026030792, DOI: 10.20944/preprints202603.0792.v1.

[20] A. A. Razborov and S. Rudich, Natural proofs, J. Comput. System Sci. 55(1) (1997), 24–35.

[21] R. M. Karp and R. J. Lipton, Turing machines that take advice, Enseign. Math. (2) 28 (1982), 191–209.

[22] D. Bhattacharjee, Block-Twist Field Geometry for UAP-Class Propulsion: Reconfiguration Metrics and Space-Route Compression Phenomenology, PhilArchive ID: BHABFG, 2026.

[23] T. Baker, J. Gill, and R. Solovay, Relativizations of the $P = ?NP$ question, SIAM J. Comput. 4(4) (1975), 431–442.

[24] S. Aaronson and A. Wigderson, Algebrization: a new barrier in complexity theory, ACM Trans. Comput. Theory 1(1) (2009), Article 2.

[25] D. Bhattacharjee, P. Nandi, S. Ghosh, and O. Frederick, Topology Kills the Ghost: Brane-Cluster UV Completion of Quantum Gravity, PhilArchive ID: BHATKT, 2026.

[26] Y. Ollivier, Ricci curvature of Markov chains on metric spaces, J. Funct. Anal. 256(3) (2009), 810–864.

[27] J. Ding, A. Sly, and N. Sun, Proof of the satisfiability conjecture for large k, Proc. 47th ACM Symp. Theory of Computing (STOC 2015), 59–68.

[28] D. Bhattacharjee, P. Nandi, O. Frederick, P. Samal, and S. Bhattacharjee, Higher Topos-Theoretic Closure of the Riemann Hypothesis, PhilArchive ID: BHAHTC, 2026.

[29] D. Achlioptas and A. Coja-Oghlan, Algorithmic barriers from phase transitions, Proc. 49th IEEE Symp. Foundations of Computer Science (FOCS 2008), 793–802.

[30] M. Mézard and A. Montanari, Information, Physics, and Computation, Oxford University Press, 2009.

[31] R. Williams, Non-uniform ACC circuit lower bounds, J. ACM 61(1) (2014), Article 2.

[32] P. Beame, A switching lemma primer, Technical Report UW-CSE-95-07-01, University of Washington, 1994.

[33] L. G. Valiant, Graph-theoretic arguments in low-level complexity, in Proc. 6th Symp. Mathematical Foundations of Computer Science, Lecture Notes in Comput. Sci. 53, Springer, 1977, 162–176.

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