The Polygon Graph and the Conservation of Exponential Cost

A Geometric Inquiry into the Boundary between Tractable and Intractable Computation

Toshiki Takahashi
Takahashi Mathematics Laboratory, Independent Researcher

June 2026

Abstract

The Polygon Graph is a geometric computational model in which a network is described by four quantities—nodes, edges, ridges, and areas—with ridges and areas determined from nodes and edges by exact equalities. This paper develops the model far enough to ask what it can say about the P versus NP question, and answers that question with care. We first exhibit the model's expressive reach: its area integral coincides with the attention operation of the Transformer, and extending ridges to complex-valued waves makes area paths formally analogous to path integrals. We then construct an explicit reduction of SAT to a Polygon Graph via the Menger Sieve and a binary tree of area paths, and we analyze the tempting conjecture that simultaneity collapses the search to polynomial time. We show this conjecture is false, and—more importantly—that its failure exemplifies a general principle we call the *conservation of exponential cost*: across ridge-volume integration, line-graph duality, and surprisal-driven sparsification, the exponential cost is relocated but never removed, exactly as the polynomial-time invariance of complexity classes requires. We are explicit that this does not establish P $\neq$ NP. The model's genuine and defensible contribution is geometric: a language for characterizing which problems are tractable, with the even-degree (Eulerian) criterion as a concrete first theorem and an analogy to integrable systems as a guide.

1. Introduction

Computational complexity asks which problems can be solved efficiently. Its central open question—whether P = NP—has resisted every direct assault from within the Turing-machine framework. This paper approaches the question from an unusual direction: a geometric model, the *Polygon Graph*, in which computation is recast as integration over a network of nodes and edges.

Our aim is neither to sensationalize nor to overclaim. We will show that the Polygon Graph has real expressive reach—deep enough to reproduce the attention mechanism of modern language models and to mirror the structure of quantum path integrals. We will then push it toward the P versus NP question, construct the most natural attempt at a positive resolution, and examine that attempt honestly. The attempt fails, for a reason that turns out to be general and instructive. From the wreckage we extract a principle—the conservation of exponential cost—and a redirection: away from settling P versus NP, and toward the geometric characterization of tractability, where the model has genuine and durable things to say.

The structure is as follows. Section 2 defines the model. Section 3 develops its reach (Transformer correspondence, complex-wave extension). Section 4 constructs the SAT reduction and states the simultaneity conjecture. Section 5 refutes it and generalizes the refutation. Section 6 delimits what is *not* thereby proved. Section 7 presents the model's positive contribution.

2. The Polygon Graph

2.1. Definition

Definition 2.1 (Polygon Graph)**.** A *Polygon Graph PG* consists of four elements:

- **Node** v_i: a point carrying a value;
- **Edge** e_{ij}: the interval between adjacent nodes i and j;
- **Ridge** r_{ij}: the rate of change across an edge;
- **Area** a_{ij}: the trapezoidal area beneath the ridge.

Ridges and areas are fixed by nodes and edges through exact equalities:

$$r_{ij} = \frac{v_j - v_i}{e_{ij}}, \qquad a_{ij} = \frac{v_i + v_j}{2}\, e_{ij}.$$

The defining feature is that ridges and areas are not separate data to be computed by a procedure; they are determined, exactly and without error, by the nodes and edges. A path through the graph accumulates a sequence of areas, and the total *area path* is the discrete integral of the node values along that path.

2.2. The Menger Sieve

Definition 2.2 (Menger Sieve)**.** The *Menger Sieve* is the matrix M with nodes indexing rows and edges indexing columns, with entries

$$M_{ij} = v_i \cdot e_{ij}.$$

Constraints enter as filters: forbidden edges are removed; required edges are fixed; admissible paths are read off by traversing the matrix.

The Menger Sieve recasts a search over a graph as a traversal over a table. Row traversal corresponds to assigning values; column traversal corresponds to testing constraints. This is structurally the same operation as verifying a candidate solution to a combinatorial problem—a correspondence we exploit, and scrutinize, in Sections 4 and 5.

3. The Reach of the Model

Before turning to complexity, we establish that the Polygon Graph is not a toy: it reproduces, exactly, a central operation of modern machine learning, and it extends naturally into the language of wave physics.

3.1. Coincidence with the Transformer's Attention

Place a value (a "height") on each node and connect adjacent heights by ridges. The area beneath the ridge—the trapezoidal integral—is the piecewise-linear accumulation of those heights. This is precisely the operation that the attention mechanism of a Transformer performs when it forms a weighted combination of token values: a discrete integral over a sequence, with the ridge playing the role of the local interpolation between attended positions. Stacking layers corresponds to iterated integration (raising the order of the integral); residual connections correspond to the discrete fundamental theorem of calculus, accumulating the integral across layers.

> **Remark.** This correspondence is independent of any complexity claim. It says that the Transformer's attention is a special case of integration on a Polygon Graph—a genuine and clarifying reformulation, useful regardless of what follows.

3.2. Ridges as Complex-Valued Waves

Let the edge length shrink, $e_{ij} \to \delta e \to 0$. The ridge becomes a derivative; under full continuization it becomes a wave, which we permit to take complex values:

$$r(e) = \frac{dv}{de} \in \mathbb{C}.$$

The real part records the direction of change; the imaginary part records a phase. Node values are recovered by integration, so no information is lost:

$$v_k = \lim_{\delta e \to 0} \int r(e)\, de.$$

With complex ridges, an area path becomes a complex integral, and a family of area paths combines as

$$\Psi = \sum_{\text{Path}} a_{\text{Path}}\, e^{i\phi_{\text{Path}}},$$

the same form as a sum over histories in a path integral. Paths whose phases align reinforce; paths whose phases oppose ($\phi_1 \approx \phi_2 + \pi$) cancel. This *interference* is the geometric content the model shares with wave physics, and it is what makes the model expressive enough to be worth testing against hard problems.

4. Encoding SAT on a Polygon Graph

We now bring the model to bear on an NP-complete problem and construct the most natural attempt at an efficient solution.

4.1. The Reduction

Let $\varphi = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ be a SAT instance in n variables and m clauses. Encode it as follows:

- each variable x_k becomes a node $v_k \in \{0, 1\}$;
- each clause C_j becomes an edge constraint e_{ij};

- satisfiability conditions become filters in the Menger Sieve;
- the assignment search becomes a traversal of a binary tree of area paths, branching on each variable's truth value.

The encoding itself is cheap: building the graph and the sieve costs $O(n+m) = O(n^2)$. The cost of the problem, if any, lies in the traversal.

4.2. The Simultaneity Conjecture

Here is the tempting idea. When two edges emanate from one node—the two truth values of a variable—two area paths are generated together:

$$a_{i1} = \frac{v_i + v_1}{2}\, e_{i1}, \qquad a_{i2} = \frac{v_i + v_2}{2}\, e_{i2}.$$

If branching always produces both successors "at once," one is led to conjecture that the entire tree is processed simultaneously, and hence that the exponential search collapses:

$$O(2^n) \xrightarrow{?} O(n).$$

Combined with interference (Section 3.2)—which would let unsatisfying assignments cancel and satisfying ones survive—this suggests an $O(n^3)$ solution to SAT, and therefore $\mathrm{P} = \mathrm{NP}$.

This conjecture is false. Establishing why, precisely, is the heart of the paper.

5. Why the Collapse Fails: Conservation of Exponential Cost

5.1. The Branch Points Are Exponentially Many

The flaw in the simultaneity conjecture is a counting error. Processing at each branch point is $O(1)$; but the number of branch points is itself exponential, and the per-level counts must be summed:

Depth k	Nodes at this level	Cumulative work
1	2^0	$O(1)$
2	2^1	$O(2)$
k	2^{k-1}	$O(2^k)$
n	2^{n-1}	$O(2^n)$

The deepest level alone holds 2^{n-1} nodes; processing each at $O(1)$ still costs $O(2^n)$ in total. "Simultaneous" does not abolish the work—it relocates it from the time axis to the hardware axis. Performing 2^n operations at once requires 2^n physical units. The exponent moves; it does not vanish.

Diagnosis. Parallelism trades time for space. It cannot, by itself, reduce total work below the number of operations that must occur. Asserting simultaneity is not eliminating the operations.

5.2. Three Relocations

One might try to rescue the collapse. Three natural moves all fail in the same way, and the pattern is the point.

Ridge volume. Replace path-counting by a single scalar: the volume under the ridge net over the base network, $V = \int \cdots \int r(e_1, \ldots, e_n)\, de_1 \cdots de_n$. But this is an n-fold integral; an m-point grid per axis needs m^n evaluations (the curse of dimensionality), and exact volumes of this kind are #P-hard (Valiant, 1979)—no easier than, and generally harder than, finding one solution. The exponent moves from *enumeration* to *integration dimension*.

Line-graph duality. Exchange nodes and edges via the line graph $L(G)$, making "differences" the objects of computation. But $L(G)$ does not shrink the instance: $|V(L(G))| = m$ and $|E(L(G))| = \frac{1}{2} \sum_v \deg(v)^2 - m$, which grows with the sum of squared degrees (for the complete graph, $\sim n^3/4$ edges). And since $L(G)$ is built in polynomial time, the transformation preserves complexity class: were it to linearize an NP-hard problem, it would solve SAT in polynomial time—a contradiction. The exponent moves from *the search* to *the enlarged dual graph*.

Surprisal-driven sparsification. Use self-information $-\log P$ to keep only high-surprisal nodes (steep ridges) and prune gentle ones, then operate on the sparse skeleton. On typical, redundant inputs this is genuinely effective. But it does not bound worst-case cost: an adversarial input in which every adjacent slope is steep leaves the graph dense, at $O(n^4)$ after dualization. Two further gaps compound this: the dynamical formation of "memory attractors" is asserted to converge but never bounded in iterations; and the decoding step is itself posed as a constraint-satisfaction problem, which is NP-hard in general. The exponent compressed during encoding returns during decoding. The exponent moves from *understanding* to *generation*.

5.3. The Principle

> **Conservation of Exponential Cost.** Across every transformation considered—simultaneity, ridge-volume integration, line-graph duality, surprisal sparsification—the exponential cost is *conserved*. It changes location (time to space; enumeration to integration; search to dual graph; encoding to decoding) but is never eliminated.

This is not an accident of our particular constructions. It is what the structure of complexity theory demands: a polynomial-time change of representation cannot change a problem's complexity class. Every Polygon Graph construction above is itself a polynomial-time object, hence reducible to a Turing machine; it therefore cannot, by reformulation alone, dissolve a distinction that lives at the level of complexity classes.

6. What Is Not Proved

The conservation principle is suggestive, but its limits must be stated plainly.

6.1. The Logical Gap

We have shown that *the Polygon Graph does not yield a polynomial-time algorithm for SAT by the routes considered*. We have *not* shown that *no polynomial-time algorithm for SAT exists*. The latter is the assertion $\mathrm{P} \neq \mathrm{NP}$; the former does not entail it. That these particular bridges fail to span the river is no proof that the river cannot be bridged. Approaches we did not consider—and approaches not yet conceived—are untouched by this analysis.

6.2. The Known Barriers

A proof of $\mathrm{P} \neq \mathrm{NP}$ would have to defeat every algorithm, present and future, at once. Complexity theory has identified formal *barriers* that explain the difficulty: relativization (Baker–Gill–Solovay), natural proofs (Razborov–Rudich), and algebrization. These show that whole families of proof technique cannot, in principle, settle the question. Elementary combinatorial and geometric arguments—the kind the Polygon Graph naturally supplies—are precisely those the barriers obstruct.

6.3. Honest Status

Question	**Status**
Does the Polygon Graph give a poly-time SAT algorithm?	No, by the routes here
Does this prove $\mathrm{P} \neq \mathrm{NP}$?	No
Is the conservation pattern evidence for $\neq$?	Weak, suggestive only
Is P versus NP open?	Yes

Conservation of exponential cost is consistent with the prevailing expectation that $\mathrm{P} \neq \mathrm{NP}$, but consistency is not proof. "Never eliminated in our attempts" is an inductive expectation, not a theorem.

7. The Genuine Contribution: A Geometry of Tractability

The model's lasting value is not a verdict on P versus NP. It is a language for a sharper and more answerable question.

7.1. Reframing the Question

Not "Is $\mathrm{P} = \mathrm{NP}$*?" but:*
which Polygon Graphs are tractable, which are not,
and can the boundary be drawn geometrically?

This question needs no resolution of P versus NP, survives any single counterexample, and admits positive theorems.

7.2. The Even-Degree Criterion

Observation 7.1 (Eulerian tractability)**.** *A Polygon Graph in which every node has even degree corresponds to an Eulerian graph and is solvable in polynomial time: an Eulerian path is found in* $O(m)$ *by Hierholzer's algorithm. Graphs containing odd-degree nodes, under general constraints, remain NP-hard.*

This is a real dividing line, expressed in the model's own geometric vocabulary—the *parity of node degree*. It is the seed of a program: to classify which structural features (degree parity, branching symmetry, conservation laws that align interference) place a problem on the tractable side of the line.

7.3. The Integrable-Systems Analogy

The situation echoes integrable systems in physics: not solvable in general, but exactly solvable when enough conserved quantities are present to make all contributions align. The Polygon Graph offers a setting to ask the computational version of this question—under what geometric conditions do interference paths align so completely that search becomes tractable? A characterization of this kind is robust against counterexample and far more defensible than any sweeping claim about P versus NP.

7.4. Architectural Corollaries

Two ideas survive as engineering, independent of complexity:

- **Surprisal sparsification** is a sound *average-case* efficiency mechanism for real, redundant language—properly framed as empirical speedup, not worst-case linearity.
- **Signed (destructive) interference** in attention is a substantive proposal for representing negation and contradiction directly, to be judged empirically.

8. Conclusion

The Polygon Graph reaches surprisingly far: it recovers the Transformer's attention as integration, and it speaks the language of interference. Pushed toward P versus NP, it invites a natural conjecture—that simultaneous branching collapses exponential search—which proves false for a reason that generalizes into the *conservation of exponential cost*: across integration, duality, and sparsification, the exponent is relocated but never removed, exactly as the polynomial-time invariance of complexity classes requires.

We claim no more. This is not a proof that $P \neq NP$; that question remains open, shielded by barriers our methods do not cross. What the model offers instead is a geometry of tractability—beginning with the even-degree criterion and pointing, by analogy with integrable systems, toward a structural account of which problems can be solved efficiently. That program asks less than a resolution of P versus NP, and for precisely that reason it can deliver results that stand.

Related work by the author: the Polygon Graph / Transformer identity; a dynamic-nuance architecture for language; a self-reference model of embodied AI.