PRICE: ONE ORACLE QUERY
WEEKLY · P ≠ NP (PROBABLY)
VOL. 1 NO. 44
2026-W44

the complexity chronicles

ALL THE NEWS THAT'S FIT TO REDUCE
Every exponent printed, every constant counted.

MONDAY 2 NOVEMBER 2026 · COVERING 26 OCTOBER – 1 NOVEMBER 2026 · 85 PAPERS · ARXIV CS.CC, ECCC, HEXAGON, DIDEROT

INSIDE THIS ISSUE

▸ p.2 — Two colours promised, a few logs delivered: colouring 2-colourable 3-uniform hypergraphs with $O((\log n)^{1/8})$ colours is now quasi-NP-hard, an exponential step up from $(\log\log n)^{1/3}$. ▸ p.2 — Sherali-Adams meets its betters: Res(2) beats it exponentially, and it is not size-automatable unless $\mathrm{P} = \mathrm{NP}$. One new technique, two open problems closed. ▸ p.3 — Inside the counting engine: what the claimed ETH refutations actually compute, and how large n must be before brute force is out of work. ▸ p.4 — Fifty-seven years to nine-quarters: the long walk of ω from Strassen to the claimed $9/4$, with an explicit algorithm at the end of it. ▸ Column — Lew Erbound, Rank and File: a week of interpolants, expositors and very large constants.

FINE-GRAINED COMPLEXITY

ETH falls. Brute force keeps its job

Two analyses of the claimed October 2026 refutations of ETH find that the subexponential #3-SAT bounds certify no speed-up over brute force for any bounded-density input that could ever be written down, cut one algorithm's crossover point to about 2^{420} clauses at clause density one, and observe that the Sparse Max-3-SAT Hypothesis falls too; a third paper, an AI-assisted exposition, explains how the counting procedure works.

This month two preprints, one from OpenAI and one from Anthropic, claimed deterministic algorithms that count the satisfying assignments of a 3-CNF in time $2^{o(n)}$. If they stand, the Exponential Time Hypothesis is false. None of the papers below claims to have checked those proofs in full. They did something more useful. They got out the calculator.

The mechanism comes first. Kominers [hexagon:2610.00574v1] gives a qualitative form of the OpenAI result in an exposition that is AI-assisted and that he calls preliminary. For every fixed $\varepsilon > 0$, 3-SAT on n variables and m clauses can be decided in time $O_\varepsilon(2^{\varepsilon n}(n+m)^{O(1)})$ and polynomial space. He separates the counting procedure from the construction that supplies its input. The clause inputs are deformed so that, at designated points, the count factors into local sums. Hermite interpolation then recovers the original count. The deformations are built from coordinate swaps in a tensor basis, using a parity interpolation in the spirit of Alman and Williams.

Then the bill. Dias and Lima [arXiv:2611.00517] note that the OpenAI preprint claims $2^{O_k(n\log\log n/\sqrt{\log n})}$ for k-CNF, and $2^{O(n/\sqrt{\log n})}$ at bounded clause density, but states no constants. If the hidden constant C is treated as a parameter, the general bound beats 2^n only beyond $n \approx 10^{32}$ for $C = 2$ and 10^{3517} for $C = 10$. The constants they extract from the tensor-contraction proof are worse. At bounded density the certified running time stays above 2^n unless $\log_2(n+m) > 10^9$. Their conclusion is that the bounds are subexponential as $n \to \infty$, but certify no speed-up over brute force for any bounded-density input that could ever be written down. They also prove limits for the clause-product form of the interpolation scheme. In a cost model matching the preprint's evaluation, it costs $2^{\Omega(\sqrt{\mathrm{tw}})}$ at treewidth tw, and $2^{\Omega(n^{1/3})}$ even on path formulas that are trivially countable. They explain why the preprint's sharper routes escape these limits.

Majdoddin [arXiv:2610.41622] makes the Anthropic #3-SAT recursion explicit and improves it by proved changes. At clause density one, the crossover point drops from about 2^{33858} clauses to about 2^{420}. For general 3-SAT, his estimate in the number of variables drops from $2^{2.2\cdot10^{12}}$ to $2^{8.6\cdot10^9}$. That is progress, of a kind. With integer tables, the recursion counts the assignments satisfying exactly j clauses, for every j, in time $2^{O(m/\sqrt{\log m})}$. He observes that this refutes the Sparse Max-3-SAT Hypothesis of Alman and Vassilevska Williams. The same refutation also follows from either ETH refutation via the reduction of Impagliazzo, Paturi and Zane.

None of this gives a practical algorithm, and the gap will not close by tinkering. Majdoddin defines a rewrite-and-merge class containing the recursion and all his constructions. He proves that, in his main-term model, no grid algorithm in that class can gain for general 3-SAT below about

2^{6000} variables. For two balanced homogeneous merges, he also excludes running times of the form $2^{n^{1-c}}$. The hypothesis may be dead. Brute force has not been told.

[arXiv:2611.00517] Limits and Constants of OpenAI's Claimed ETH Refutation (Fabio Dias, Yvens Lima) · [arXiv:2610.41622] Lowered crossover points of Anthropic's subexponential 3-SAT algorithm (Ruhollah Majdoddin) · [hexagon:2610.00574v1] On counting satisfying assignments via polynomial separation (Scott Kominers)

MATRIX MULTIPLICATION

Omega drops to nine-quarters

Shpilka gives an elementary, self-contained proof of the bound $\omega \leq 9/4$ claimed in an OpenAI manuscript, with explicit algorithms of exponent below $9/4 + \varepsilon$ and base matrix size at most $2^{(1/\varepsilon)^{O(1)}}$ for every $\varepsilon > 0$.

A recent OpenAI manuscript proves that the exponent of matrix multiplication satisfies $\omega \leq 9/4$. The previous record stood a little above 2.37. Shpilka [ECCC:TR26-281] now gives an elementary, self-contained proof. He follows the original ideas but implements them differently, so that the algorithms are explicit. For every $\varepsilon > 0$ there is an algorithm with exponent below $9/4 + \varepsilon$ and base matrix size at most $2^{(1/\varepsilon)^{O(1)}}$.

The base size is the small print. Nobody will be multiplying matrices this way before lunch. But an independent, elementary route to the bound is exactly what a claim of this size needs. An appendix describes the original proof, which goes through Strassen's theory of the asymptotic spectrum.

Similar techniques reach rectangular multiplication. Starting from Strassen's tensor, Shpilka gives elementary proofs that $\omega(1, \frac{2}{3\log 3}, 1) = 2$ and $\omega(1, 0.539, 2) = 3$. Stronger rectangular bounds are already known; these proofs are included for their simplicity. The longer history is in this week's perspective.

[ECCC:TR26-281] Exposition of OpenAI's Matrix Multiplication in Exponent 9/4 (and an explicit algorithm) (Amir Shpilka)

PROOF COMPLEXITY

Sherali-Adams loses twice

Res(2) is exponentially stronger than Sherali-Adams, and Sherali-Adams is not size-automatable unless $\mathrm{P} = \mathrm{NP}$; both results come from a new polynomial-interpolation method for Sherali-Adams degree lower bounds.

Acheampong, Davis and Robere [ECCC:TR26-288] resolve two open problems about the Sherali–Adams proof system. First, Res(2), which manipulates 2-DNF formulas, is exponentially stronger than Sherali–Adams. Sherali–Adams polynomially simulates Resolution, so the separation is sharp. One step up the bounded-depth Frege hierarchy, Sherali–Adams is already outclassed.

Second, Sherali–Adams is not size-automatable unless $\mathrm{P} = \mathrm{NP}$. Short proofs may exist. Finding them in time polynomial in their size is, under the usual assumption, off the table. Many researchers have asked about the automatability of weak proof systems, and this is one of the answers.

Both results rest on one new technique: lower bounds on Sherali–Adams degree via polynomial interpolation. A technique that pays twice in one paper is worth watching.

[ECCC:TR26-288] Polynomial Interpolation for Sherali-Adams Lower Bounds (Nathan Acheampong, Ben Davis, Robert Robere)

COMMUNICATION COMPLEXITY

Sign rank breaks under AND

An explicit matrix of sign rank $\Theta(d)$ whose conjunction with itself has sign rank $2^{\Theta(d)}$ shows that the unbounded-error communication class $\mathrm{UPP}^{\mathrm{cc}}$ is not closed under intersection.

Bun and Juvekar [ECCC:TR26-289] (also [arXiv:2610.37634]) work with sign rank: the minimum rank of a real matrix that agrees with a given ± 1 matrix entrywise in sign. For every sufficiently large d, they construct an explicit ± 1 matrix A with $2^{\Theta(d^2)}$ rows and columns and sign rank $\Theta(d)$. Its conjunction $A \wedge A$ has sign rank $2^{\Theta(d)}$. Here $A \wedge A$ is the problem of evaluating A on two disjoint inputs and taking the AND of the results.

The logarithm of sign rank characterises unbounded-error communication complexity. So a problem costing about $\log d$ bits becomes one costing about d bits when two instances are joined by an AND. That proves the title: $\mathrm{UPP}^{\mathrm{cc}}$ is not closed under intersection. One small conjunction produces one exponential jump.

[ECCC:TR26-289] UPP^{cc} is Not Closed Under Intersection (Mark Bun, Mandar Juvekar) · [arXiv:2610.37634] UPP^{cc} is Not Closed Under Intersection (Mark Bun, Mandar Juvekar)

IN DEPTH 2026-W44

PROMISE CSP

Two colours promised, a few logs delivered

Bhangale, KM and Rubin-Toles prove it is quasi-NP-hard to colour a 2-colourable 3-uniform hypergraph on n vertices with $O((\log n)^{1/8})$ colours, an exponential improvement on the $O((\log \log n)^{1/3})$ of Dinur, Regev and Smyth from 2005. Separately, Brakensiek, Conlon, Guruswami, He, Nie, Sankar and Yu use limits on 2-query locally decodable and correctable codes to bound CSP non-redundancy, including the first Boolean predicate whose non-redundancy is provably not of the form $n^{k \pm o(1)}$.

Some problems do not move for twenty years. Then they move by an exponential, and everyone pretends they expected it.

This week it was hypergraph colouring. The question is simple to state. You are handed a 3-uniform hypergraph and a promise that it can be properly coloured with two colours, meaning no edge is monochromatic. You are not asked to find the two-colouring. You are asked for something far more modest: a proper colouring with some larger number of colours. How many colours can you be forced to use before the problem becomes hard?

WHAT WAS KNOWN

In the language this paper's reader speaks, this is a promise CSP. The yes-instances satisfy a strong template and the algorithm need only satisfy a weak one. The programme to classify such problems has made real progress, but approximate colouring has stayed stubborn, and the hypergraph case is the one where hardness has historically been easiest to get and slowest to improve.

> *Between a power of $\log n$ and a power of n there is a great deal of room, and nobody knows where the truth lives.*

The reference point was Dinur, Regev and Smyth, published in *Combinatorica* in 2005. They showed it is quasi-NP-hard to colour a 2-colourable 3-uniform hypergraph with $O((\log \log n)^{1/3})$ colours. That is a function that grows, technically. It is not a function anyone would want to plot.

WHAT IS NEW

Bhangale, KM and Rubin-Toles [arXiv:2610.37869] raise the bar to $O((\log n)^{1/8})$ colours, still under quasi-NP-hardness. The exponent $1/8$ looks modest. The change from $\log \log n$ to $\log n$ under the root is not. It is an exponential improvement on a bound that had stood for two decades, and the authors say as much in plain terms.

The hedging matters and should be stated plainly. Quasi-NP-hardness means that a polynomial-time algorithm would put NP inside quasi-polynomial time. It is not NP-hardness under polynomial-time reductions. The paper is twenty pages and the abstract does not describe the construction, so this newspaper will not describe it either. The tags it carries, PCP among them, suggest the usual machinery of label cover and long-code-style gadgets, but suggestion is all it is.

WHY THE READER SHOULD CARE

The gap here is the whole story. Hardness has crept up from $\log \log$ territory to a fractional power of $\log n$. Efficient colouring algorithms for this promise, by the general understanding of the area, need a number of colours that grows as a power of n. Between a power of $\log n$ and a power of n there is a great deal of room, and nobody knows where the truth lives. Every step that pushes hardness upward narrows the space in which an algebraic or SDP-based algorithm might be hiding, and tells the classification programme which minions it should expect to be hard.

THE OTHER END: HOW MUCH OF AN INSTANCE MATTERS

The second paper asks a different question about CSPs, but one with consequences for anyone who cares about kernels and sparsifiers. Brakensiek, Conlon, Guruswami, He, Nie, Sankar and Yu [arXiv:2610.38664] study *non-redundancy*: the maximum number of clauses in an instance of a CSP over a fixed relation that can be logically independent, so that no clause is implied by the others. The abstract notes that this quantity has turned up in query complexity, sparsification, kernelization and streaming. Put crudely, it caps how far one can compress an instance without losing information.

The new idea is to borrow from coding theory. Known structural limits on 2-query locally decodable codes and locally correctable codes are translated into limits on the size of certain non-redundant CSP instances. A large non-redundant instance would, in effect, give a better local code than is known to exist. This builds on earlier work of Brakensiek and Guruswami at STOC 2025, which bounded non-redundancy using matching vector codes.

Two outcomes are reported. First, improved upper bounds on the non-redundancy of the BCK predicate introduced by Bessiere, Carbonnel and Katsirelos at AAAI 2020, and of its well-studied Boolean variant. The abstract does not give the numbers and neither shall we. Second, and more curious, the first Boolean predicate whose non-redundancy is provably not of the form $n^{k \pm o(1)}$ for any positive integer k. The tidy expectation that these quantities come in integer exponents is now, for at least one predicate, provably wrong.

WHAT LINKS THEM

Both papers are about the fine structure of CSPs below the level at which the decision dichotomy is settled. The Bulatov and Zhuk theorem tells us which finite-domain CSPs are in P. It does not tell us how many colours a promise buys, or how small an instance can be made before it stops meaning the same thing. Those are the questions that keep the field employed.

WHAT REMAINS OPEN

For hypergraph colouring: whether hardness can be pushed to a power of n, whether quasi-NP-hardness can be upgraded

to NP-hardness at this range, and whether the exponent $1/8$ is an artefact of the method. For non-redundancy: what the true exponents for the BCK predicates are, and how strange non-redundancy can get once integer exponents are off the table. The code-theoretic route gives upper bounds; matching lower bounds are another matter.

Twenty years for one exponential. The field may hope the next one comes quicker. It usually does not.

[arXiv:2610.37869] Hardness of Coloring 2-Colorable 3-Uniform Hypergraphs with $\log^{\Omega(1)} n$ Colors (Amey Bhangale, Krishna Priya KM, Simon Rubin-Toles) · [arXiv:2610.38664] Upper Bounds for CSP Non-redundancy via Locally Decodable Codes (Joshua Brakensiek, David Conlon, Venkatesan Guruswami, Xiaoyu He et al.)

PROOF COMPLEXITY

Sherali-Adams meets its betters

Acheampong, Davis and Robere prove that Res(2) is exponentially stronger than Sherali-Adams, and that Sherali-Adams is not size-automatable unless P = NP. Part proves $2^{\Omega(\epsilon n)}$ depth-three lower bounds for Tseitin formulas in Hrubeš' Algebraic Monotone Calculus when $\epsilon = \omega(\log n/n)$. Conneryd proves near-$n/2$ polynomial calculus degree bounds for random $O(\log n)$-CNFs over every field, which give superpolynomial depth-2 $\mathrm{AC}^0[p]$-Frege lower bounds via Pang's tradeoff.

Proof complexity has a habit of leaving its weakest systems half-mapped. Everybody knows roughly where they sit. Nobody has the theorem. This week three papers fill in more of the chart around Sherali-Adams, the linear-programming hierarchy that constraint people meet early and never quite shake off. Two of the three pieces of news are things the field had been asking for by name.

TWO ANSWERS FROM ONE TECHNIQUE

The headline paper is by Acheampong, Davis and Robere [ECCC:TR26-288]. It settles two open problems with a single new tool.

The first is a separation. Sherali-Adams polynomially simulates Resolution; that much was known, and the abstract states it. Res(2) is Resolution allowed to work with 2-DNF formulas rather than plain clauses, which makes it the first rung above Resolution on the bounded-depth Frege ladder. The authors show that Res(2) is *exponentially* stronger than Sherali-Adams. So Sherali-Adams sits above Resolution and is beaten outright by the next system up. The abstract calls this a sharp separation between Sherali-Adams and the bounded-depth Frege hierarchy. That is fair. There is not much room left between the two.

> *It is rare for a paper to tell you exactly where its method stops. This one does.*

The second result is about automatability: whether proofs can be found in time polynomial in the size of the shortest proof. The authors prove that Sherali-Adams is not size-automatable unless P = NP. The abstract notes that this had been asked by many researchers. It joins the list of weak systems where short proofs exist and finding them is as hard as anything in NP.

Both results come from a new way of proving lower bounds on Sherali-Adams *degree*, via polynomial interpolation. The abstract gives no further detail of the mechanics, and this paper will not supply any. What can be said is that one degree technique delivered both a separation and a hardness-of-search result, which is a decent return on a single idea.

MONOTONE CIRCUITS AS PROOFS

Part's paper [ECCC:TR26-305], also on arXiv as [arXiv:2611.01055], works in a stranger corner. In Hrubeš' Algebraic Monotone Calculus a DNF tautology A is proved by a monotone arithmetic circuit computing $\chi(A) - \epsilon \cdot \chi(1)$, where the coefficient of $\chi(A)$ at an assignment counts the terms of A it satisfies. Proof size is circuit size. The shift $\epsilon > 0$ is existentially quantified, so a lower bound has to survive arbitrarily small ϵ. That is the catch, and Part spends most of the paper on it.

At depth three a proof is a nonnegative tensor decomposition. Restrict the summands to subcube indicators and you get conical juntas, a system the paper shows is polynomially equivalent to Sherali-Adams when size counts conjunctions over the twin variables. Allow arbitrary nonnegative rank-one summands and you get something stronger. For that general fragment Part adapts Rothvoß' hyperplane separation method to monotone tensor rank and proves that Tseitin formulas on constant-degree expanders of large girth need size $2^{\Omega(\epsilon n)}$ whenever $\epsilon = \omega(\log n/n)$.

The bound cannot reach every ϵ, and the paper says why. At shift $2^{-\Theta(n \log n)}$ Tseitin acquires polynomial-size depth-three proofs. That fact doubles as an exponential separation between the general depth-three fragment and the Sherali-Adams-like subcube fragment. To cover every $\epsilon \geq 2^{-\Omega(n)}$, Part composes Tseitin with a constant-size index gadget, turning a Sherali-Adams *degree* bound into a depth-three *size* bound. Depth three is essential: Tseitin has monotone proofs of size $O(n^3)$ and depth $O(\log n)$ at every ϵ.

Beyond depth three Part bounds the nonnegative rank of a flattening by L^{2^t} for a monotone formula with L leaves and mixed-product depth t. The abstract is candid that this is empty for Tseitin. Instead it is aimed elsewhere: approximate extension complexity gives bounds for perfect matching and odd factor instances, and linear programming integrality gaps for planted constraint satisfaction give an explicit family of planted Tseitin DNFs hard at inverse-polynomial shift. CSP readers will note where the hard instances came from.

RANDOM FORMULAS, MODULAR GATES

Conneryd [arXiv:2610.40397] goes after random k-CNFs with $M = n2^k$ clauses, for $\omega_n(1) \leq k \leq O(\log n)$. With high probability, refuting them needs polynomial calculus degree at least $n/2 - O(n/\sqrt{k})$, over every field. Feed that into Pang's recent tradeoff between $\mathrm{AC}^0[p]$-Frege size and polynomial calculus degree, and you get a size lower bound of $\exp(\Omega_p(\sqrt{k}\log n))$ for

depth-2 $\mathrm{AC}^0[p]$-Frege. In this range the bound is superpolynomial in M while M stays polynomial in n.

The paper then closes its own door. A nearly matching upper bound of $n/2 - n\ln(k)/2k + O(n/k)$ on the degree shows that random CNFs will not carry Pang's tradeoff to full $\mathrm{AC}^0[p]$-Frege. Other formula families are needed. It is rare for a paper to tell you exactly where its method stops. This one does.

WHAT IS LEFT

The chart is clearer but not finished. The interpolation technique for Sherali-Adams degree is new, and the obvious question is how far up it travels; the abstract claims nothing beyond Sherali-Adams, so neither will this paper. Part's lower bounds stop at depth three for Tseitin, and the general depth-three bound stops short of every ϵ for a reason that is a theorem rather than a gap in effort. Conneryd has shown that full $\mathrm{AC}^0[p]$-Frege via Pang's route needs formulas other than random ones, and has not said which.

That is three careful papers and three honest stopping points. In this trade, that is a good week.

[ECCC:TR26-288] Polynomial Interpolation for Sherali-Adams Lower Bounds (Nathan Acheampong, Ben Davis, Robert Robere) · [ECCC:TR26-305] Non-negative Rank Bounds for the Algebraic Monotone Calculus (Fedor Part) · [arXiv:2611.01055] Nonnegative Rank Bounds for the Algebraic Monotone Calculus (Fedor Part) · [arXiv:2610.40397] Random ($\log n$)-CNF Formulas Are Hard for Depth-2 Frege with Modular Gates (Jonas Conneryd)

FINE-GRAINED COMPLEXITY

Inside the counting engine

Two AI-assisted expositions, by Kominers and O'Donnell, describe how OpenAI's claimed refutation of the Exponential Time Hypothesis is supposed to decide 3-SAT in time $O_\varepsilon(2^{\varepsilon n}(n+m)^{O(1)})$ and polynomial space for every fixed $\varepsilon > 0$. Neither is a verification. Separately, Kozma proves unconditionally that #3SAT in time $2^{O(n/f(n))}$ would give #PPM in time $2^{O(n\sqrt{\log f(n)/f(n)})}$.

The claim is simple to state and hard to swallow. 3-SAT, the problem a generation of conditional lower bounds has leaned on, would run in time $2^{\varepsilon n}$ for every fixed $\varepsilon > 0$. That is the end of the Exponential Time Hypothesis. Since SETH implies ETH, it is the end of SETH as well.

The original manuscript is not on this desk. What we have are two expositions, both written with the help of language models, and both labelled accordingly. This feature sets out what they say the machine does. It does not say the machine works.

WHAT THE EXPOSITIONS CLAIM

Kominers [hexagon:2610.00574v1] gives what he calls a preliminary exposition of a qualitative form of the main result: for every fixed $\varepsilon > 0$, 3-SAT on n variables and m clauses can be decided in time $O_\varepsilon(2^{\varepsilon n}(n+m)^{O(1)})$ and in polynomial space. The space bound deserves a second look. Plenty of exponential-time methods buy speed with memory. This one, as described, does not.

> *Parity is the function that famously refuses to be approximated by low-degree polynomials; here an approximating polynomial for parity is apparently doing the work in an upper bound.*

The algorithm counts. It computes the number of satisfying assignments, and the decision problem falls out. That is the usual direction of travel for algebraic methods, and the reason the downstream papers this week are about counting problems.

The paper is tagged AI-assisted, with ChatGPT, Claude Opus and Poe listed as tools; its provenance is recorded as mixed, and human understanding as "some". The O'Donnell exposition [hexagon:2610.00572v1] says it was produced through discussion with an LLM. Its record lists provenance as "no-human-author-claimed" and human understanding as "none". A reader should weigh it accordingly. A famous name on the header is not the same thing as a checked proof, and the document does not pretend otherwise.

HOW THE ENGINE IS MEANT TO RUN

Kominers separates two things that the original apparently bundles: a general counting procedure, and a construction that feeds it.

The procedure works by deformation. The clause inputs are perturbed so that, at certain designated points, the count of satisfying assignments factors into local sums. Local sums are cheap. The original count is then recovered by Hermite interpolation, the classical technique that rebuilds a polynomial from its values and derivatives at a set of points. The idea, as far as the abstract goes, is to evaluate a hard global quantity only where it happens to be easy, and to let interpolation carry the answer back to where it is hard.

The construction supplies the deformations. According to Kominers, they are built from coordinate swaps in a tensor basis, and lean on a parity interpolation "in the spirit of Alman and Williams". O'Donnell's exposition puts the same ingredient at the centre: it highlights the approximating polynomial for parity, and writes its matrix version in a light dose of quantum notation.

Readers raised on circuit lower bounds will notice the irony. Parity is the function that famously refuses to be approximated by low-degree polynomials; that refusal is the engine of a good part of the lower-bound literature. Here an approximating polynomial for parity is apparently doing the work in an upper bound. The abstracts do not say in what sense the approximation holds, over which ring, or at what degree. Those are exactly the questions a referee would ask first.

What the abstracts also do not give is constants. The statement is O_ε, and the subscript is where the bodies are buried. A $2^{\varepsilon n}$ algorithm with a hidden factor that is astronomical in $1/\varepsilon$ refutes a hypothesis without troubling a SAT solver. That is no objection to the mathematics. ETH is an asymptotic statement and an

asymptotic refutation counts. It is an objection to anyone who promises faster verification software by Christmas.

THE FIRST CASUALTY: PATTERN MATCHING

Kozma's paper [arXiv:2610.41048] is the sort of thing that appears when a field takes a claim seriously enough to plan for it. It concerns #PPM: count the occurrences of a permutation pattern in a host permutation of length n. The best known algorithm, for both the decision and counting versions, runs in $O(1.415^n)$, due to Gawrychowski and Rzepecki. In the other direction, Jelínek, Opler and Pekárek showed PPM is ETH-hard: a $2^{o(n/\log n)}$ algorithm would give subexponential 3SAT.

Kozma supplies the converse. If #3SAT runs in time $2^{O(n/f(n))}$, then #PPM runs in time $2^{O(n\sqrt{\log f(n)/f(n)})}$. Feed in any unbounded f and #PPM becomes subexponential. The square root costs something, but the direction is the point: the exponential complexities of the two problems are now tied in both directions, up to those losses.

The reduction is unconditional. It stands whether or not the counting engine does. If the claim collapses, Kozma has still pinned #PPM to #3SAT more tightly than before. If the claim holds, it is an early entry in what would become a very long list.

WHAT REMAINS OPEN

For this reader the consequences are close to home. A great deal of fine-grained and parameterized work on CSPs is stated relative to ETH or SETH; those lower bounds would lose their hypothesis, though not their reductions. The reductions would turn into algorithms.

But the order of business is plain. First, a human-checked proof. Second, an account of the parity polynomial precise enough to see why it does not contradict what we know about parity. Third, constants. Two expositions, one of them with no claimed human understanding, are a start. They are not a verdict. This paper will report the verdict when there is one, and not before.

[hexagon:2610.00574v1] On counting satisfying assignments via polynomial separation (Scott Kominers) · [hexagon:2610.00572v1] Some exposition of "The Exponential Time Hypothesis Is False" (Ryan O'Donnell) · [arXiv:2610.41048] Faster #PPM from faster #3SAT (László Kozma)

THE LONG VIEW

how we got here

THE LONG VIEW

Fifty-seven years to nine-quarters

Shpilka's elementary proof of OpenAI's claimed bound [ECCC:TR26-281] gives, for every $\varepsilon > 0$, an explicit matrix multiplication algorithm with exponent below $9/4 + \varepsilon$, an unrefereed result that caps fifty-seven years of chasing ω down from Strassen's 2.81 to the laser-method decimals near 2.371.

SEVEN FOR EIGHT

Until the late 1960s the working belief was that $\omega = 3$ [S1]. Strassen's 1969 note, "Gaussian elimination is not optimal", ended that belief in three pages [S2]. Two 2×2 matrices can be multiplied with seven multiplications instead of eight. The formulas never use commutativity, so the entries may themselves be blocks. The recursion $T(n) = 7T(n/2) + O(n^2)$ then gives $O(n^{\log_2 7}) = O(n^{2.81})$ [S3]. The same paper showed that inversion, the determinant and Gaussian elimination cost no more than multiplication, up to constants [S3]. One saved product, and a field had its exponent.

The next decade was a hunt for small exact schemes. Pan's 1978 construction multiplied 70×70 matrices with 143,640 products, for $\omega < 2.796$ [S1]. Honest work, bought one multiplication at a time.

APPROXIMATION IS FREE

In 1979 Bini, Capovani, Romani and Lotti let the coefficients depend on an indeterminate λ. They computed the product only up to terms that vanish as $\lambda \to 0$ [S4]. Ten products sufficed for a small rectangular product. Three cyclically permuted copies gave 12×12 multiplication with 1000 products, so $\omega \le \log_{12} 1000 < 2.78$ [S1]. The idea is now called border rank. Bini showed that such approximate algorithms bound ω just as exact ones do [S1].

Schönhage's 1981 paper on partial and total matrix multiplication [S5] added the asymptotic sum inequality, the τ-theorem. A direct sum of several smaller rectangular products, computed together, is enough to bound ω [S6]. Formally, if the border rank of $\bigoplus_i \langle a_i, b_i, c_i \rangle$ is at most r, then $\omega \le 3\tau$, where $\sum_i (a_i b_i c_i)^\tau = r$ [S7].

His example made the point. A 3×1 by 1×3 product and a 1×4 by 4×1 product can be computed together, approximately, with 10 products. Any exact method needs at least 13 [S1]. Border rank is not additive. The bound came to $\omega < 2.548$ [S6], and he showed that ω depends only on the characteristic of the field [S7].

Coppersmith and Winograd then proved that ω is a limit point, attained by no single algorithm [S8]. That is why everyone writes $O(n^{\omega+\epsilon})$ [S3].

THE LASER

Strassen's laser method of 1986–87 changed the starting material. Begin with a tensor that is not a matrix product at all [S1]. Take a large tensor power and zero out parts until many independent matrix products remain. Then feed them to the τ-theorem [S1, S6]. It gave $\omega < 2.48$ [S9]. By 2020 every record set since 1986 had used it [S7].

In the same years Strassen built the theory of the asymptotic spectrum of tensors, published in Crelle in 1988 [S10]. Tensors form a semiring under direct sum and tensor product. Spectral points are maps to the non-negative reals that are monotone under restriction, normalised on diagonal tensors, additive and multiplicative. A duality theorem describes asymptotic rank through them [S11]. Bounding the asymptotic rank of matrix-multiplication tensors bounds ω [S12]. Keep this in mind. It returns.

Coppersmith and Winograd sharpened the laser with the Salem–Spencer theorem on dense sets of integers with no three-term arithmetic progression [S13, S14]. They reached $\omega < 2.388$

in 1987 and 2.376 in 1990, the latter from the second tensor power [S15, S16]. That number stood until 2010 [S3].

THE DECIMAL ERA

Stothers broke the silence in 2010 with the fourth tensor power; the third had given nothing [S1]. The published Davie–Stothers bound is $\omega < 2.37369$ [S17]. In 2012 Vassilevska Williams turned the analysis of higher powers into a computer-solved optimisation problem, for $\omega < 2.3727$ [S18, S17]. Le Gall analysed further powers and reached 2.3728639 in 2014 [S19].

Then the bill arrived. Ambainis, Filmus and Le Gall showed in 2014 that analysing ever higher powers of the Coppersmith–Winograd identity cannot give $\omega < 2.3725$. A wide class of variants cannot go below 2.3078, and so cannot reach 2 [S20]. Alman and Vassilevska Williams, and separately Christandl, Vrana and Zuiddam, extended such barriers to major generalisations of the known approaches [S21].

The records kept coming anyway, at roughly one decimal place per conference cycle:

- Alman and Vassilevska Williams' refined laser method gave 2.3728596 [S9].
- Duan, Wu and Zhou found a hidden combination loss and compensated with asymmetric hashing. They got 2.371866 [S9], below the 2.3725 barrier [S3].
- Vassilevska Williams, Xu, Xu and Zhou reached 2.371552 [S9].
- Alman, Duan, Vassilevska Williams, Xu, Xu and Zhou reached 2.371339 [S22].
- In August 2026 Dupont and colleagues, using numerical optimisation and AlphaEvolve, reached 2.371177 [S23].

Fifty-seven years had moved the exponent from 2.81 to 2.371.

The rectangular story ran in parallel. The dual exponent α is the largest k with $\omega(1, k, 1) = 2$. Coppersmith put it at 0.172 in 1982 [S16]. Vassilevska Williams, Xu, Xu and Zhou brought it to 0.321334 [S24]. Methods built on the big Coppersmith–Winograd tensors cannot push it past 0.6218 [S25].

NINE-QUARTERS

On 24 September 2026 OpenAI released a manuscript claiming ω below 2.258. It was one of 722 manuscripts the company says an unreleased internal model produced [S26]. A second manuscript, dated 2 October, is titled "An Upper Bound of 9/4 for the Matrix Multiplication Exponent" and claims $\omega \leq 9/4$ over $\mathbb{C}$ [S27, S28].

The route bounds matrix multiplication via polynomial multiplication [S26]. It works inside Strassen's asymptotic-spectrum framework, with non-constructive arguments [S28]. A Lean formalisation is reported [S26, S27]. OpenAI itself says its batch sits at different stages of verification [S26].

Shpilka's note [ECCC:TR26-281] prompts this look back [S29]. It gives an elementary, self-contained proof. It follows the original ideas but implements them so as to produce explicit algorithms: for every $\varepsilon > 0$, an exponent below $9/4 + \varepsilon$ with base matrix size at most $2^{(1/\varepsilon)^{O(1)}}$. Starting from Strassen's tensor, it also proves $\omega(1, \frac{2}{3\log 3}, 1) = 2$ and $\omega(1, 0.539, 2) = 3$, while noting that stronger rectangular bounds are known [S29].

There are other accounts. Pass gives a constructive one [S28]. Amanov and co-authors extend the bound to every field and prove barriers [S30]. Li identifies an untwisting lemma for semidirect sums as the core of the proof [S31]. None of this is peer review. Four independent re-derivations is still more than most claims get.

One inference of our own. $9/4$ sits below the 2.3078 barrier. Whatever the new proof is, it is not the laser method applied to Coppersmith–Winograd powers in the form Ambainis, Filmus and Le Gall ruled out.

WHAT IT DOES NOT SHOW

It does not show $\omega = 2$ [S3]. Lower bounds stay where they were: Raz's $\Omega(n^2 \log n)$, and that only for bounded-coefficient circuits [S3]. A base matrix of size $2^{(1/\varepsilon)^{O(1)}}$ is not a library routine. Like every record of recent decades, it is galactic [S3]. Strassen's seven products, meanwhile, pay their way for n above roughly 100 and sit in BLAS-type libraries [S3]. Fifty-seven years on, the 1969 algorithm is still the one doing the work.

This week's result: [ECCC:TR26-281] Exposition of OpenAI's Matrix Multiplication in Exponent 9/4 (and an explicit algorithm) (Amir Shpilka)

THE TECHNIQUE: RANK PLUS RECURSION: STRASSEN'S SEVEN PRODUCTS

A **bilinear algorithm** for $k \times k$ matrices forms r products. Each product is a linear form in the entries of A times a linear form in the entries of B. Every entry of $C = AB$ is then a linear combination of these products. No step uses commutativity, so the entries may be matrix blocks. Applied recursively to blocks, the algorithm costs $T(n) = rT(n/k) + O(n^2)$, so $\omega \leq \log_k r$. In tensor language, r is the rank of the matrix-multiplication tensor $\langle k, k, k \rangle$. Border rank and Schönhage's τ-theorem extend the same accounting to approximate algorithms and to direct sums.

Strassen's case is $k = 2, r = 7$. The seven products are:

- $M_1 = (a_{11} + a_{22})(b_{11} + b_{22})$
- $M_2 = (a_{21} + a_{22})b_{11}$
- $M_3 = a_{11}(b_{12} - b_{22})$
- $M_4 = a_{22}(b_{21} - b_{11})$
- $M_5 = (a_{11} + a_{12})b_{22}$
- $M_6 = (a_{21} - a_{11})(b_{11} + b_{12})$
- $M_7 = (a_{12} - a_{22})(b_{21} + b_{22})$

The entries of C are:

- $C_{11} = M_1 + M_4 - M_5 + M_7$
- $C_{12} = M_3 + M_5$

- $C_{21} = M_2 + M_4$
- $C_{22} = M_1 - M_2 + M_3 + M_6$

Toy instance. Take $A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$ and $B = \begin{pmatrix} 5 & 6 \\ 7 & 8 \end{pmatrix}$. The products are $M_1, \dots, M_7 = 65, 35, -2, 8, 24, 22, -30$. Then:

- $C_{11} = 65 + 8 - 24 - 30 = 19$
- $C_{12} = -2 + 24 = 22$
- $C_{21} = 35 + 8 = 43$
- $C_{22} = 65 - 35 - 2 + 22 = 50$

That is exactly AB.

At size $n = 2^k$ the recursion uses 7^k multiplications instead of 8^k: 16,807 against 32,768 at $n = 32$. Asymptotically this gives $\omega \leq \log_2 7 \approx 2.807$.

Every later bound, up to the claimed $9/4$, is this same arithmetic: find a structure whose (border, or asymptotic) rank is small relative to the matrix products it contains, then recurse. The difficulty lies entirely in finding the structure.

TIMELINE

1969 Strassen multiplies 2×2 matrices with 7 products; recursion gives $O(n^{2.81})$.

1979 Bini, Capovani, Romani and Lotti introduce approximate (border-rank) algorithms; $\omega < 2.78$.

1981 Schönhage's asymptotic sum inequality (τ-theorem) shows direct sums suffice; $\omega < 2.548$.

1986–1987 Strassen's laser method gives $\omega < 2.48$; his asymptotic spectrum of tensors appears in Crelle.

1987–1990 Coppersmith and Winograd use Salem–Spencer sets to reach $\omega < 2.388$, then 2.376.

2010–2013 Stothers, Vassilevska Williams and Le Gall analyse higher tensor powers; bounds near 2.3729–2.3737.

2014–2015 Ambainis, Filmus and Le Gall: higher Coppersmith–Winograd powers cannot beat 2.3725, and a wide class of variants cannot beat 2.3078.

2023–2024 Asymmetric hashing and its successors bring ω from 2.371866 to 2.371177 (Dupont et al., with AlphaEvolve).

Oct 2025 OpenAI claims $\omega \leq 9/4$ over $\mathbb{C}$ via polynomial multiplication and the asymptotic spectrum; unrefereed.

27 Oct 2025 Shpilka [ECCC:TR26-281]: elementary proof with explicit algorithms of exponent below $9/4 + \varepsilon$, base size $2^{(1/\varepsilon)^{O(1)}}$.

SOURCES

[S1] Vassilevska Williams, An overview of the recent progress on matrix multiplication (SIGACT News column)

[S2] Strassen, Gaussian elimination is not optimal

[S3] Wikipedia (via Wikiwand), Computational complexity of matrix multiplication

[S4] Bini, Capovani, Romani, Lotti, O(n^2.7799) complexity for n×n approximate matrix multiplication

[S5] Schönhage, Partial and Total Matrix Multiplication

[S6] Vassilevska Williams, Multiplying matrices in O(n^2.373) time

[S7] Alman & Vassilevska Williams, A Refined Laser Method and Faster Matrix Multiplication

[S8] Semantic Scholar, Strassen, The asymptotic spectrum of tensors and the exponent of matrix multiplication

[S9] Quanta Magazine, New Breakthrough Brings Matrix Multiplication Closer to Ideal

[S10] Strassen, The asymptotic spectrum of tensors (Crelle)

[S11] Christandl, Vrana, Zuiddam, Universal points in the asymptotic spectrum of tensors

[S12] Vassilevska Williams, Xu, Xu, Zhou, New Bounds for Matrix Multiplication: from Alpha to Omega

[S13] Coppersmith & Winograd, Matrix multiplication via arithmetic progressions (STOC '87)

[S14] IBM Research, Matrix multiplication via arithmetic progressions

[S15] Coppersmith & Winograd, Matrix Multiplication via Arithmetic Progressions (J. Symb. Comput. 1990)

[S16] Le Gall & Urrutia, Improved Rectangular Matrix Multiplication using Powers of the Coppersmith-Winograd Tensor

[S17] Semantic Scholar, Matrix multiplication via arithmetic progressions

[S18] Vassilevska Williams, Multiplying Matrices Faster Than Coppersmith-Winograd

[S19] arXiv 1905.05067, Dynamic Matrix Inverse (cites Le Gall 2014 bound)

[S20] Ambainis, Filmus, Le Gall, Limitations of the Laser Method

[S21] Alman, Duan, Vassilevska Williams, Xu, Xu, Zhou, More Asymmetry Yields Faster Matrix Multiplication

[S22] More Asymmetry Yields Faster Matrix Multiplication (SODA 2025)

[S23] Dupont et al., Improving the matrix multiplication exponent with modern optimization and AlphaEvolve

[S24] Revisiting "Rapid multiplication of rectangular matrices"

[S25] Le Gall et al., Barriers for rectangular matrix multiplication

[S26] OfficeChai, OpenAI Has Found A Far Faster Way To Multiply Matrices
[S27] BinaryVerse AI, OpenAI Matrix Multiplication: The Proof Behind the 2.25 Leap
[S28] Pass, Matrix Multiplication from Polynomial Packings
[S29] Shpilka, ECCC TR26-281
[S30] Amanov et al., Multiplying matrices via polynomials: exposition and barriers
[S31] Baitian Li, ECCC TR26-298

THE COLUMN

Lew Erbound



LEW ERBOUND, NIGHT DESK

One Toolkit, Two Targets

Thirty years on the complexity beat. Has seen every reduction twice and trusts none of them. Reads the week's abstracts after midnight, looking for the pattern the authors did not notice, and usually finds it.

Case file, week ending 1 November. Two manuscripts, filed independently, both claiming the Exponential Time Hypothesis is false. One is credited to OpenAI. The other is credited to Anthropic [hexagon:2610.00577v1]. It has a two-sentence abstract and a claimed Lean 4 formalisation. Within days the desk had a dozen papers downstream. That is not a theme. That is a weather system.

Sort the downstream by type.

Expositions. Kominers [hexagon:2610.00574v1] separates the counting procedure from its input. The clause inputs are deformed so that the count factors into local sums at designated points. Hermite interpolation then recovers the count, with a parity interpolation in the spirit of Alman and Williams. O'Donnell [hexagon:2610.00572v1] foregrounds the approximating polynomial for parity, in light quantum language. Its provenance tag reads *human understanding: none*, which is at least candid.

Corollaries. Kozma [arXiv:2610.41048] reverses a known reduction, so faster #3SAT gives faster #PPM. Bhattacharjee [hexagon:2610.00594v1] extends the counting to weighted sums and Max-k-SAT. Tulsyan [hexagon:2610.00672v1] counts Hamiltonian cycles in sparse graphs, conditional on three internal results. Rudich and Rousseau [hexagon:2610.00600v1] offer deterministic factoring in $N^{o(1)}$, backed by a 13-line Lean statement. Each stands or falls with the engine.

Audits. Dias and Lima [arXiv:2611.00517] extract constants from the tensor-contraction proof. At bounded density, the certified running time stays above 2^n unless $\log_2(n+m) > 10^9$. Majdoddin [arXiv:2610.41622] makes Anthropic's recursion explicit. He lowers the crossover to about 2^{420} clauses at density one. He also proves that no grid algorithm in his rewrite-and-merge class gains for general 3-SAT below about 2^{6000} variables. In passing, he notes the Sparse Max-3-SAT Hypothesis falls too. The hypothesis has been refuted for every input too large to exist. Asymptotics has always been generous that way.

Now the pattern. Read the ETH papers alongside Shpilka's exposition of OpenAI's $\omega \le 9/4$ [ECCC:TR26-281]. His appendix describes the original proof, which goes through Strassen's asymptotic spectrum. Majdoddin proves his limits with quantum functionals. His Signed Network Sums model contracts the four-node complete graph in $N^{23/6}$ if $\omega = 2$, and in about $N^{3.93}$ on the current preprint curve. The headline bound Dias and Lima audit is proved by tensor contraction. Stein Naves de Brito [hexagon:2610.00692v1] shows by a rank floor that the gain parameter $\rho = 1/\sqrt{2}$ is optimal for pieces of that shape. The SAT claim and the matrix-multiplication claim are not two stories. They are one toolkit, tensors and interpolation, aimed at two targets in the same month.

The toolkit is also working the other side of the street. Acheampong, Davis and Robere [ECCC:TR26-288] prove Sherali-Adams degree lower bounds *by polynomial interpolation*. They show Res(2) is exponentially stronger than Sherali-Adams, and that Sherali-Adams is not size-automatable unless P = NP. Part [ECCC:TR26-305] bounds depth-three Algebraic Monotone Calculus proofs by monotone tensor rank. Garlik, Gryaznov, Lu, Santhanam and Tzameret [ECCC:TR26-277] formalise rank-method circuit lower bounds in VNC^2. Under an algebraic Krajíček–Razborov analogue, they suggest that general circuits will need stronger reasoning. So the same week has interpolation for algorithms, interpolation for lower bounds, and rank for both.

My reading. Attention has moved from proving to auditing, and the audits are written in the language of the thing audited. The constants have been extracted, and the gap is large enough that Majdoddin calls it structural. If limits on these methods come, they will likely look like the rank floor Stein Naves de Brito computed. Watch the question he leaves open, whether $\rho \ge 1/2$.

Meanwhile, with nobody looking, Bun and Juvekar [arXiv:2610.37634] showed $\mathrm{UPP}^{\mathrm{cc}}$ is not closed under intersection. Bhangale, KM and Rubin-Toles [arXiv:2610.37869] exponentially improved the hardness of colouring 2-colourable 3-uniform hypergraphs. An ordinary week would have led with either. Nothing above is verified. Everything is filed.

— Lew Erbound, Night Desk — still waiting for an input that could actually be written down.

IN BRIEF

everything else that appeared this week

Meta-Mathematics of Algebraic Complexity

Michal Garlik, Svyatoslav Gryaznov, Jiaqi Lu, Rahul Santhanam et al. [ECCC:TR26-277]

Garlik, Gryaznov, Lu, Santhanam and Tzameret [ECCC:TR26-277] show that rank-method algebraic circuit lower bounds, including the constant-depth bounds of Limaye, Srinivasan and Tavenas and of Forbes, can be formalised in the bounded arithmetic theory VNC^2, while Polynomial Calculus Resolution cannot efficiently prove superpolynomial algebraic circuit lower bounds for any family of polynomials.

Improved NP-Hardness of Approximation for Max-3-Cut and for Max-k-Cut with Small k

Ziang Ni [arXiv:2611.01402]

Ni [arXiv:2611.01402] proves that Max-3-Cut is NP-hard to approximate within $27/29 + \varepsilon \approx 0.931$, improving the $16/17$ of Austrin, O'Donnell, Tan and Wright, though unlike theirs the new result does not hold for 3-colourable graphs; the paper also improves NP-hardness of Max-k-Cut for every k from 4 to 20, using gadget reductions from Håstad's linear equations.

Enumeration Complexity of Permutation Classes

Manuel Bodirsky, Mohamed Elsheikh [arXiv:2610.37940]

Bodirsky and Elsheikh [arXiv:2610.37940] study the problem of listing all permutations of length n that avoid a fixed finite set of patterns, which is output-polynomial in every case, and seek polynomial delay by building permutations through repeated insertion of a new maximum, which reduces the question to an associated insert-maximum extension problem.

The Complexity of Choosing Parity Projections

Karthik Sheshadri, Nicholas Smirnov [arXiv:2611.00289]

Sheshadri and Smirnov [arXiv:2611.00289] show that finding a smallest cube with parity sum 1 can be done in $2^{O(k^2)}\mathrm{poly}(N)$ time for cube size k, whereas the sum-0 analogue is not fixed-parameter tractable unless $\mathrm{FPT} = \mathrm{W}[1]$, and that without a bound on cube size both problems sit at the $\Sigma_2^{\oplus\mathrm{P}}$ level.

Separations between #BIS and #IS

Ijay Narang [arXiv:2610.38164]

Narang [arXiv:2610.38164] gives complexity-theoretic separations between approximating the hard-core partition function $Z_G(\lambda)$ on Δ-regular bipartite graphs (#BIS) and on general Δ-regular graphs (#IS).

Factor Closure in Noncommutative Algebraic Complexity

V. Arvind, Pushkar S. Joglekar [arXiv:2611.02039]

Arvind and Joglekar [arXiv:2611.02039] prove that VP_{nc} is closed under taking factors over every field, using a small weighted finite automaton: for a size-s noncommutative circuit of degree at most D, left and right factors have circuits of size $O(sD^3)$ and central factors have circuits of size $O(sD^6)$.

Upper and Lower Bounds on the Complexity of Unit-Distance Graph Realization in $\mathbb{Z}^d$

Eric Binnendyk [arXiv:2610.36516]

Binnendyk [arXiv:2610.36516] shows, assuming ETH, that deciding unit-distance realisability of an n-vertex graph requires $2^{\Omega(\sqrt{n})}$ time in $\mathbb{Z}^2$ and $2^{\Omega(n^{2/3})}$ in $\mathbb{Z}^3$, and gives a $2^{O(n^{(d-1)/d}(\log n)^2)}$ divide-and-conquer algorithm in $\mathbb{Z}^d$, so the bounds match up to a logarithmic factor in dimensions 2 and 3.

Extending OpenAI's cubic lower bound for the permanent to odd characteristic, Pfaffian representations, and polynomial entries

Karthik Sheshadri [ECCC:TR26-287]

Sheshadri [ECCC:TR26-287] extends the cubic lower bound on the border determinantal complexity of the permanent, previously proved only over $\mathbb{C}$, to every algebraically closed field of odd characteristic, and also proves the bound for Pfaffian representations.

Cubic lower bounds for the permanent in odd characteristic, for Pfaffian representations, and for polynomial entries

Karthik Sheshadri [arXiv:2610.40484]

In a stand-alone version of the same work, Sheshadri [arXiv:2610.40484] proves that the cubic determinantal-complexity lower bound for the permanent, which holds even for border representations, carries over to every algebraically closed field of odd characteristic, to Pfaffian representations and to matrices with polynomial entries.

Polynomial Identity Testing and Lower Bounds for Sparse Affine Orbits

Vishwas Bhargava, Akash Sengupta [ECCC:TR26-283]

Bhargava and Sengupta [ECCC:TR26-283] study polynomial identity testing and lower bounds for polynomials that are sparse after some invertible affine change of variables, seeking small explicit hitting sets for such polynomials and explicit polynomials that have many monomials in every affine coordinate system.

Irreducibility Testing of Sparse Polynomials over Integers is NP-hard and Probabilistic Polynomial Time

Yichuan Cao, Xiao-Shan Gao, Qiao-Long Huang, Jiawei Nie [arXiv:2610.37179]

Cao, Gao, Huang and Nie [arXiv:2610.37179] prove that testing irreducibility of sparse polynomials in $\mathbb{Z}[X]$ is coNP-hard under deterministic polynomial-time many-one reductions, and that the problem nevertheless lies in $\mathsf{PP} \subseteq \mathsf{PSPACE}$.

Classical Verification of Quantum Computation Beyond the Polynomial Hierarchy

Atul Mantri [arXiv:2610.37635]

Mantri [arXiv:2610.37635] shows that, relative to a classical oracle, a language outside the polynomial hierarchy that is decided exactly in quantum polynomial time admits classical verification with an efficient quantum honest prover; the public-coin protocol has perfect completeness and soundness error at most 1/4, and no language with such a proof system was previously shown to lie outside AM relative to an oracle.

On the complexities of Interactive Proofs of Proximity

Noga Amit, Oded Goldreich [ECCC:TR26-286]

Amit and Goldreich [ECCC:TR26-286] study interactive proofs of proximity in which the verifier's computation is left unrestricted, so that query complexity and communication complexity are the measures that count, and examine those complexities directly.

The Exact Communication Threshold for Parallel Repetition

Daiki Suruga [arXiv:2611.01339]

Suruga [arXiv:2611.01339] proves that, for every finite total relation and every input distribution, the probability of winning all parallel repetitions decays exponentially at any fixed communication rate below the zero-error internal information complexity, identifying the exact communication threshold for parallel repetition.

Reachability in Five-Dimensional Vector Addition Systems Is PSPACE-Complete

Weijun Chen, Bo Fu, Yuxi Fu [arXiv:2610.40012]

Chen, Fu and Fu [arXiv:2610.40012] prove that reachability in five-dimensional vector addition systems is PSPACE-complete, settling the complexity of a fixed-dimensional case whose picture had been largely unclear.

An FPRAS for Mixed Volumes Xiaoyu Chen, Kuikui Liu [arXiv:2611.01368]

Chen and Liu [arXiv:2611.01368] claim the first FPRAS for the mixed volume of n arbitrary compact convex sets in $\mathbb{R}^n$ given by weak membership oracles, which would disprove a 2009 conjecture of Gurvits; the authors credit the main ideas to an AI system and present the paper as a human-written exposition, so it awaits checking.

Computational complexity of the H-coloring problem in cubic graphs Vahan Mkrtchyan [arXiv:2610.35794]

Mkrtchyan [arXiv:2610.35794] shows that S_6-colouring of connected cubic graphs is NP-complete, disproving a conjecture of Hakobyan and the author, and gives an $O^*(\sqrt{8}^n)$ algorithm for H-colouring.

Improved Quasi-Polynomial-Time Approximation for Maximum Independent Set Zhao Song [arXiv:2610.35228]

Song [arXiv:2610.35228] gives a deterministic algorithm that, on an N-vertex graph with an independent set of size $2^{-t}N$, finds one of size 2^t in time $N^{O(t)}$, giving an $N/2^{\Omega(\sqrt{\log N})}$ approximation in time $N^{O(\sqrt{\log N})}$ and improving the exponent $1/3$ of Lovett to $1/2$; Tulsiani's integrality gaps limit this form of analysis.

Treedepth is NP-Hard on Planar Graphs of Maximum Degree 3 Kolja Kühn [arXiv:2610.42033]

Kühn [arXiv:2610.42033] proves that computing treedepth is NP-hard under polynomial-time Turing reductions even on planar graphs of maximum degree 3, which is best possible since degree 2 is easy.

On the complexity of isomorphism problems for tensors, groups, and polynomials V: over commutative rings Joshua A. Grochow, Youming Qiao, Katherine E. Stange, Xiaorui Sun [arXiv:2611.00735]

The full version of Grochow, Qiao, Stange and Sun's STOC'25 paper [arXiv:2611.00735] sets up complexity classes for tensor isomorphism over commutative rings, shows that the classification of Calabi–Yau threefolds is complete for the corresponding class over $\mathbb{Z}$, and gives an $N^{O((\log N)^{8/9})}$ isomorphism test for odd-order class-2 nilpotent groups.

Linear Hashing is Poissonesque Vinayak Kumar [ECCC:TR26-284]

Kumar [ECCC:TR26-284] shows that when balls are hashed into bins by a random $\mathbb{F}_2$-linear map the bin loads still follow a Poisson law, in a necessarily weaker sense, so that, for example, the expected maximum load is within a $3/2 + o(1)$ factor of fully random hashing.

Linear Hashing Is Poissonesque Vinayak M. Kumar [arXiv:2610.34750]

The arXiv version of Kumar's linear hashing paper [arXiv:2610.34750] gives the same Poisson-type description of bin loads under random $\mathbb{F}_2$-linear maps, together with a strengthened Leftover Hash Lemma for linear maps.

Half-Optimal Induced Matchings in Random Graphs: Online and Low-Degree Algorithms Abhishek Dhawan, Eren C. Kızıldağ, Neeladri Maitra, Cheng Mao et al. [arXiv:2610.40369]

Dhawan, Kızıldağ, Maitra, Mao and Wee [arXiv:2610.40369] show that in $G(n,p)$ with $C/n \le p \le 1-\kappa$ online and low-degree algorithms find induced matchings of only about half the optimal size, a sharp factor-two gap of the kind familiar from independent set.

Exponential quantum-classical query separations for cyclic Forrelation Yuxin Zhang, Changpeng Shao [arXiv:2610.41694]

Zhang and Shao [arXiv:2610.41694] give nearly tight nonadaptive randomised classical query bounds for cyclic Forrelation, with an exponent that depends on the parity of k, and a polylogarithmic-query algorithm at $k = 3$ that disproves a conjectured lower bound.

Catalytic space-efficient algorithms for Distance and Determinant Srijan Chakraborty, Samir Datta, Kushagra Shandilya [arXiv:2610.41723]

Chakraborty, Datta and Shandilya [arXiv:2610.41723] extend recent sublinear-catalytic-space, polynomial-time reachability algorithms to shortest distance, and give near-linear catalytic space algorithms for the determinant.

The Runtime-Moment Frontier for Universal Sampling Sebastian Ben Daniel [arXiv:2611.01884]

Ben Daniel [arXiv:2611.01884] shows that for every oracle and rational $a > 0$ a single sampler with polynomial ath runtime moment dominates every source with polynomial rth moment with power $B(a,r) = 1 + (a - \min\{r,1\})_+/r$, and that this power is optimal relative to an oracle.

Explicit Reed-Solomon Subspace Designs: Optimal List and Line Decoding Omar Alrabiah, Rohan Goyal, Venkatesan Guruswami [arXiv:2610.36215]

Alrabiah, Goyal and Guruswami [arXiv:2610.36215] give an explicit $n^{O_\varepsilon(1)}$-time construction of Reed–Solomon codes over polynomial-size extension fields that are list-decodable from a $1 - R - \varepsilon$ fraction of errors with optimal list size $O((1-R)/\varepsilon)$, plus a simpler strongly explicit construction with list size $(1/\varepsilon)^{O(1/\varepsilon)}$.

Decomposing Directed Graphs into Partial Orders is Hard Grzegorz Gutowski, Dominik Jilg, Antonio Lauerbach, Mikołaj Rams [arXiv:2610.37379]

Gutowski, Jilg, Lauerbach and Rams [arXiv:2610.37379] prove that deciding whether a directed acyclic graph can be covered by k strict partial orders is NP-complete for every fixed $k \ge 2$.

Mutual-visibility colouring with two colours is NP-complete for every fixed diameter Siddhesh Chaubal [arXiv:2611.01012]

Chaubal [arXiv:2611.01012] shows that deciding whether a graph can be partitioned into two mutual-visibility sets is NP-complete for graphs of every fixed diameter $d \ge 2$, extending the known diameter-four result, via a reduction from 2-colouring 3-uniform hypergraphs.

Mathematics and Efficiency of Song's CVP Hardness Reduction Shuhong Gao [arXiv:2610.39211]

Gao [arXiv:2610.39211] gives an expository and more efficient version of Song's $N^{1-\varepsilon}$ NP-hardness for the closest codeword problem, which also yields $N^{1/2-\varepsilon/2}$ hardness for Euclidean CVP, cutting the construction's dimension from $n^{8/\varepsilon^2+O(1/\varepsilon)}$ to $n^{4/\varepsilon^2+O(1/\varepsilon)}$ and extending it to any finite field.

On the NP-Hardness of Rank-Metric Minimum Distance and Syndrome Decoding, at Every Rate Venkata Gandikota, Vatsal Jha, S. Venkitesh [arXiv:2610.40825]

Gandikota, Jha and Venkitesh [arXiv:2610.40825] derandomise the NP-hardness of minimum distance and syndrome decoding for $\mathbb{F}_{q^m}$-linear rank-metric codes of length N with $m = O(N^{1+\varepsilon})$, and show that hardness holds at every fixed rate.

An Improved Johnson Bound for List-Decoding with Hints Swastik Kopparty, Harry Sha [ECCC:TR26-282]

Kopparty and Sha [ECCC:TR26-282] show that a code of distance δ is list-decodable beyond the Johnson radius $J(\delta)$ when the received word is known to have some fairly close codewords (hints), characterise the optimal radius, and decode Reed–Solomon codes efficiently up to it.

Matrix Samplers and Quantum-Proof Extractors: Their Correspondence and Almost-Optimal Constructions Kuan Cheng, Ruiyang Wu [arXiv:2610.37122]

Cheng and Wu prove that matrix samplers and quantum-proof randomness extractors are equivalent, a matrix analogue of Zuckerman's classical sampler-extractor correspondence, and use it to build explicit strong matrix samplers with optimal randomness overhead and near-optimal sample complexity.

Compressed Oracle Duality: A simple framework for knowing what quantum adversaries know Matthew Ding, Mark Zhandry [arXiv:2611.00609]

Ding and Zhandry introduce Compressed Oracle Duality, a framework for tracking what a quantum adversary querying a random oracle actually knows, which fixes the original compressed oracle's habit of recording information the adversary has since uncomputed, a defect that matters in the space-bounded setting.

A Decomposition Theorem and a Simple Algorithm for Contraction Fixed Points Xi Chen, Yuhao Li, Mihalis Yannakakis [arXiv:2611.01302]

Chen, Li and Yannakakis prove a decomposition theorem for computing approximate fixed points of ℓ_∞-contraction maps on $[0,1]^k$, which Fedorov and Haslebacher used to turn poly$(k, \log(1/\varepsilon))$-query algorithms running in $(k\log(1/\varepsilon))^{O(k)}$ time into $(k\log(1/\varepsilon))^{O(\sqrt{k})}$-time algorithms, and they give a simpler algorithm.

Quantum Lower Bounds for Collision Problems Pan Peng, Jingyu Wu [arXiv:2611.01767]

Peng and Wu prove tight quantum query lower bounds for finding a k-to-one collision in an exactly k-to-one string: $\Omega_k(n^{1/2-1/(2(2^k-1))})$ for adaptive algorithms, matching Liu and Zhandry, and $\Omega_k(n^{1-1/k})$ for one-round non-adaptive algorithms, matching classical birthday sampling.

The Complexity of Value Optimization in Sporadic Task Systems Christopher Harbourne, Nathan Fisher, Kecheng Yang, Pontus Ekberg et al. [arXiv:2610.34725]

Harbourne, Fisher, Yang, Ekberg, Baruah and Guo show that choosing sporadic real-time tasks to reach a value target while meeting every deadline on one preemptive processor is Σ_2^P-complete, for both constrained and arbitrary deadlines, even with all numbers in unary.

Directed Cut-Query Complexity of Finding a King Ziad Ismaili Alaoui, Oded Lachish, Nikhil S. Mande [arXiv:2610.40521]

Ismaili Alaoui, Lachish and Mande show that finding a king in an n-vertex tournament takes $\Theta(n)$ deterministic directed cut queries, first proving that any correct deterministic algorithm in this model must output a vertex of maximum out-degree.

A Numerical Refinement of an Exact Discrete Fourier Transform Bound Zhihao Yao, Sichao Liu [arXiv:2611.01257]

Yao and Liu tighten the numerical bounds in Kartal's instantiation of Sussman's exact discrete Fourier transform construction, giving work $O(n(\log n)^{1-\delta})$ with $\delta = 0.0030201659$ in an exact complex-arithmetic RAM model, an improvement of 5.9×10^{-9} on the previous parameter.

Impossibility of No-Regret in Multi-Agent Bidding Zongjun Yang, Rachitesh Kumar, Christian Kroer [arXiv:2610.36666]

Yang, Kumar and Kroer show that in repeated second-price auctions with i.i.d. item arrivals it is computationally impossible for all budget- or ROI-constrained pacing bidders to achieve no-regret at the same time, and they show that high variance in online bidding forces high regret.

What Can Be Computed with Indefinite Causal Structure? Omri Shmueli [arXiv:2610.40573]

Shmueli studies quantum computation without a definite causal order in the process-matrix formalism, with cost measured by the size of the circuit generating the process, and gives polynomial-time algorithms in two settings: for NP with general process matrices and for unique-sink orientation with pure processes.

Clifford synthesis is NP-complete Søren Fuglede Jørgensen [arXiv:2610.40886]

Jørgensen proves that synthesising a given Clifford operator within a resource budget is NP-complete for several practical models, including counting CNOT gates with single-qubit gates free, assuming full connectivity and no ancillas.

FIXP-Completeness in Second-Price Auctions Nobutaka Shimizu [arXiv:2611.00543]

Shimizu proves that exact equilibrium computation is FIXP-complete in three second-price auction models, including pacing with return-on-investment constraints at multiplier cap $1+\delta$ for every fixed rational $\delta > 0$ and budget-constrained generalised second-price sponsored shopping.

Approximate clearing in financial networks with credit default swaps: hardness thresholds, a box algorithm, barriers, and exact solvency at bounded degree Michal Čertík [arXiv:2611.00855]

Čertík proves PPAD-hardness of weak approximate clearing in financial networks with credit default swaps below a threshold of about 0.2530, and gives a polynomial-time interval algorithm above about 0.2834; on networks whose swap writers are solvent ex ante the thresholds are about 0.1577 and 0.1716.

The Value of Steiner Hubs in Two-Hop Distance Labeling Sichen Wang [arXiv:2611.01106]

Wang determines how much Steiner hubs help two-hop distance labelings on ℓ_1 products of betweenness-free metrics, a class that includes every Hamming graph: they give no saving when the factors have at most three letters and a polynomial saving from four letters upwards.

The Movement Complexity of Sorting Sichen Wang [arXiv:2611.01157]

Wang determines, for every grid dimension d and number of rounds r, the expected total head movement needed to sort n keys on a d-dimensional grid with two comparison heads, showing it is n^β up to polylogarithmic factors.

Deterministic Polynomial-Time Primality Testing with a Single Binomial Congruence Joseph M. Shunia [arXiv:2611.02280]

The arXiv version of Shunia's deterministic polynomial-time primality test built on the single binomial congruence $(1+X)^n = 1+X^n$; an earlier version appeared as a revision of ECCC TR23-200.

Quantum Query Complexity of Tournament Solutions Pragya Arora, Udit Kumar [arXiv:2610.39702]

Arora and Kumar study the quantum query complexity of tournament solutions and show that the top cycle needs only $\tilde{\Theta}(N)$ quantum queries, whereas classically it, the Copeland set and the uncovered set all need $\Theta(N^2)$.

Ancilla-free randomized Toffoli gates and the T-count of Boolean oracles Arul Rhik Mazumder [arXiv:2610.38189]

Mazumder shows that computing the sampled parities in place, on the gate's own controls, removes every ancilla from the Gosset-Kothari-Zhang randomised Toffoli construction once the gate is large enough, with each random circuit using about 8 T gates per bit of precision.

Quantum query complexity of Hurwitz stability testing

Maxwell Fitzsimmons, Mohamed Serry, Meenu Kumari, Jun Liu [arXiv:2610.41180]

Fitzsimmons, Serry, Kumari and Liu characterise the quantum query complexity of deciding whether a matrix is Hurwitz stable, given an exact normalised block-encoding, in terms of the dimension n and a supplied lower bound ν, including an $\Omega(\sqrt{n}/\nu)$ lower bound.

Lower Bound for Region of Non-Distillable Magic States

Akul Sethi [arXiv:2610.37807]

Sethi extends Campbell and Browne's proof that bound magic states exist outside the stabiliser octahedron, and uses it to give a lower bound on the region of single-qubit states that cannot be distilled for universal quantum computation.

Diagonalization and Undecidability in Semi-Thue Systems

Vesa Halava, Reino Niskanen [arXiv:2610.40105]

Halava and Niskanen recast classical diagonalisation in terms of derivation trees of semi-Thue systems and prove, by direct Post-style diagonalisation, that the complements of the assertion and word problems are not recursively enumerable.

Retrieval Sequencing in a Compact Storage System with Side Access

Yağmur Gül [arXiv:2610.40169]

Gül studies how to order retrievals in a compact storage system with side access so as to minimise the number of lifted bins, and proves that for every number of stacks the total energy splits into an instance constant plus a pairwise penalty.

Stochastic and Noisy Oracles in Smooth Continuous Optimization

Albert S. Berahas, Liyuan Cao, Katya Scheinberg [arXiv:2610.41086]

Berahas, Cao and Scheinberg survey deterministic and stochastic models of inexact zeroth-, first- and second-order oracles in smooth continuous optimisation, and what they imply for convergence, computational cost and attainable accuracy.

Leaving Quantum Bias on the Table: Pushforwards as a Path to Quantum Advantage in Constrained Optimization

Chinonso Onah, Stuart Hadfield, Kristel Michielsen [arXiv:2611.00759]

Onah, Hadfield and Michielsen propose a post-measurement framework that builds classical repair maps from constraint-violation statistics and uses them to map infeasible outputs of constrained quantum samplers onto feasible solutions, deriving exact identities for the resulting success probability and expected cost.

FROM THE AI-ASSISTED REPOSITORIES

Work deposited with AI systems credited as authors or co-authors, not screened like arXiv or ECCC. Read with care.

A $2^{o(n)}$ Time Algorithm for 3-SAT

AI-ASSISTED Anthropic [hexagon:2610.00577v1]

A manuscript credited only to Anthropic [hexagon:2610.00577v1] claims a deterministic $2^{o(n)}$-time algorithm for #3-SAT, which would refute ETH; it says the result was found by an internal Claude model and formalised in Lean 4, and its two-sentence abstract gives no technical detail.

The rank of 3×3 matrix multiplication over F2 is 23

AI-ASSISTED Tejasvi Singh Tomar [hexagon:2610.00548v1]

Tomar [hexagon:2610.00548v1], in an AI-assisted paper, proves with the substitution method that the rank of the 3×3 matrix multiplication tensor over $\mathbb{F}_2$ is exactly 23, matching Laderman's algorithm; this makes Laderman's algorithm optimal among bilinear algorithms over $\mathbb{F}_2$ and among those with integer coefficients.

Subexponential Weighted Counting and Exact Optimisation

AI-ASSISTED Deep Bhattacharjee [hexagon:2610.00594v1]

Bhattacharjee [hexagon:2610.00594v1], in an AI-assisted paper, reproduces OpenAI's claimed $2^{o(n)}$ #k-CNF algorithm and extends it to sums of local weights over commutative $\mathbb{Q}$-algebras; this would give $2^{o(n)}$ algorithms for exact Max-Cut and Max-k-SAT on sparse instances, maximum independent set and Vertex Cover in $2^{o(k)} + \mathrm{poly}(n)$, all conditional on the unverified original.

Deterministic Factoring in Time $N^{o(1)}$

AI-ASSISTED Isaac Rudich, Louis-Martin Rousseau · with Lean Lean, Anthropic Claude Opus [hexagon:2610.00600v1]

Rudich and Rousseau [hexagon:2610.00600v1] present an AI-assisted Lean certificate claiming deterministic factoring in $N^{o(1)}$ time, specifically $C_k 2^{n/k}$ steps for every k; the method reduces factoring to counting divisors with prescribed leading bits and relies on the counting engine of the claimed ETH refutation.

Counting Hamiltonian cycles in sparse graphs in subexponential time, assuming the new #SAT algorithms

AI-ASSISTED Vaibhav Tulsyan · with Anthropic Claude Fable, Anthropic Claude Opus [hexagon:2610.00672v1]

Tulsyan [hexagon:2610.00672v1], in an AI-assisted paper, derives a corollary of the Anthropic manuscript: assuming three of its internal results, the Hamiltonian cycles of a sparse graph can be counted in subexponential time, a case the claimed #k-SAT algorithms do not cover directly.

The partition exponent in the graded pieces of the Claude 3-SAT algorithm is optimal for pieces of that shape

AI-ASSISTED Carlos Stein Naves de Brito · with Anthropic Claude Fable, SymPy SymPy [hexagon:2610.00692v1]

In an AI-assisted note, de Brito [hexagon:2610.00692v1] argues that the rank floor stated in Appendix H of the claimed Anthropic subexponential 3-SAT algorithm already rules out partition-type growth exponents $\alpha < 1/2$, so the gain parameter $\rho = 1/\sqrt{2}$ is optimal for graded pieces of that shape.

A degree-three local-weight cap and an improved relevant-variable bound

AI-ASSISTED Ricardo Pinho · with OpenAI Codex [hexagon:2610.00514v1]

In an AI-assisted (Codex) preprint, Pinho [hexagon:2610.00514v1] caps the local-degree potential of Boolean functions of degree at most three at $5/4$ and, feeding this into Wellens's degree–block-sensitivity recurrence, lowers the constant to $C_{44} \le 4.3576549226$, an improvement of $1/128$ checked by two exact-arithmetic programs.

The Euclidean algorithm is not optimal among uniform processes

AI-ASSISTED Jason Davies · with OpenAI Codex, Anthropic Claude Code, Lean Lean [hexagon:2610.00604v1]

In an AI-assisted, Lean-formalised paper, Davies [hexagon:2610.00604v1] shows that $\gcd(a, b)$ can be derived from $a \ge b$ with $O(\log a / \log\log a)$ applications of the remainder function alone, refuting Moschovakis's conjectured logarithmic lower bounds for uniform processes, while the Main Conjecture for algorithms stays open.

Deterministic Polynomial-Time Primality Testing with a Single Binomial Congruence

AI-ASSISTED Joseph Shunia · with OpenAI ChatGPT [hexagon:2610.00697v1]

Shunia gives an AI-assisted deterministic polynomial-time primality test whose main check is the single congruence $(1 + X)^n = 1 + X^n$ in an extension ring, with the modulus built deterministically from Jacobi sums and the Cureton-Thorner zero-free theorem.

Normality of smooth lattice three-polytopes AI-ASSISTED

Mateusz Michalek · with Chat GPT Sol, Claude Fable and Opus [diderot:3d8534c6211a]

Michałek gives an AI-assisted proof that every smooth lattice three-polytope is normal, settling Oda's normality question in dimension three, by reducing to 2-normality and excluding the remaining cases with Howe's theorem, Pick's formula, Ball's plank theorem and a sharp 3×3 matrix inequality.

FOURTEEN POINTS IN THREE-SPACE AI-ASSISTED

Mateusz Michalek · with Chat GPT Astra and Sol, Anthropic Fable and Opus [diderot:67aceeceed8f]

In an AI-assisted paper, Michalek shows that every finite subscheme of length at most fourteen in complex affine three-space is a limit of reduced subschemes, so the Hilbert schemes of twelve, thirteen and fourteen points in $\mathbb{C}^3$ are irreducible.

Nineteen points in three-space AI-ASSISTED

Mateusz Michalek · with Chat GPT Astra and Sol, Claude Opus and Fable [diderot:4afe1493dff9]

Michalek's AI-assisted sequel uses linkage to extend smoothability to all finite subschemes of $\mathbb{C}^3$ of degree at most 19, so the corresponding Hilbert schemes are irreducible, with degrees up to 18 formally verified in Lean.

Twenty points in three-space AI-ASSISTED

Mateusz Michalek · with Chat GPT Astra and Sol, Claude Fable and Opus [diderot:ed3ef6793489]

Continuing the same AI-assisted induction, Michalek proves that every length-twenty subscheme of $\mathbb{C}^3$ is smoothable, so the Hilbert scheme of twenty points in $\mathbb{C}^3$ is irreducible.

Twenty-one points in three-space AI-ASSISTED

Mateusz Michalek · with Chat GPT Astra and Sol, Claude Fable and Opus [diderot:712abac1ba30]

Michalek's AI-assisted paper introduces linkage at infinity to prove that the Hilbert scheme of twenty-one points in $\mathbb{C}^3$ is irreducible, placing the least number of points with a reducible Hilbert scheme between twenty-two and seventy-eight.

Fixed-window golden-ratio subset sums and Fibonacci representation counts AI-ASSISTED

openai/gpt-6-astra, anthropic/claude-opus-5.5 · with openai/gpt-6-astra, anthropic/claude-opus-5.5 [diderot:4c32cc8f9f1c]

This AI-authored note proves that the number of subsets of $\{-2m, \dots, 2m\}$ whose golden-ratio powers sum to an integer t equals the number of representations of $F_{2m+2}t$ as a sum of distinct Fibonacci numbers for $0 \le t \le L_{2m+1}$, and is zero for all other t; it then uses this correspondence to attack the associated extremal problem.

Extremal coefficients of Pell representation polynomials and integer fibres in the silver ratio AI-ASSISTED

openai/gpt-6-luna, anthropic/claude-opus-5.5 · with openai/gpt-6-luna, anthropic/claude-opus-5.5 [diderot:4624e972e436]

This AI-authored note proves that the largest coefficient of the Pell representation polynomial $\prod_{j=2}^{\ell+1}(1 + x^{P_j} + x^{2P_j})$ is $2^{\lfloor(\ell-1)/2\rfloor}$ for every $\ell \ge 1$, and determines the full set of exponents at which it is attained.

ODDS AND ENDS

Complexity Chronicles

NUMBERS OF THE WEEK

$9/4$	2^{420}	10^9	$27/29$	23	5.9×10^{-9}
CLAIMED BOUND ON ω, NOW WITH AN ELEMENTARY PROOF	CLAUSES BEFORE THE #3-SAT RECURSION WINS, AT DENSITY ONE	WHAT $\log_2(n+m)$ MUST EXCEED BEFORE ONE CERTIFIED BOUND BEATS 2^n	NEW NP-HARDNESS FACTOR FOR MAX-3-CUT, DOWN FROM $16/17$	RANK OF 3×3 MATRIX MULTIPLICATION OVER $\mathbb{F}_2$ (AI-ASSISTED)	IMPROVEMENT IN AN EXACT DFT EXPONENT SAVING

COMPLEXITY FORECAST

Fine-grained: ETH low, falling, crossover points remaining astronomical; brute force good, occasionally very good. Proof complexity: Sherali-Adams weakening, Res(2) moderate, interpolation spreading from the east. Algebraic: ω steady at nine-quarters, expositors expected in all sea areas. Lean: certificates plentiful, human understanding poor, becoming some later.

CLASSIFIEDS

- FOR SALE: subexponential 3-SAT algorithm. Overtakes brute force from about 2^{420} clauses. Buyer collects.
- WANTED: constants, any size. Preprint will ask no questions.
- LOST: automatability of Sherali-Adams. Not recoverable unless $\mathrm{P} = \mathrm{NP}$. No reward offered.
- FOR SALE: Laderman's algorithm, 23 multiplications. Now certified optimal over $\mathbb{F}_2$. One careful owner.
- TO LET: 2-colourable 3-uniform hypergraph. Decorators quoting $O((\log n)^{1/8})$ colours need not apply.
- FOR SALE: matrix, sign rank $\Theta(d)$. Do not AND with itself.
- SITUATION VACANT: human reader for 13-line Lean statement on factoring. Understanding not supplied.
- EXCHANGE: exponent saving of 5.9×10^{-9}. Will swap for a hot coffee.

Complexity Chronicles is written by a language model (anthropic: claude-opus-5-5) from the abstracts of 85 papers (arXiv cs.CC, ECCC, Hexagon, Diderot) that appeared during 26 October – 1 November 2026, selected according to a reader profile. Generated 2026-11-06 08:59 UTC. Headlines are headlines; standfirsts are meant to be accurate; nothing here replaces reading the paper. Typeset with LuaLaTeX in Merriweather, Montserrat, Playfair Display, Roboto Condensed and Fira Mono.