PRICE: ONE ORACLE QUERY
WEEKLY · P ≠ NP (PROBABLY)
VOL. 1 NO. 46
2026-W46

the complexity chronicles

Every bound stated, every assumption named.

MONDAY 16 NOVEMBER 2026 · COVERING 9–15 NOVEMBER 2026 · 75 PAPERS · ARXIV CS.CC, ECCC, HEXAGON, DIDEROT, ZENODO

INSIDE THIS ISSUE

p.2: Two analyses of the claimed subexponential 3-SAT algorithms. One moves the savings to NTIME; the other moves them to proof complexity. p.2: Deciding whether a finite structure has a Maltsev polymorphism is in polynomial time. Rectangularity carries the proof. p.3: If NP ⊆ P/poly, the hierarchy collapses to a class inside every earlier collapse class. p.4: How the Razborov–Smolensky polynomial method reached $AC^0[2]$-Frege. Column: Shell Scott, Taking Inventory.

PROOF COMPLEXITY

First hard instances for AC0[2]-Frege at every depth

For every fixed $d \geq 2$, explicit depth-d formulas need superpolynomial-size semantic depth-d $AC^0[2]$-Frege refutations. Semantic depth-3 and depth-4 polynomial calculus over $\mathbb{F}_2$ also receive exponential and superpolynomial lower bounds, which settles the Res(log) question for the weak pigeonhole principle.

Alekseev, Gryaznov, Itsykson, Pitassi, Robere and Tzameret prove size lower bounds against a ladder of algebraic-Boolean proof systems [ECCC:TR26-316]. Every bound is for the semantic version of the system, over $\mathbb{F}_2$.

The first rung is ΣΠΣ-PC, polynomial calculus whose lines are depth-3 circuits with a sum gate on top. Here the bounds are exponential: for the bit pigeonhole principle, and $2^{\Omega(n)}$ for random 3-CNFs on n variables and for the bit and functional pigeonhole principles with $n+1$ pigeons and n holes. The second rung is ΣΠΣΠ-PC with unrestricted bottom fan-in. Here the bounds are superpolynomial, for the weak unary pigeonhole principle and for lifted random CNFs of logarithmic width. With $2n$ pigeons and n holes the bound is $\exp(\Omega((\log n)^2/\log\log n))$. This answers the open question of a superpolynomial Res(log) lower bound for that principle. If the bottom fan-in is held to logarithmic, the bounds for random CNFs of logarithmic width become exponential.

The top rung is $AC^0[2]$-Frege. For each fixed $d \geq 2$ the authors build unsatisfiable formulas of depth d. These need superpolynomial-size semantic depth-d refutations, yet have polynomial-size refutations at depth $d+1$. They are the first unconditional hard instances against $AC^0[2]$-Frege of arbitrary depth. The instances are deep formulas, not CNFs. For CNFs, the strongest system the paper reaches is depth-four polynomial calculus.

The method starts from Braun's recent AI-generated proof of an exponential lower bound for resolution over parities. The authors distil the annihilating-polynomial argument implicit in that proof into a general framework built on pseudo-reduction operators. Within it, rank lower bounds for these operators give proof-size lower bounds. The Razborov–Smolensky polynomial approximation method extends the framework to bounded bottom degree, and lifting with random restrictions removes the bound on bottom fan-in. The $AC^0[2]$-Frege result adds a lifting construction based on constant-depth formulas for the coin problem.

The weak pigeonhole bound is superpolynomial, which is what the question asked for. It is not exponential.

[ECCC:TR26-316] Hard Instances for Strong Semantic Proof Systems: CNFs for Depth-Four PC and Deep Instances for AC0[2]-Frege (Yaroslav Alekseev, Svyatoslav Gryaznov, Dmitry Itsykson, Toniann Pitassi et al.)

FINE-GRAINED

Claimed 3-SAT savings carry over to NTIME

If the subexponential 3-SAT algorithms reported by OpenAI and Anthropic are correct, then $\mathrm{NTIME}[t] \subseteq \mathrm{DTIME}[2^{O(t\log\log t/\sqrt{\log t})}]$, and a stronger consequence of their proof gives $\mathrm{NTIME}[t] \subseteq \mathrm{DTIME}[2^{O(t/\sqrt{\log t})}]$ for multitape machines.

Strassle and Subercaseaux show that the 3-SAT algorithms recently reported by OpenAI and Anthropic imply $\mathrm{NTIME}[t] \subseteq \mathrm{DTIME}[2^{O(t\log\log t/\sqrt{\log t})}]$ [arXiv:2611.11832]. A stronger conse-

quence of the same proof removes the $\log \log t$ factor. More generally, any subexponential algorithm for circuit-CSPs becomes a deterministic simulation of nondeterministic time, with the same savings.

The argument rests on new generalisations of the Cook–Levin theorem, which the authors offer as of independent interest. These carry a saving on circuit-CSPs over to multitape nondeterministic Turing machines without losing it.

The result is conditional in the plain sense: it holds if the reported algorithms are correct. The saving in the exponent is a factor of order $\sqrt{\log t}$.

[arXiv:2611.11832] Simulating Nondeterministic Turing Machines in Subexponential Deterministic Time (Carmen Strassle, Bernardo Subercaseaux)

VALUED CSP

One number fixes the treewidth cost of valued CSPs

For finite-valued VCSPs whose language contains all crisp unary constraints, instances of treewidth tw can be solved in time $(c_\Gamma)^{\mathrm{tw}}$ times a polynomial. When $c_\Gamma \geq 2$, no $(c_\Gamma - \varepsilon)^{\mathrm{tw}}$ algorithm exists unless the Primal Treewidth SETH, and hence CNF-SETH, fails.

Greilhuber, Marx, Osipov, Piecyk, Rzążewski and Živný study valued CSPs over a fixed finite domain [arXiv:2611.16729]. The language contains every crisp unary constraint, called a hard list, and all its other cost functions are finite-valued. The parameter that matters is c_Γ. This is the size of the largest set S on which the language can express a binary disequality, using the usual VCSP closure operations.

Given a tree decomposition of width tw of the primal graph, the problem is solvable in time $(c_\Gamma)^{\mathrm{tw}} \cdot (|I| + |\mathcal{T}|)^{O(1)}$. For $c_\Gamma \geq 2$, the base cannot be lowered to $c_\Gamma - \varepsilon$ for any $\varepsilon > 0$ unless the Primal Treewidth SETH fails. This generalises a result of Esmer et al. on the optimisation version of list homomorphism for graphs.

The lower bound needs only the Primal Treewidth SETH, which is weaker than CNF-SETH. That choice matters in a season of reported ETH refutations. Because SETH implies ETH, a refutation of ETH would also refute CNF-SETH. Whether the weaker hypothesis survives is a separate question.

[arXiv:2611.16729] Tight Algorithms for Conservative Finite-Valued CSPs Parameterized by Treewidth (Jakob Greilhuber, Dániel Marx, George Osipov, Marta Piecyk et al.)

PARAMETERISED SAT

Forest backdoors are FPT; treewidth two is not

A strong backdoor of size at most k to the class of incidence forests can be found, or shown not to exist, in time $2^{2^{O(k)}}|\varphi|$ on any non-tautological CNF. For treewidth $t \geq 2$ the problem is W[1]-hard even on 3-CNFs, and any better dependence on k would give a $2^{o(N/\log N)}$-time 3-SAT algorithm.

In 2013 Gaspers and Szeider gave an FPT approximation for strong backdoors to bounded incidence treewidth. When a backdoor of size k exists, their algorithm finds one of size at most 2^k, in time $f(k,t)|\varphi|^3$. Whether exact detection is FPT in $k + t$ was left open, even for forests and even for 3-CNF inputs. Bilyk, Munduzbaev, Sharma and Siam settle it [arXiv:2611.12321].

For forests the answer is yes. On arbitrary non-tautological CNFs, their algorithm runs in time $2^{2^{O(k)}}|\varphi|$, which is linear in the input. For treewidth $t \geq 2$, the problem is W[1]-hard, even on non-tautological 3-CNFs.

The double-exponential dependence on k is also shown to be needed. A better one would give a $2^{o(N/\log N)}$-time algorithm for 3-SAT on N variables. The authors note that this would beat even the recently claimed subexponential 3-SAT algorithms.

[arXiv:2611.12321] Strong Backdoor to Forests is FPT: Tight Linear-Time FPT for Essentially General SAT (Olesia Bilyk, Aidar Munduzbaev, Roohani Sharma, Mamnoon Siam)

IN DEPTH 2026-W46

FINE-GRAINED COMPLEXITY AND PROOF COMPLEXITY

What subexponential 3-SAT buys

If the reported OpenAI and Anthropic algorithms are correct, then every multitape nondeterministic machine running in time t can be simulated deterministically in time $2^{O(t/\sqrt{\log t})}$. Separately, under three stated assumptions from the Anthropic preprint, every unsatisfiable 3-CNF on n variables has an Ideal Proof System refutation of size $2^{o(n)}$ over any field of characteristic other than 2. Known lower bounds rule out the same for resolution, Res(k), polynomial calculus and sum-of-squares. Neither result beats brute force at any size that could be run.

Two preprints, one from OpenAI and one from Anthropic, report deterministic algorithms that count the satisfying assignments of a 3-CNF on n variables in time $2^{o(n)}$. If they are correct, ETH is false. Neither preprint has been refereed. This week two papers ask what follows. One looks at nondeterministic time. The other looks at proof complexity and knowledge compilation. Both are conditional, and both say so.

WHY COOK–LEVIN WAS NOT ENOUGH

The first instinct is to push NTIME[t] through Cook–Levin and run the new algorithm on the result. Strassle and Subercaseaux [arXiv:2611.11832] explain why this fails. The best known Cook–Levin reduction produces a 3-CNF with $O(t \log t)$ variables. A $2^{O(n/\sqrt{\log n})}$ algorithm on that many variables gives time $2^{O(t\sqrt{\log t})}$, which is worse than the trivial $2^{O(t)}$. Making Cook–Levin linear is a long-standing open problem.

Their answer is to avoid needing it. They prove NTIME[t] $\subseteq$ DTIME[$2^{O(t/\sqrt{\log t})}$], given a version of the reported algorithms that solves circuit-CSPs: constraint problems in which each constraint is an arbitrary Boolean circuit on its variables. From the 3-SAT statement alone they get the weaker $2^{O(t \log\log t/\sqrt{\log t})}$.

More generally, any subexponential saving for circuit-CSPs of bounded density and suitable arity transfers to NTIME with the same saving.

> *There are $2^{O(t/b)}$ possible patterns, and the simulation tries them all.*

This answers, conditionally, a conjecture attributed to Lipton that NTIME$[t] \subseteq$ DTIME$[2^{\varepsilon t}]$ for every $\varepsilon > 0$. There is an oracle relative to which that conjecture fails, so the argument cannot relativize. The authors note that the statement also fails relative to a random oracle.

CUTTING THE RUN INTO EPOCHS

The idea is to encode the computation in pieces. Split the run into epochs of length b and each tape into blocks of b cells. By a classical result of Hopcroft, Paul and Valiant, the machine may be assumed to move its heads between blocks only at epoch boundaries. Once the pattern of block moves is fixed, each epoch becomes one constraint circuit of size $O(b \log b)$. That circuit links the block contents before and after the epoch with the b nondeterministic bits used during it. Each variable is shared by at most two constraints. There are $2^{O(t/b)}$ possible patterns, and the simulation tries them all.

Each pattern gives a circuit-CSP with $O(t)$ variables, arity $O(b)$ and density at most 2. Converting each circuit to clauses gives a 3-CNF with $O(t \log b)$ variables rather than $O(t \log t)$. With $b = t$ the classical reduction comes back. The authors compare the construction to the sparsification lemma. The cost is a family of $2^{O(t/b)}$ formulas instead of one. That is affordable precisely when b is large compared with the savings of the solver.

REFUTATIONS FROM COUNTING IDENTITIES

Szeider [zenodo:24243988] starts from the identities the Anthropic preprint uses to prove its counting algorithm correct, not from the algorithm itself. He takes three statements from that proof as assumptions. Under them, he shows that the weighted model count of every 3-CNF on n variables has an arithmetic circuit of size $2^{o(n)}$ over $\mathbb{Z}[1/2]$. The circuit can be chosen smooth and decomposable.

The proof-complexity consequence uses a certificate from Grochow and Pitassi. Its coefficient polynomials, one per clause, are weighted model counts of prefixes of the formula. Small circuits for those counts give Hilbert-like IPS refutations of size $2^{o(n)}$ for every unsatisfiable 3-CNF, over any field of characteristic other than 2. Characteristic 0 is only sketched, from a weaker algorithm in the OpenAI preprint. Extended Frege inherits the bound only under a further simulation hypothesis.

Lean 4 checks the circuit construction, the IPS size bound from the three assumptions, the DNNF lemma and the monotone lower bound for the clique polynomial. The corollaries on Extended Frege and on Tseitin formulas are proved on paper.

WHERE THE BOUND CANNOT HOLD

The note lists the systems that cannot do this. Resolution, Res(k) for fixed k, polynomial calculus and sum-of-squares all have $2^{\Omega(n)}$ lower bounds on 3-CNFs with $O(n)$ clauses. Cutting planes and bounded-depth Frege have no such known bound, and the note leaves them as open problems. This week Alekseev and coauthors [ECCC:TR26-316] prove new lower bounds for semantic polynomial calculus over depth-3 and depth-4 circuits and for $AC^0[2]$-Frege. Read together, the two papers narrow the gap from both sides.

The knowledge-compilation result separates monotone from non-monotone circuits. A monotone circuit for a weighted model count yields a complete DNNF with no more gates. Tseitin formulas of degree-3 expanders need $2^{\Omega(n)}$ gates as DNNFs. The new circuits for them are smaller, so they must use negative constants. Cancellation also explains why the natural soundness invariant fails. Intermediate values need not be counts of anything, and a test found a network whose value is -1.

WHAT IS NOT SHOWN

Both algorithms are galactic. By the parameter choices in the preprints, OpenAI's algorithm is exhaustive search on every formula with fewer than 2^{131071} variables and clauses. Anthropic's proved bound does not beat exhaustive search unless $n > 10^{10000}$. The simulation and the refutations inherit this. Szeider states that he establishes no improvement over exhaustive search at physical sizes. A deterministic simulation of nondeterministic time is a large statement to rest on two unrefereed preprints. Its value at present lies in the reduction, not in the bound.

The open questions are concrete. Does Extended Frege refute every unsatisfiable 3-CNF in size $2^{o(n)}$ without the extra simulation hypothesis? Which is the weakest standard proof system that does? Can the circuits avoid all constants except 0, 1 and -1? And can Cook–Levin be made linear, which would let a 3-SAT saving alone transfer without loss?

[arXiv:2611.11832] Simulating Nondeterministic Turing Machines in Subexponential Deterministic Time (Carmen Strassle, Bernardo Subercaseaux) · [zenodo:24243988] Consequences of the ETH refutations for proof complexity and knowledge compilation (Szeider, Stefan) · [ECCC:TR26-316] Hard Instances for Strong Semantic Proof Systems: CNFs for Depth-Four PC and Deep Instances for AC0[2]-Frege (Yaroslav Alekseev, Svyatoslav Gryaznov, Dmitry Itsykson, Toniann Pitassi et al.)

ALGEBRAIC APPROACH

Maltsev recognition is in polynomial time

Delic and Syed give a polynomial-time algorithm that decides whether a finite relational structure, with relation arities bounded by a fixed constant, admits a Maltsev polymorphism. The algorithm needs no promise and no prior knowledge of the structure's polymorphisms.

A Maltsev polymorphism is a ternary operation m that preserves every relation of a structure and satisfies $m(x, x, y) = m(y, x, x) = y$. Its presence has a well-known consequence: every CSP with that structure as template is solvable in polynomial time. So the question "does this template have a Maltsev polymorphism?" is a natural way to recognise one class of tractable constraint languages. In [arXiv:2611.14379] Delic and Syed show that this question can itself be answered in polynomial time.

WHAT WAS KNOWN

Questions of this kind are called metaproblems. The input is a template, not an instance, and the task is to decide whether it has a polymorphism satisfying given identities. The paper's own survey of the area is short. Deciding whether a structure has a semilattice polymorphism, a conservative semilattice polymorphism or a commutative semigroup polymorphism is NP-complete. Majority and near-unanimity polymorphisms of structures and digraphs have been studied in several papers. On the algebraic side, there is work on deciding whether an idempotent algebra, given by its operations, has a Maltsev term. What was missing, the authors say, was a polynomial procedure that starts from a relational structure alone.

HOW THE ALGORITHM WORKS

The starting point is standard. The ternary polymorphisms of **A** are exactly the solutions of a CSP instance whose variables are the triples in A^3 and whose domain is A. Adding unary constraints that fix the value at every triple of the form (a, a, b) or (b, a, a) to b leaves exactly the Maltsev polymorphisms. The metaproblem thus becomes a single CSP instance of size polynomial in $|A|$ for bounded arity. A reduction due to Barto and Kozik then turns it into an equivalent instance with only binary constraints, over a power of A; a structure has a Maltsev polymorphism exactly when its powers do. The difficulty is that no known algorithm for Maltsev CSPs applies directly. Those algorithms use the polymorphisms they are given. Here none are known in advance.

> *A property of relations, checkable in logarithmic space, is made to stand in for algebraic information that the algorithm never sees.*

The authors replace that knowledge with rectangularity. If a Maltsev operation exists, every binary constraint is rectangular: whenever (a, c), (a, d) and (b, d) are in it, so is (b, c). This can be checked directly, and a failure ends the search. Once rectangularity holds, each constraint splits into a disjoint union of blocks, and the associated equivalence relations behave like congruences without any algebra being written down. Between simple pieces, a constraint is either a full product or the graph of an isomorphism. The algorithm builds a "triple graph" whose vertices are a variable with a pair of distinct values, and whose edges record constraints of the isomorphism kind. Connected components group the variables that must move together. The algorithm then tests each variable–pair combination by restricting domains along these components and running a symmetric $(1, 2)$-Datalog consistency check, discarding what fails, and repeating until nothing changes. The number of rounds is polynomial in $|A|$.

WHY IT MATTERS

For anyone who uses polymorphisms to classify templates, the gap between "the template is tractable if it has a Maltsev polymorphism" and "we can tell whether it does" has been real. This result closes that gap for Maltsev operations, at least for templates of bounded arity. It also contrasts with the semilattice case, where the corresponding question is NP-complete. The structure of the argument is of separate interest. A property of relations, checkable in logarithmic space, is made to stand in for algebraic information that the algorithm never sees.

This week's other CSP item points the same way. The tight treewidth algorithm for finite-valued hard-list VCSPs [arXiv:2611.16729] is governed by a single parameter that is also read off what the language can express, namely the largest set on which it can express a binary disequality.

WHAT REMAINS OPEN

The authors are explicit about the bound. Maltsev CSPs themselves can be solved in #L, by an algorithm that counts within known minimal subalgebras. Here the minimal subalgebras are unknown, and the authors place the metaproblem only in polynomial time, which they call the sharpest bound they can obtain. Whether it lies in a logspace-type class is open.

Several other points are worth noting. The analysis fixes a bound on the arity of the relations, so the case where arity is part of the input is not covered by the stated bound. The stated aim of producing an explicit Maltsev polymorphism, not just a yes or no, is described as an ideal rather than a theorem. And the correctness argument reasons about minimal subalgebras of an algebra the algorithm cannot see; that is the step most worth checking. Whether the rectangularity technique carries over to other conditions, such as near-unanimity, is not addressed.

[arXiv:2611.14379] Complexity of Deciding Existence of Maltsev Polymorphisms of Finite Relational Structures (Dejan Delic, Ali Syed) · [arXiv:2611.16729] Tight Algorithms for Conservative Finite-Valued CSPs Parameterized by Treewidth (Jakob Greilhuber, Dániel Marx, George Osipov, Marta Piecyk et al.)

STRUCTURAL COMPLEXITY

Karp–Lipton, pushed to the floor

Hirsch and Volkovich define an input-oblivious generator–tester–decoder class OM, prove that it lies inside each of the three deepest known Karp–Lipton collapse classes, and show that NP $\subseteq$ P/poly implies PH $=$ OM. They also prove $\mathrm{P}^{\mathrm{prS}_2^{\mathrm{P}}} \subseteq \mathrm{P}^{\mathrm{prSBP}}$. Fixed-polynomial circuit lower bounds for OM remain open.

Karp and Lipton proved that if NP $\subseteq$ P/poly, the polynomial hierarchy collapses to its second level. The theorem matters less for itself than for what it feeds. Kannan showed that PH contains, for every k, a language with no circuits of size n^k. Combine the two and a win–win follows. Either NP lacks polynomial-size circuits, or PH falls into the collapse class, which then inherits Kannan's lower bound. The smaller the collapse class, the stronger the lower bound that falls out of it.

A LINE OF SMALLER CLASSES

So the collapse has been pushed down. The learning algorithm of Bshouty, Cleve, Gavaldà, Kannan and Tamon, as used by Köbler and Watanabe, gave $\mathrm{ZPP}^{\mathrm{NP}}$. An argument of Sengupta gave the symmetric alternation class $\mathrm{S}_2^{\mathrm{P}}$. Chakaravarthy and Roy then gave two collapses: to $\mathrm{O}_2^{\mathrm{P}}$, the input-oblivious version of $\mathrm{S}_2^{\mathrm{P}}$, in which certificates may depend on the input length but not on the input; and to $\mathrm{P}^{\mathrm{prMA}}$. The two looked incomparable. Hirsch and Volkovich had already united them with a collapse to $\mathrm{P}^{\mathrm{prOMA}}$, and with Williams went slightly lower.

The new paper [arXiv:2611.16096] counts three deepest collapse classes, pairwise incomparable as far as anyone knows. The first is $\mathrm{P}^{\mathrm{prOMA}}$ and its variants. The second is $\mathrm{ZPP}^{\mathrm{ONP}}$, the input-oblivious form of Köbler–Watanabe. The third is $\mathrm{OMIN} = \mathrm{OL}_2^{\mathrm{P}}$, the oblivious version of the Korten–Pitassi class in which acceptance is read off the minimum of a polynomial-time decidable linear order. The collapse to OMIN is itself new here. It matters because OMIN is already known to have fixed-polynomial circuit lower bounds, and no smaller collapse class with that property is known.

> *The smaller the collapse class, the stronger the lower bound that falls out of it.*

The main theorem defines a class OM, shows it sits inside all three, and shows that NP $\subseteq$ P/poly implies PH $=$ OM. The collapse holds even when the certified advice string is required to be unique.

The authors add their own caution. A Karp–Lipton class inside all three is easy to build by a closure construction over tally oracles. Their claim is that OM comes from a natural model of computation and is, for instance, closed under polynomial-time Turing reductions. Naturalness is an argument rather than a theorem, and the paper presents it as one.

GENERATOR, TESTER, DECODER

The model has three parts. A polynomial-time generator sees only 1^n and, with an oracle, writes down an advice string. A tester, also blind to the input, must accept everything the generator can legally output and nothing that would mislead the decoder. A polynomial-time decoder then reads the actual input and the advice, and decides. In OM the generator's oracle is a strict oblivious Merlin–Arthur class, the tester is in co-ONP, and the decoder is in P. The tester's job rests on a simple observation: the minimum of an oblivious linear order is certified by the statement that nothing is smaller, and the minimum itself refutes every rival.

The three containments come from running the same generator three ways. Run it with its oracle and the class lands in the Merlin–Arthur collapse class. Replace the oracle by a fixed language in $\mathrm{L}_2^{\mathrm{P}}$ and the advice becomes canonical, which places the class in OMIN. Sample the oracle's answers at random and let the tester check the result, and the class lands in $\mathrm{ZPP}^{\mathrm{ONP}}$.

THE UPPER END

The second result bounds the whole picture from above. The paper proves $\mathrm{P}^{\mathrm{prS}_2^{\mathrm{P}}} \subseteq \mathrm{P}^{\mathrm{prSBP}}$, where a prSBP oracle amounts to approximate counting and SBP lies between MA and AM. This strengthens the Chakaravarthy–Roy containment $\mathrm{S}_2^{\mathrm{P}} \subseteq \mathrm{P}^{\mathrm{prAM}}$. The tool is a version of the Goldreich–Wigderson lemma: every Boolean matrix has a small set of rows that together beat every column, or a small set of columns that together beat every row. Chakaravarthy and Roy found such sets with an AM-type oracle. Here approximate counts decide, round by round, whether to pick a row or a column, choosing the denser side of what remains.

Two consequences follow. Every language in $\mathrm{O}_2^{\mathrm{P}}$ has polynomial-size circuits constructible in $\mathrm{FP}^{\mathrm{prSBP}}$. And if approximate counting can be done in $\mathrm{FP}^{\mathrm{NP}}$, then $\mathrm{P}^{\mathrm{prS}_2^{\mathrm{P}}} = \mathrm{P}^{\mathrm{NP}}$.

WHAT IT DOES NOT SHOW

The win–win does not yet reach OM. Input-oblivious classes lie in P/poly by definition, so they cannot contain NP unless NP $\subseteq$ P/poly already holds. For them the collapse and the lower bound are separate problems. For $\mathrm{O}_2^{\mathrm{P}}$ the lower bound came long after the collapse, through range avoidance. For OM it is open.

For a reader who follows circuit lower bounds, the result settles the location of the floor more than it moves the bounds. Every lower bound obtained this way remains fixed-polynomial. Whether ONP $\subseteq \mathrm{OL}_2^{\mathrm{P}}$ holds unconditionally is open. So is whether $\mathrm{P}^{\mathrm{prSBP}}$ contains $\mathrm{ZPP}^{\mathrm{NP}}$. The collapse is now as deep as anyone can currently make it. Whether it buys a lower bound is the next question.

[arXiv:2611.16096] The Ultimate Karp-Lipton Collapse (Edward A. Hirsch, Ilya Volkovich)

THE LONG VIEW how we got here

THE LONG VIEW

The polynomial method reaches ACo[2]-Frege

Alekseev, Gryaznov, Itsykson, Pitassi, Robere and Tzameret prove exponential lower bounds for semantic $\Sigma\Pi\Sigma$-PC over $\mathbb{F}_2$, superpolynomial bounds for semantic $\Sigma\Pi\Sigma\Pi$-PC (including the first superpolynomial Res(log) bound for the weak pigeonhole principle), and, for each fixed $d \geq 2$, depth-d formulas that are hard for semantic depth-d $\mathrm{AC}^0[2]$-Frege but easy at depth $d+1$; the result caps nearly forty years of attempts to carry the Razborov–Smolensky method from circuits to proofs.

THE PIGEONHOLE AS A TEST CASE

Propositional proof complexity took its first hard instance from counting. In 1985 Haken showed that resolution needs size $\exp(\Omega(n))$ to refute the pigeonhole principle. Razborov's survey records the common view that this result started the area [S1]. In 1987 Buss showed that the same principle has polynomial-size Frege proofs, so it can only separate weaker systems [S1].

In 1988 Ajtai obtained the first superpolynomial bound for bounded-depth Frege, using non-standard models of arithmetic [S1]. Bellantoni, Pitassi and Urquhart made the argument combinatorial in 1992, with a bound of $n^{\epsilon_d \log n}$. Pitassi, Beame and Impagliazzo, and independently Krajíček, Pudlák and Woods, then raised it to $\exp(n^{\epsilon_d})$ [S1].

CIRCUITS WITH PARITY GATES

In 1987 Razborov and Smolensky proved exponential lower bounds for constant-depth circuits with AND, OR, NOT and mod-p gates [S2]. Smolensky's algebraic version shows that $\mathrm{AC}^0[p]$ cannot compute MOD_q for $q \neq p$ [S3]. The method approximates each gate by a random low-degree polynomial over a finite field. Over $\mathbb{F}_2$, parity gates cost nothing, because addition already is parity.

The proof-system analogue, $\mathrm{AC}^0[p]$-Frege, became a target soon after. In 1996 Buss, Impagliazzo, Krajíček, Pudlák, Razborov and Sgall connected algebraic proof systems with bounded-depth Frege with modular counting. This week's paper describes the transfer of Razborov–Smolensky to proofs as "a significant challenge since the 1990s" [S4]. The circuit bound has stood for nearly four decades. The proof bound did not follow from it.

POLYNOMIAL CALCULUS AND DEGREE

In 1996 Clegg, Edmonds and Impagliazzo introduced a Gröbner-basis proof system, now called polynomial calculus (PC) [S4]. In 1998 Razborov proved that every PC proof of the functional pigeonhole principle needs degree $\Omega(n)$, over any field [S5]. Impagliazzo, Pudlák and Sgall simplified the argument in 1999 [S1]. Alekhnovich and Razborov gave a general route to degree lower bounds in 2001–03, which Mikša and Nordström later generalised [S4].

A common tool in this work is the pseudo-reduction operator. This is a linear map on polynomials of degree at most B with two properties: it sends every axiom to zero, and its kernel is closed under multiplication by a variable while the degree stays within B. If the operator does not send 1 to zero, no refutation of degree B exists [S4]. The machinery produced bounds on degree. It did not produce bounds on size.

THE WEAK PIGEONHOLE PRINCIPLE

With $2n$ pigeons and n holes, the principle is much harder to prove hard. In 1988 Buss and Turán showed that resolution needs $\exp(\Omega(n))$ for the onto-functional version. Atserias, Bonet and Esteban extended this to Res(2) [S1].

On the other side, in 2002 Maciel, Pitassi and Woods gave quasipolynomial $n^{O(\log n)}$ proofs in resolution over polylogarithmic-width DNFs [S1]. These yield quasipolynomial Res(log) refutations [S4]. The same year, Segerlind, Buss and Impagliazzo proved a switching lemma for small restrictions. It gave exponential bounds in Res(k) for $k \leq \sqrt{\log n / \log\log n}$ [S6]. In 2015 Razborov removed the square root, reaching $k = \epsilon \log n / \log\log n$ [S7]. Width $\log n$ stayed out of reach.

RESOLUTION OVER PARITIES

In 2008 Raz and Tzameret introduced resolution over linear equations [S4]. Itsykson and Sokolov studied its mod-2 version, Res(⊠), and proved lower bounds for the tree-like form [S8]. Further restricted forms followed:

- **Regular** Res(⊠): Efremenko, Garlík and Itsykson, 2023 [S9].
- **Bounded-depth** Res(⊠): Alekseev and Itsykson, by lifting, 2025 [S10].
- **Polynomial-depth** Res(⊠): Alekseev and Gaevoy, January 2026 [S10].

In September 2026 Braun proved the first lower bound for unrestricted DAG-like Res(⊠). Refutations of the bit pigeonhole principle need $2^{\Omega(n/\log^2 n)}$ clauses [S10]. The proof is formalised in Lean 4 and, in Braun's words, was "developed with substantial AI assistance" [S10].

In October two related results appeared. Pang gave a degree–size relation for Res(PC) over prime fields. Conneryd used it to prove $\exp(\Omega_p((\log n)^{3/2}))$ bounds for syntactic depth-2 $\mathrm{AC}^0[p]$-Frege on random $(\log n)$-CNFs [S4, S11].

TWO PIECES ALREADY IN PLACE

Two earlier results set up the target.

The first is Grochow and Pitassi's Ideal Proof System (IPS) of 2018. It converts depth-d $\mathrm{AC}^0[2]$-Frege refutations into $\mathbb{F}_2$-IPS refutations of polynomially related size and depth $O(d)$ [S4]. In 2024 Hakoniemi, Limaye and Tzameret obtained constant-depth IPS lower bounds. These hold only for individual degree $O(\log\log n)$, which is where the Frege simulation breaks [S4].

The second is a construction for the coin problem. Limaye, Sreenivasaiah, Srinivasan, Tripathi and Venkitesh built monotone AC^0 formulas for the δ-coin problem, of depth d and size $\exp(O(d(1/\delta)^{1/(d-1)}))$. They used these to prove a fixed-depth size hierarchy for $\mathrm{AC}^0[\oplus]$ formulas [S12].

THIS WEEK

Alekseev, Gryaznov, Itsykson, Pitassi, Robere and Tzameret [ECCC:TR26-316] read Braun's argument as an implicit annihilating-polynomial argument. They recast it in terms of pseudo-reduction operators [S4].

The argument runs as follows. Multiply every line of a semantic refutation by a low-degree polynomial f that vanishes on the support of each high-degree term. The refutation becomes a low-degree derivation of f. If the operator R satisfies $R(f) \neq 0$, that derivation cannot exist. Ruling out low degree needs only $R \neq 0$. Ruling out small size needs R to have large rank, and the bound is $S(\pi) \geq \mathrm{rank}(R_{\leq k})/H_{n-r}(k)$ [S4].

Three further tools extend this [S4]:

- Razborov–Smolensky approximation carries the argument to depth four.
- Lifting and random restrictions remove the bound on bottom fan-in.
- Coin-problem formulas supply an approximate-majority gadget for $\mathrm{AC}^0[2]$-Frege, where the bound is $\exp(\Omega_d((\log N)^{d/(d-1)}))$.

The results are these:

- $2^{\Omega(n)}$ for semantic $\Sigma\Pi\Sigma$-PC on the bit and functional pigeonhole principles and on random 3-CNFs.
- $\exp(\Omega((\log n)^2/\log\log n))$ for $\Sigma\Pi\Sigma\Pi$-PC on PHP^{2n}_n, which settles the Res(log) question for the weak pigeonhole principle.
- For each fixed $d \geq 2$, depth-d formulas that are hard for semantic depth-d $\mathrm{AC}^0[2]$-Frege but easy at depth $d+1$.

The authors state that they used OpenAI models as interactive tools, and that they wrote and verified the paper themselves [S4].

WHAT IT DOES NOT SHOW

The hard instances are deep formulas, not CNFs. There is a reason: depth-5 semantic $\Sigma\Pi\Sigma\Pi\Sigma$-PC refutes every unsatisfiable CNF in polynomial size [S4]. Syntactic $\mathrm{AC}^0[p]$-Frege lower bounds for CNFs remain open [S13]. The route through Pang's condition needs near-maximal PC degree bounds, which the authors call "a form of SETH for PC" [S4]. The work covers $\mathbb{F}_2$ only. In Res(log), the upper and lower bounds for the weak pigeonhole principle are now both quasipolynomial, but their exponents do not meet [S4].

This week's result: [ECCC:TR26-316] Hard Instances for Strong Semantic Proof Systems: CNFs for Depth-Four PC and Deep Instances for AC0[2]-Frege (Yaroslav Alekseev, Svyatoslav Gryaznov, Dmitry Itsykson, Toniann Pitassi et al.)

THE TECHNIQUE: RAZBOROV–SMOLENSKY POLYNOMIAL APPROXIMATION

The method replaces each AND or OR gate of a constant-depth circuit by a *random* low-degree polynomial over $\mathbb{F}_2$. On any fixed input, this polynomial agrees with the gate with high probability. Parity gates need no approximation, since addition over $\mathbb{F}_2$ is parity. Applied to a whole circuit, the method yields one low-degree polynomial that is correct on most inputs. A lower bound then follows by showing the target function has no such approximation.

OR of n bits. Choose random subsets $S_1, \ldots, S_k \subseteq [n]$ and set $p(x) = 1 + \prod_{j=1}^{k}(1 + \sum_{i \in S_j} x_i)$.

This polynomial has three properties:

- Its degree is at most k.
- At $x = 0$, every subset sum is 0, so $p(0) = 0$, which is correct.
- For a fixed nonzero x, each subset sum equals 1 with probability exactly $1/2$. So p is wrong only if all k sums vanish, which happens with probability 2^{-k}.

A toy instance. Take $n = 3$, $k = 2$, $S_1 = \{1,2\}$ and $S_2 = \{2,3\}$. Then $p = 1 + (1 + x_1 + x_2)(1 + x_2 + x_3)$. Multilinearising ($x_i^2 = x_i$) and reducing mod 2 gives $p = x_1 + x_2 + x_3 + x_1x_2 + x_1x_3 + x_2x_3$.

This agrees with OR on seven of the eight inputs. It fails only at 111, where both subset sums are even.

Whole circuits. Approximate the gates from the bottom up. Each AND or OR layer multiplies the degree by at most k, so the degree grows to about k^d at depth d. A union bound over s gates gives error at most $s \cdot 2^{-k}$ on each fixed input.

How the new paper uses it. A random error bound is not enough for a proof, because the approximation must hold everywhere. The new paper therefore asks for one polynomial f that annihilates the difference between every term and its approximation. This makes the approximations exact along the entire refutation.

TIMELINE

1985	Haken: resolution refutations of the pigeonhole principle need size $\exp(\Omega(n))$.
1987	Razborov and Smolensky: exponential lower bounds for $\mathrm{AC}^0[p]$ circuits by polynomial approximation; $\mathrm{MOD}_q \notin \mathrm{AC}^0[p]$ for $q \neq p$.
1988	Ajtai: first superpolynomial bounded-depth Frege lower bound for the pigeonhole principle.
1996	Clegg, Edmonds and Impagliazzo: the Gröbner-basis proof system now called polynomial calculus.
1998	Razborov: PC proofs of the functional pigeonhole principle need degree $\Omega(n)$ over any field.

2002 Segerlind, Buss and Impagliazzo: weak pigeonhole principle exponentially hard for Res(k) with $k \leq \sqrt{\log n / \log\log n}$; Razborov reaches $\epsilon \log n / \log\log n$ in 2015.

2008–2014 Raz and Tzameret introduce resolution over linear equations; Itsykson and Sokolov prove tree-like Res(⊕) lower bounds.

2019 Limaye, Sreenivasaiah, Srinivasan, Tripathi and Venkitesh: explicit constant-depth formulas for the coin problem and a fixed-depth hierarchy for $\mathrm{AC}^0[\oplus]$.

2026 (Sep) Braun: first DAG-like Res(⊕) lower bound, $2^{\Omega(n/\log^2 n)}$ for the bit pigeonhole principle, Lean-formalised and developed with substantial AI assistance.

2026 (Nov) Alekseev, Gryaznov, Itsykson, Pitassi, Robere and Tzameret: $\Sigma\Pi\Sigma$- and $\Sigma\Pi\Sigma\Pi$-PC lower bounds, Res(log) for the weak pigeonhole principle, and hard deep instances for every fixed-depth $\mathrm{AC}^0[2]$-Frege.

SOURCES

[S1] Proof Complexity of Pigeonhole Principles (Razborov survey)

[S2] Functional lower bounds for arithmetic circuits and connections to boolean circuit complexity

[S3] On teaching the approximation method for circuit lower bounds (Goldreich)

[S4] ECCC – TR26-316

[S5] Lower bounds for the polynomial calculus

[S6] A Switching Lemma for Small Restrictions and Lower Bounds for k-DNF Resolution

[S7] Pseudorandom generators hard for k-DNF resolution and polynomial calculus resolution

[S8] Resolution Over Linear Equations: Combinatorial Games for Tree-like Size and Space

[S9] Lower bounds for regular resolution over parities

[S10] An exponential lower bound for the bit pigeonhole principle in resolution over parities

[S11] Random (log n)-CNF Formulas Are Hard for Depth-2 Frege with Modular Gates

[S12] A Fixed-Depth Size-Hierarchy Theorem for AC0[⊠] via the Coin Problem

[S13] Hard CNF Instances for Ideal Proof Systems

THE COLUMN

Shell Scott



SHELL SCOTT, NIGHT DESK

Taking Inventory

The late column. The week's papers set against each other, looking for the pattern nobody wrote down. Plain words. No favourites.

The week had no headline of its own. It had one from late October, and a stack of papers doing the accounting.

That headline is the pair of subexponential 3-SAT algorithms reported by OpenAI and Anthropic, and with them the refutation of ETH. Separately, Alman and Vassilevska Williams have refuted 3SUM and APSP. Nobody this week re-proves any of it. The papers ask what it touches.

Start with the direct consequence. Strassle and Subercaseaux [arXiv:2611.11832] show the algorithms imply $\mathsf{NTIME}[t] \subseteq \mathsf{DTIME}[2^{O(t \log\log t/\sqrt{\log t})}]$, and $2^{O(t/\sqrt{\log t})}$ from a stronger consequence of the proof. The route is a generalisation of Cook–Levin. Any subexponential circuit-CSP algorithm transfers to NTIME with the same savings. That is the honest size of the thing: a $\sqrt{\log t}$ in the exponent.

Szeider [zenodo:24243988] takes three statements of the Anthropic preprint as given. From them he derives $2^{o(n)}$ smooth decomposable arithmetic circuits for the weighted model count of every 3-CNF. Hence $2^{o(n)}$ IPS refutations of every unsatisfiable 3-CNF, over every field of characteristic other than 2. Lean checks the derivation from the three statements. It does not check the three statements. The note says plainly that both algorithms are galactic and that it establishes no improvement over exhaustive search at physical sizes.

Then the repairs. Bilyk, Munduzbaev, Sharma and Siam [arXiv:2611.12321] settle the Gaspers–Szeider question: strong backdoors to forests in $2^{2^{O(k)}}|\varphi|$ time on general non-tautological CNFs, with W[1]-hardness for treewidth $t \geq 2$. Their tightness argument for the double exponential no longer rests on ETH. It rests on 3-SAT having no $2^{o(N/\log N)}$ algorithm, a bar placed below the claimed ones on purpose. Karthik and Minzer [arXiv:2611.14892] repair a k-means reduction whose hypothesis, XXH, does not hold in the regime where it was used, and rebuild it on Almost ETH. Cheng, Huang and Wu [arXiv:2611.17690] prove read-k BPL = BPL for constant k, and note what follows if the announced BPL = L is true.

Tomar [hexagon:2611.00388v1] does the job for 3SUM and APSP. When a load-bearing wall comes out, someone has to walk the building and check which ceilings are still up. Most set-disjointness, dynamic-graph and triangle bounds survive, because their reductions produce balanced instances, or lopsided ones with a middle part too large for the new algorithms. A Goldstein–Lewenstein–Porat bound for Set Disjointness over small universes is false. Some Kopelowitz–Pettie–Porat trade-offs have lost their support. And the surviving hypotheses drew

their evidence largely from the refuted ones. For several bounds the hypothesis is the bound itself, restated for the instances that remain.

The ceilings are up. Nobody has checked what holds them.

So the pattern. The conditional lower bound was for a long time the cheapest instrument the field owned. This week it acquired a price. Hypotheses are being weakened by hand to sit just under whatever was claimed last. The premises downstream are no longer conjectures about the world. They are statements lifted from a preprint which, by Tucker's account [zenodo:24227146], Anthropic itself says no person is known to have checked.

The other half of the week runs the other way. Alekseev, Gryaznov, Itsykson, Pitassi, Robere and Tzameret [ECCC:TR26-316] prove unconditional lower bounds: $2^{\Omega(n)}$ for random 3-CNFs against semantic $\Sigma\Pi\Sigma$-PC over $\mathbb{F}_2$; superpolynomial bounds for semantic $\Sigma\Pi\Sigma\Pi$-PC that settle Res(log) for the weak pigeonhole principle; and the first unconditional hard instances for $AC^0[2]$-Frege at every fixed depth, using deep formulas. The method is Braun's AI-generated argument for resolution over parities, distilled into a framework of pseudo-reduction operators. Dvir [arXiv:2611.13281] writes up constant-query LCCs over the reals, found by ChatGPT Astra, previously conjectured not to exist.

Set the two proof-complexity papers side by side. Lower bounds for random 3-CNFs climb from below, in characteristic 2. A conditional subexponential upper bound sits at IPS, in characteristic other than 2. Szeider lists where such an upper bound is ruled out: resolution, Res(k) for fixed k, polynomial calculus, sum-of-squares. The interval between is where proof complexity now works, and both of this week's markers on it trace back, in part, to machine output.

Only one of the two needs re-inspection if a preprint changes. *— Szeider's note runs to fourteen pages; the companion that states its assumptions runs to eighteen.*

IN BRIEF

everything else that appeared this week

arXiv cs.CC 49 papers

The arXiv listing for cs.CC, moderated for scope, not refereed.

Polynomial Swap Coboundary Expansion with Applications to LTCs and PCPs Yotam Dikstein, Izhar Oppenheim, Inga Valentiner-Branth [arXiv:2611.16295]

Proves that suitable skeleta of Kac–Moody–Steinberg complexes have r-swap coboundary expansion at least $r^{-O(1)}$, the first inverse-polynomial bound for bounded-degree complexes; substituted into an existing analysis, this gives two-query PCPs with perfect completeness, soundness ε and proof length $n(\log n)^{\exp(\mathrm{poly}(1/\varepsilon))}$.

Towards Understanding the Complexity of CAQE: A Proof-Theoretic Analysis of its Core Procedure Benjamin Böhm, Olaf Beyersdorff [arXiv:2611.19613]

Formalises runs of the basic CAQE QBF algorithm as proofs in a system CAQE$^{\mathrm{Core}}$ and shows that, on QCNFs of bounded quantifier alternation, Q-resolution with an NP oracle simulates it; QCDCL simulates it on Σ_3^b QCNFs after preprocessing, and in that setting both systems are strictly stronger.

Almost Optimal Lower Bounds under Almost ETH for Approximating Euclidean k-means Karthik C. S., Dor Minzer [arXiv:2611.14892]

Proves near-optimal lower bounds under an Almost-ETH hypothesis for $(1+\varepsilon)$-approximation of high-dimensional Euclidean k-means, ruling out running times of the form $2^{(k/\varepsilon)^{1-\beta}} \cdot \mathrm{poly}(n,d)$ and so nearly matching the known $2^{\tilde{O}(k/\varepsilon)}$ algorithms.

A short proof of the Aaronson-Ambainis Conjecture via the completion framework Thomas Huffstutler, David Miloschewsky, Supartha Podder [arXiv:2611.12795]

Gives a short proof of the Aaronson–Ambainis conjecture: a partial Boolean function defined on a δ-fraction of inputs can be computed by a deterministic algorithm, correctly on a $1-\varepsilon$ fraction of its domain, with $O(Q(f)^6/(\varepsilon\delta)^2)$ queries.

Constant query LCCs over the reals Zeev Dvir [arXiv:2611.13281]

Constructs constant-query locally correctable codes over the real numbers, which had been conjectured not to exist; the codes use low-degree polynomial evaluations on an integer grid, and decoding samples random quadratic curves through the grid.

On Non-Deterministic Representation of Boolean Functions Sudharshan Rajagopalan, Nitin Saurabh [arXiv:2611.11966]

Proves that almost all Boolean functions on n variables have non-deterministic degree at most $n/2 + O(\sqrt{n \log n})$, confirming a conjecture of Iyer et al.

Weak Composition Theorems for Randomized Communication Pachara Sawettamalya [arXiv:2611.16192]

Proves that, up to additive terms, computing $g \circ f^n$ in randomised communication requires, information-theoretically, solving at least $\mathrm{Inf}(g)$ instances of f; this is the first composition lower bound that applies to every choice of inner function f and outer function g.

Exact Quantum Algorithms for Composed Boolean Functions Gagan Chiruvella, Chandrima Kayal, Rajat Mittal, Manaswi Paraashar et al. [arXiv:2611.16412]

Proves that $Q_E(f \circ g) = O(mD + Q_E(f)Q_{1/(4D)}(g))$ for an outer function f of degree D and an inner function g on m bits, by combining approximate inner evaluations with a finite exact correction.

The multiplicative complexity of the three-bit low-half multiplier Bshara Alhodali, Gregory Morse [arXiv:2611.15992]

Proves, without enumeration and with a full Lean 4 formalisation, that the multiplicative complexity of the three-bit low-half multiplier is exactly six, a value previously established by SAT search.

Simulating Space Bounded Read-k PTMs with Space Bounded Read-Once PTMs Kuan Cheng, Shengtang Huang, Ruiyang Wu [arXiv:2611.17690]

Shows that read-k probabilistic Turing machines with space S and time T can be simulated by read-once ones in space $O(k^2(S + \log(n+T)) + \log(1/\varepsilon))$, so read-$k$ BPL equals BPL for every constant k.

Independent set reconfiguration is hard at treewidth four Hussein Abbas, Haydar Hamieh, Amer E. Mouawad, Ali Slim [arXiv:2611.20205]

Proves that Independent Set Reconfiguration is hard on graphs of treewidth four under both token sliding and token jumping, narrowing the long-open gap between polynomial-time solvability on trees and known hardness at some unspecified constant treewidth.

EQP $\subseteq$ IP Does Not Relativize Sinan Oral, Supartha Podder [arXiv:2611.12746]

Constructs a recursive oracle relative to which a unary language lies in EQP but outside both IP and PH, extending the BQP versus IP oracle separation of Bouland et al. to exact quantum computation through a Forrelation variant with correlation exactly $\pm 1/2$.

Go under Chinese Rules with a Succinct Game History is Complete for Exponential Space Markus Bläser, Felix Czerner [arXiv:2611.18201]

Proves that Go under Chinese rules, with its no-repetition rule, is complete for exponential space when the game history is given succinctly.

Polynomial-Space Quantum Algorithms from Classical Branching: Colouring, Domination and Monotone Local Search Sebastian Senge [arXiv:2611.12714]

Gives bounded-error, polynomial-space quantum algorithms without quantum random access memory, built from classical branching, including 4-Colouring in time $O(1.3032^n)$ (previously $O(1.3231^n)$) and improved bounds for k-Colouring for several values of k between 5 and 18.

Uniform Quantum Average-Case Fine-Grained Hardness of Arithmetized k-Orthogonal Vectors Dhruv Prasad Deshpande, Abhishek Mishra [arXiv:2611.14411]

Shows that, for every fixed $k \geq 3$ and assuming the basic decision version of QSETH, no uniform quantum algorithm running in time $O(n^{k/2-\varepsilon})$ evaluates the arithmetized k-Orthogonal Vectors polynomial correctly with probability $3/4$ over uniformly random inputs.

SVP_∞ Is NP-Hard for Power-of-two Cyclotomic Ideals Yansong Feng, Jiaqi Liu, Hengyi Luo, Yanbin Pan [arXiv:2611.14334]

Proves that exact decision-SVP in the coefficient ℓ_∞-norm is NP-complete for principal ideals and rank-two free modules over $\mathbb{Z}[X]/(X^n+1)$ with n a power of two, by polynomial-time reductions from Subset Sum.

Two-Stage Mixed-Integer Quadratic Programming in Polynomial Time Alberto Del Pia [arXiv:2611.14837]

Gives a polynomial-time algorithm for mixed-integer quadratic programs with two-stage structure, a Lenstra-type result that has to rely on such structure because MIQP is NP-hard even without integer variables.

Individually Fair Facility Location Rajni Dabas, Samir Khuller, Milind Maiti [arXiv:2611.16138]

Proves that for individually fair facility location, with any fixed radius violation below three, no deterministic polynomial-time algorithm achieves a sublogarithmic cost factor unless P = NP, even with uniform opening costs and radii, which contradicts a recently claimed constant-factor guarantee.

The Tournament Equilibrium Set Is PSPACE-Complete, Even under Hereditary Uniqueness Ethan Dickey, Athina Terzoglou [arXiv:2611.17312]

Proves that deciding membership in Schwartz's tournament equilibrium set is PSPACE-complete, closing the gap between the known NP-hardness and the PSPACE upper bound, and that hardness persists under hereditary uniqueness.

Local Modifications of Nearest-Neighbor and Min-Plus Representations Liliia Redina [arXiv:2611.13787]

Bounds how much changing a Boolean-valued function at d inputs can enlarge its nearest-neighbour and integer min-plus representations, showing for example that $m + 2d$ anchors suffice for nearest-neighbour representations on finite subsets of $\mathbb{R}^n$, $n \geq 2$, and that the coefficient 2 is optimal.

Type-logical grammars and formal languages Richard Moot [arXiv:2611.12618]

Shows that type-logical grammars with the standard rules for product, associativity, extraction and infixation have an NP-complete fixed recognition problem, and that first-order linear logic grammars can simulate the quasi-linear grammars of Book and Greibach.

Pair-Merge Tokenisation Is Hard to Approximate on Binary Strings and Strongly NP-Hard on Unary Datasets Kundan Baliga, Aiden Xie, Aditya Byju [arXiv:2611.13568]

Proves that, unless P = NP, no polynomial-time algorithm can choose k BPE-style pair merges that shorten a single binary string to within a factor $cn/\log n$ of optimal, and that the problem is strongly NP-hard on unary datasets.

Feedback Vertex Set, Max Cut, and Edge Coloring on the Strong Product of a Graph and a Path Jan Diester, Kolja Kühn, Manuel Rapp, Samuel Schneider et al. [arXiv:2611.13917]

Proves Max Cut NP-hard on induced subgraphs of the strong product of two paths, answering a question of Dvořák, Huynh, Joret, Liu and Wood, and shows 3-Edge Colouring NP-hard on its subgraphs and equivalent to the planar case on its induced subgraphs.

Sowing the Seeds of Intractability: Kalah is PSPACE-complete Ethan Dickey, Ethan Mader [arXiv:2611.17252]

Proves that deciding whether a player can force a win in generalised Kalah, with the number of pits per side part of the input and stone counts in unary, is PSPACE-complete.

On Applications of Polynomial-Time Group Isomorphism Algorithm Wenkang Zhang, Zhinan Wang, Jincheng Zhuang [arXiv:2611.18084]

Builds on the Qiao–Zhang polynomial-time algorithm for finite group isomorphism in the multiplication-table model to obtain polynomial-time algorithms that return the coset of all isomorphisms for finite associative rings, with further applications to bilinear-map isotopism and free three-tensor isomorphism over finite commutative rings.

Impatience Makes Matching NP-Hard Chen Yan, Sushil Mahavir Varma [arXiv:2611.20234]

Shows a phase transition in the complexity of maximising long-run throughput in a dynamic bipartite matching model where customers and suppliers arrive by Poisson processes and may leave before they are matched.

Polynomial Sample Complexity of Noninteractive Simulation over Finite Alphabets Haoran Wang [arXiv:2611.11974]

Shows that for any fixed finite correlated source, every joint law that Alice and Bob can simulate without communication from arbitrarily many samples can be approximated using a number of samples polynomial in the output alphabet sizes and the inverse accuracy.

Affine Combiners for Reed-Solomon Proximity Gaps Alessandro Chiesa, Guy Weissenberg, Eylon Yogev, Kai Zhe Zheng [arXiv:2611.15128]

Proves that affine combiners of the form $f(x) + (\alpha + \beta x)g(x)$ give smaller error bounds for Reed–Solomon proximity gaps than the $O(n/q)$ of scalar combiners, in a range of distances extending towards the Johnson radius for sufficiently long codes.

Statistical-to-computational gaps in central limit theorems for high-dimensional tensors Debsurya De, Dmitriy Kunisky [arXiv:2611.15803]

Treats high-dimensional central limit theorems as testing problems, sharpening Mikulincer's total-variation bounds for distinguishing Wishart from Wigner tensors with matching polynomial-time tests, and studies tensor settings in which the statistical and computational thresholds come apart.

Tight Quantum Query Bounds for Collision Finding over Arbitrary Bounded Types Tianci Peng, Shiqi Hou, Xiaoyang Dong [arXiv:2611.17804]

Determines the bounded-error quantum query complexity of finding an r-collision in a function with prescribed equality-class sizes, as the maximum of scales set by the numbers of classes of size at least $2, \dots, r$, with the lower bound holding for a uniformly random function of that type.

Parameterized Complexity and Geometric Approximation of Riesz Energy Subset Selection Michael Emmerich, Soumen Atta [arXiv:2611.13600]

Gives an efficient parameterized approximation scheme for choosing k of n sites to minimise Riesz s-energy in fixed Euclidean dimension and fixed $s > 0$, sharpened by a farthest-first packing argument; exact selection for $s = 2$ is polynomial on a line and NP-hard in the plane.

Computation with Backward Feedback: Power, Observation, and Cost Saman Farhat, Farzad Didehvar [arXiv:2611.16779]

Introduces a model of deterministic programs whose runs are read at stages chosen by a random clock that can move backward and feed register contents to earlier stages, and studies how much power this feedback adds and how soon an observer can obtain the answer.

Succinct Infinite-Horizon Lyapunov and LQR Value Problems Are BQP-Complete Mohamed Serry, Jun Liu [arXiv:2611.12757]

Proves that promise problems for evaluating infinite-horizon Lyapunov output energies and LQR optimal values are BQP-complete when the sparse system matrices are given succinctly by Boolean circuits.

Computational Power of Clifford and Ising Quantum Computers Sawyer Dobson, Zhenghan Wang [arXiv:2611.13347]

Extends the efficient classical simulation of Ising quantum computers from braidings to measurements, and shows that Clifford quantum computers whose circuit generator sees only the input length decide exactly the so-called affine languages, while efficient classical simulation with a generator that sees the whole input would separate P from PSPACE.

Tokenization: Hardness and Bicriteria Approximation Algorithms Atieh Bastan, Xiaoling He, A. Pavan, N. V. Vinodchandran [arXiv:2611.15074]

Proves that choosing at most K substrings to minimise the token count of a binary corpus has no polynomial-time $n^{1-\varepsilon}$-approximation for any fixed $\varepsilon \in (0, 1)$ unless P = NP, and gives bicriteria approximation algorithms.

Small Maximum Likelihood on Fixed Cavender-Farris-Neyman Trees: An Approximation Scheme and Limits to Fully Polynomial Approximation Mohammad-Hadi Foroughmand-Araabi, Sama Goliaei [arXiv:2611.15784]

Gives a deterministic polynomial-time approximation scheme for fitting edge substitution probabilities to an alignment on a fixed binary Cavender–Farris–Neyman tree, and identifies limits to a fully polynomial scheme.

Topological Entropy of Group Cellular Automata is Effectively Computable Niccolò Castronuovo, Alberto Dennunzio, Jarkko Kari, Luciano Margara [arXiv:2611.16746]

Shows that the topological entropy of cellular automata whose global map is a group endomorphism is effectively computable, in dimension one by splitting it into left and right components computed from finite families of traces.

Tight Quantum Query Bounds for Minimum Cut and Multicut on Dense Graphs Tianhang Lu, Runtian Ren, Shengcai Liu [arXiv:2611.17845]

Shows that exact minimum multicut with at most k demand pairs on dense simple graphs has quantum query complexity $\tilde{O}(n + \min\{m, n^{3/2}\sqrt{k}\})$ with a matching lower bound, giving $\tilde{O}(n^{3/2})$ queries for a prescribed s–t cut.

Who Observes Whom: Network Rationalizability and Its Complexity Mojtaba Fayaz-Bakhsh, MohammadAmin Fazli [arXiv:2611.20110]

Defines solution concepts induced by a directed observation network, lying between Nash equilibrium and rationalizability, and studies the complexity of computing them.

A Fast Algorithm for Subspace Convolution Shlomo A. Perles, Omri Weinstein [arXiv:2611.13134]

Shows that summation over all linear subspaces of $\mathbb{F}_q^n$, the analogue of the subset-zeta transform, can be computed in $\Theta_q(N_q \log N_q)$ time, where N_q is the number of subspaces.

The Monophonic Directed Convexities Júlio Araújo, Ana Karolinna Maia, Thiago Marcilon, Nicolas Nisse et al. [arXiv:2611.14191]

Defines the monophonic and strong monophonic convexities on directed graphs, based on minimal and induced directed paths, and studies the complexity of their basic decision problems.

Planar PArcK is PSPACE-complete Craig Tennenhouse, Kanae Yoshiwatari [arXiv:2611.20138]

Tennenhouse and Yoshiwatari prove that Partizan Arc Kayles, a generalisation of Domineering played on edge-coloured graphs, is PSPACE-complete even on planar graphs of maximum degree three, by a reduction from Generalized Hex.

From Worst-Case to Smoothed Complexity in Continuous Optimization Daniel Cortild, Coralia Cartis [arXiv:2611.14282]

Cortild and Cartis develop a smoothed-analysis framework for classical continuous optimisation algorithms, showing that methods suited to non-degenerate functions, applied to random perturbations of degenerate ones, recover the faster rates known for the non-degenerate class.

Solomonoff Induction with Reflective Oracles Damini Kusum [arXiv:2611.14454]

Kusum shows that a variant of Solomonoff induction built on reflective oracles succeeds against computable predictive strategies that specify only part of the data-generating process, including incomplete models given by computable interval restrictions on next-bit probabilities.

ECCC 4 papers

The Electronic Colloquium on Computational Complexity, the field's own report series. Screened for scope, not refereed.

On derandomizing the line-point low-degree test

Prahladh Harsha, Swastik Kopparty, Mrinal Kumar, Ramprasad Saptharishi et al. [ECCC:TR26-315]

Studies the derandomised bivariate line-point low-degree test in the high-error regime and finds unexpected behaviour: a counterexample when the field size is around k^2, alongside a positive result for larger fields.

Polynomial-size hitting sets for multilinear any-order read-once algebraic branching programs and for set-multilinear depth-3 circuits with constant-size variable sets

Amir Shpilka [ECCC:TR26-314]

Gives polynomial-size hitting sets, over fields of any characteristic, for multilinear any-order read-once oblivious algebraic branching programs and for set-multilinear depth-3 circuits with constant-size variable sets, including the border of both classes.

Explicit Subspace Evasive Sets via Subspace Designs

Joshua Brakensiek, Yeyuan Chen, Mahdi Cheraghchi, Rohan Goyal et al. [ECCC:TR26-317]

Gives explicit subspace-evasive sets built from subspace designs whose intersection bound is polynomial in the dimension r, where earlier explicit constructions had bounds exponential in r, and uses them to obtain explicit list-recoverable codes.

FROM THE AI-ASSISTED REPOSITORIES

unrefereed; read with care

Hexagon 9 papers

An AI-assisted repository: AI-written and AI-co-written work with mandatory provenance labels. Nothing here is refereed. Read with care.

9	Anthropic Claude Opus	algebraic complexity
PREPRINTS THIS WEEK	MOST CREDITED SYSTEM (2 OF 9)	MOST COMMON TOPIC (2 PAPERS)

Nine papers came in, none refereed, all with provenance labels. Three touch the reader's interests. [hexagon:2611.00388v1] claims to sort the fine-grained lower bounds built on 3SUM and APSP into those that survive the reported refutations and those that do not. Most survive only on hypotheses about other instance shapes, and a few are stated to be false. [hexagon:2611.00439v1] claims a coefficient of $6/\sqrt{2\pi}$ for the average sensitivity of degree-d threshold functions; no human author claims it. [hexagon:2611.00372v1] and [hexagon:2611.00375v1] work on the rank of 4×4 matrix multiplication over $\mathbb{F}_2$ from below and above, and [hexagon:2611.00374v1] tightens selection-network counts. The rest lies outside complexity. Most of the batch moves small constants; the survey is the exception.

THEMES rank of 4×4 matrix multiplication over F2 · exact constants and small-case lower bounds

CLAIM OF THE WEEK The paper claims that most 3SUM- and APSP-based lower bounds survive only on hypotheses about instances outside the new algorithms' range, that a few, including a Goldstein–Lewenstein–Porat Set Disjointness bound over small universes, are false, and that the lopsided-triangle method extends to $\varepsilon < 0.17843$ when at most $n^{2-\varepsilon/2}$ prescribed edges are known in advance. [hexagon:2611.00388v1]

Which Fine-Grained Lower Bounds Survive the Refutation of 3SUM and APSP?

Tejasvi Singh Tomar [hexagon:2611.00388v1]

Claims to sort the main conditional lower bounds derived from the refuted 3SUM and APSP hypotheses into those that survive, those that have lost their support and those that are false, among the last a Set Disjointness bound of Goldstein, Lewenstein and Porat over small universes; it also extends the lopsided-triangle method to every $\varepsilon < 0.17843$ when the prescribed edges are known in advance.

An edge comparison with constant 3/2 and the coefficient 6/sqrt(2 pi) for the average sensitivity of polynomial threshold functions

David Niedbala Giraudin · with Anthropic Claude Opus [hexagon:2611.00439v1]

Claims that every degree-d polynomial threshold function on $\{-1,1\}^n$ has average sensitivity at most $(6/\sqrt{2\pi})d\sqrt{n} < 2.394d\sqrt{n}$, improving the coefficient from $8/\sqrt{2\pi}$ by means of an edge-comparison lemma with constant $3/2$ in place of 2.

The rank of 4×4 matrix multiplication over F2 is at least 35

Tejasvi Singh Tomar [hexagon:2611.00372v1]

Claims that the rank of the 4×4 matrix multiplication tensor over $\mathbb{F}_2$ is at least 35, up from 34, by proving that the rank of $\langle 3,3,4\rangle$ over $\mathbb{F}_2$ is at least 27 and then substituting.

Exact sizes of networks that select the three or four largest of n elements

Tejasvi Singh Tomar [hexagon:2611.00374v1]

Claims to raise Alekseev's lower bound $(n-t)\lceil\lg(t+1)\rceil$ on comparators in (t,n)-selection networks by $k-1$ when $t=2^k\geq4$ and by $2k-5$ when $t\geq32$, and that Yao's upper bound for $t=3$ is exact for $4\leq n\leq18$.

Local optimality of rank-47 schemes for 4×4 matrix multiplication over F2

Tejasvi Singh Tomar [hexagon:2611.00375v1]

Claims that five pairwise inequivalent rank-47 decompositions of 4×4 matrix multiplication over $\mathbb{F}_2$, and all decompositions equivalent to them, are locally optimal to depth 18, which limits how many terms a rank-46 decomposition could share with them.

The Logical Anatomy of Optimization: Reverse Mathematics, Proof Mining, and the Price of Convergence

Xiaoyu Li · with Anthropic Claude Opus, Anthropic Claude Code, Python Software Foundation Python 3 (numerical checks in anc/) [hexagon:2611.00217v1]

Claims a map of how much mathematics the existence, convergence and duality theorems of optimisation need, in terms of reverse mathematics, Weihrauch complexity and proof mining, including a proof in RCA_0 that convex functions on $[0,1]$ attain their minima.

Energy complexity and dissipation of noisy consensus Zhenyu Sun [hexagon:2611.00292v1]

Claims that any protocol for n processes to agree on L binary coordinates over independent Gaussian channels, with failure probability at most ε, needs worst-input expected signal energy $\Omega(N_0 n[L + \ln(n/\varepsilon)])$, and relates this bound to dissipation.

Non-computable derivatives of the fractional-part autocorrelation Michael Yampolsky [hexagon:2611.00243v1]

Yampolsky claims that the multiplicative autocorrelation of the fractional-part function, $A(x) = \int_0^\infty \{t\}\{xt\}dt/t^2$, is computable and differentiable almost everywhere, yet its finite derivative can be non-computable even at a computable point x.

Quantum-Information Probes of Collapse Dynamics in Causal Fermion Systems Benjamin Kaminsky [hexagon:2611.00324v1]

Kaminsky claims to translate the effective collapse model that Finster, Kleiner and Paganini derived from causal fermion systems into quantities measurable on coherent quantum processors, including an exact purity-loss identity for the associated Dirichlet form under arbitrary Hamiltonian control.

Diderot 3 papers

An AI-assisted preprint server on which AI systems may be listed as co-authors. Nothing here is refereed. Read with care.

3	openai/gpt-5.6-sol	combinatorics
PREPRINTS THIS WEEK	MOST CREDITED SYSTEM (2 OF 3)	MOST COMMON TOPIC (2 PAPERS)

Three papers came in this week, and none is complexity theory. Two by Cresswell concern Pilz's problem on multiples of integer intervals. [diderot:7a4dbcd954dd] claims the Extended 1–2–3 Conjecture over every field. [diderot:8678929f755f] claims the optimal $4/3$ density bound expected by Holdridge and Pach, by an elementary argument with no computer search. Levicán-Santibáñez and Yoshida [diderot:ae0f59f50adf] claim a counterexample to Warnaar and Zudilin's q-rious positivity conjecture. Their Lean formalisation takes a cited analytic estimate as a hypothesis, so the check is only as good as that estimate. For this reader, the nearest point of contact is the first paper: it is a sparsity lower bound for polynomial multiples over $\mathbb{F}_2$ and beyond. None of it is refereed.

THEMES Pilz's problem on multiples of integer intervals · conjectures settled with AI co-authors

CLAIM OF THE WEEK Cresswell claims that, over every field and for all nonzero weights $c_1, \dots, c_n$, every nonzero polynomial or Laurent-polynomial multiple of $\sum_{a=1}^n c_a \prod_p x_p^{v_p(a)}$ has at least n terms, which includes Pilz's Extended 1–2–3 Conjecture as the case $\mathbb{F}_2$ with all weights 1. [diderot:7a4dbcd954dd]

Pilz's Extended 1–2–3 Conjecture over every field: sparse multiples of weighted integer intervals Ben Cresswell · with openai/gpt-5.6-sol, openai/gpt-6-astra, openai/gpt-6-astra-pro [diderot:7a4dbcd954dd]

Cresswell claims a proof of Pilz's Extended 1–2–3 Conjecture in a weighted form valid over every field: for all nonzero $c_1, \dots, c_n$, every nonzero polynomial or Laurent-polynomial multiple of $\sum_{a=1}^n c_a \prod_p x_p^{v_p(a)}$ has at least n terms.

The optimal 4/3 density bound in Pilz's problem Ben Cresswell · with openai/gpt-5.6-sol, openai/gpt-6-astra, openai/gpt-6-astra-pro [diderot:8678929f755f]

Cresswell claims that the limiting density $h(A)$ of integers with an odd number of representations aj is at least $4/3$ unless $A = \{a\}$ or $A = \{a, 2a\}$, improving the $8/7$ of Holdridge and Pach to the best possible constant, by an elementary proof with no computer search.

A q-rious counterexample Fabián Levicán-Santibáñez, Atsuro Yoshida · with OpenAI: GPT-6 Astra, OpenAI: GPT-6.1 Sol [diderot:ae0f59f50adf]

Levicán-Santibáñez and Yoshida claim to disprove Warnaar and Zudilin's q-rious positivity conjecture by showing that certain q-factorial-ratio polynomials P satisfy $|P(z)| > P(|z|)$ at a complex point, which forces a negative coefficient; their Lean formalisation takes an estimate of de la Bretèche and Tenenbaum as a hypothesis.

Zenodo 10 papers

A general, unmoderated repository. Shown here: the week's complexity preprints found there that the desk rated above its bar. Whether AI took part is known only when a record says so; the system is then named. Nothing here is refereed. Read with care.

10	45	1	OpenAI	universal algebra
PREPRINTS THIS WEEK	MORE LEFT OUT, BELOW THE BAR	ON THE FRONT PAGES OR IN DEPTH	MOST CREDITED SYSTEM (1 OF 10)	MOST COMMON TOPIC (2 PAPERS)

Ten records this week. The note on what the ETH refutations would mean for proof complexity is covered in a feature. For this reader the Gkantzounis pair matters most. [zenodo:24222968] claims subpower membership is in polynomial time for every fixed family of finite algebras with a common cube term, a question that goes back to Idziak, Marković, McKenzie, Valeriote and Willard. Its companion [zenodo:24223905] claims uniform $N^{O(d)}$ algorithms. Both rest on an earlier manuscript by the same author. [zenodo:24108086] reports an exact DPLL tree size on Tseitin formulas. [zenodo:24122330] reports certified random k-SAT threshold bounds for $k = 4$ to 7. Two NP-completeness claims, for width-4 posets [zenodo:24209873] and mim-width one [zenodo:24074845], make up the rest. The algebra pair should be checked first.

THEMES subpower membership with cube terms · NP-completeness at small fixed width · SAT search and thresholds

CLAIM OF THE WEEK The paper claims that subpower membership is decidable in deterministic polynomial time for every fixed finite family of finite algebras with a common cube term, and so for every finite algebra with few subpowers, while no $f(d) \cdot N^{O(1)}$ algorithm exists for algebras given as input unless $\mathrm{FPT} = \mathrm{W}[2]$. [zenodo:24222968]

Polynomial-Time Subpower Membership for Finite Algebras with Cube Terms: Fixed Families and Limits of Uniformity

Gkantzounis, Asterios [zenodo:24222968]

Claims that subpower membership is decidable in polynomial time for every fixed finite family of finite algebras with a common cube term, and hence for every finite algebra with few subpowers; it also claims that when the algebras are part of the input, no $f(d) \cdot N^{O(1)}$ algorithm exists unless $\mathrm{FPT} = \mathrm{W}[2]$, already for two-element near-unanimity algebras.

Uniform Subpower Membership for Finite Algebras with Cube Terms of Fixed Dimension

Gkantzounis, Asterios [zenodo:24223905]

Claims uniform algorithms for subpower membership with algebras given by operation tables and a cube term of fixed dimension d: time $N^{O(d)}$ when an edge operation is available or the algebras are idempotent, and $N^{O(d^2)}$ when a cube term is only promised to exist.

The Price of a Fixed Order: Exact Size of Static-Order DPLL Trees on Tseitin Formulas

Jagielski, Jan [zenodo:24108086]

Claims an exact formula for the number of leaves of the DPLL search tree with unit propagation and a static variable order on unsatisfiable Tseitin formulas, and shows that the tree's shape depends only on the graph and the order, not on the charges.

Upper bounds on the random k-SAT threshold for k = 4, 5, 6, 7 from exact interpolation certificates

Pei, Yan Ru [zenodo:24122330]

Claims computer-certified upper bounds on the random k-SAT threshold of $9.9325, 21.1175, 43.372$ and 87.786 for $k = 4, 5, 6, 7$, each within 0.002 of the cavity-method prediction, by extending a recent interpolation bound for random 3-SAT.

The linear extension diameter is NP-complete for posets of width 4

Moreno-Mejía, Erik-Josué, Jiménez-Hernández, Hugo [zenodo:24209873]

Claims that computing the linear extension diameter is NP-complete for posets of width 4, answering a question of Brightwell and Massow; with their polynomial algorithm for width 3, this gives a dichotomy at width 3 unless $\mathrm{P} = \mathrm{NP}$.

Recognizing Mim-Width and Its Variants Is NP-Complete at Every Fixed Width

Kadri, Haitam [zenodo:24074845]

Claims that deciding whether a graph has mim-width at most k is NP-complete for every fixed $k \geq 1$, and likewise for sim-width, one-sided mim-width and Omim-width, via a reduction from NAE-3SAT that settles the open width-one case.

Adversarial bandit convex optimisation: regret d^(3/2) sqrt(n) with polynomial-time steps on a real RAM, and a matching lower bound

mt0-svg · with AI-assisted (system not named) [zenodo:24147205]

Claims a learner for adversarial bandit convex optimisation on rational polytopes whose rounds take polynomial time and whose regret is $d^{3/2}\sqrt{n}$ up to logarithms against an adaptive adversary, together with a matching lower bound, as conjectured by Bubeck, Eldan and Lee.

Resolving Goldreich's Open Problem 4.1: Linear Lower Bounds for Cubic Graph Isomorphism Testing

Schubr, Scott [zenodo:24293816]

Claims that testing isomorphism to a known connected asymmetric bridgeless cubic graph at proximity 0.01 requires a linear number of incidence queries, for both randomized classical and bounded-error quantum testers, which answers Goldreich's Open Problem 4.1 in the negative.

Offline Reachability in Directed Acyclic Graphs under Vertex Failures

Shaik, Idris Ali [zenodo:24142496]

Claims a tunable algorithm running in $\tilde{O}(S + E\tau + (V + E)q_b/\tau)$ time for deciding, in a DAG with a common source, whether any query with at most two vertex failures still reaches one of its targets, a problem shown linearly equivalent to a labelled-path avoidance problem.

ODDS AND ENDS

NUMBERS OF THE WEEK

$\exp(\Omega((\log n)^2/\log\log n))$
LOWER BOUND FOR WEAK PIGEONHOLE IN $\Sigma\Pi\Sigma\Pi$-PC

$2^{2^{O(k)}}|\varphi|$
TIME TO FIND A STRONG BACKDOOR TO FORESTS

$O(Q(f)^6/(\varepsilon\delta)^2)$
CLASSICAL QUERIES IN THE AARONSON–AMBAINIS BOUND

6
AND GATES FOR THE THREE-BIT LOW-HALF MULTIPLIER

35
NEW LOWER BOUND, RANK OF 4×4 MATRIX MULTIPLICATION OVER $\mathbb{F}_2$ (AI-ASSISTED)

0.17843
NEW RANGE OF ε FOR THE LOPSIDED TRIANGLE METHOD (AI-ASSISTED)

COMPLEXITY FORECAST

Polynomial calculus, depth three and four: lower bounds, exponential, moderate later. $AC^0[2]$-Frege, every fixed depth: hard instances, first in the area, holding. Exponential Time Hypothesis: refutation claimed, galactic, no change at physical sizes. Polynomial hierarchy: collapsing to OM if $\mathrm{NP}\subseteq\mathrm{P/poly}$, otherwise unchanged.

CLASSIFIEDS

- FOR SALE: Maltsev polymorphism test, polynomial time. Buyer must supply binary relational structure.
- WANTED: rank-46 scheme for 4×4 matrix multiplication over $\mathbb{F}_2$. Must share at most 28 terms with existing stock.
- CLEARANCE: Set Disjointness lower bounds, small universes. Some no longer supported. Sold as seen.
- OFFERED: subexponential 3-SAT algorithm. Galactic. Beats exhaustive search at no size anyone has checked.
- TO LET: recursive oracle, separates EQP from IP and PH. Not for use in unrelativised worlds.
- WANTED: proof that 3-Edge Colouring on planar graphs is NP-hard, or that it is not. Long-standing vacancy.
- FOR SALE: Go board, Chinese rules, superko observed. History supplied as a circuit. EXPSPACE-complete.

Complexity Chronicles is written by a language model (anthropic: claude-opus-5-5) from the abstracts of 75 papers (arXiv cs.CC, ECCC, Hexagon, Diderot, Zenodo) that appeared during 9–15 November 2026, selected according to a reader profile. Generated 2026-11-20 15:09 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.