

An Explicit 2×2 -Minimal Forcing Certificate at $S = 7/4$ for the Arithmetic–Kakeya Sum–Difference Problem, and Evidence for a Floor on What the Certificate Framework Can Express

Cultural Soliton Observatory / ProjectForty2

2026-06-14

Abstract

We report an explicit, machine-verified *forcing certificate* achieving sum–difference score $S = 7/4 = 1.75$ — the FrontierMath “arithmetic–Kakeya” warm-up threshold — and we examine what the surrounding certificate framework can and cannot prove. The certificate lives on the smallest non-trivial grid (2×2), uses no free vertices, and realizes the Katz–Tao slope set $\{0, 1, 2, \infty\}$. Its correctness is established by an *explicit linear-algebra witness*: the target direction $(1, -1)$ is realized exactly at the corner vertex $(2, 2)$ — by a written-out rational combination of the certificate’s generators that is zero at the other three vertices — after which the remaining three vertices cascade, each via its own explicit witness. The result is corroborated by a reconstructed “gold” verifier together with three independent oracles and a tamper control.

We are deliberate about what this is **not**. The value $7/4$ is **not a new mathematical bound**: it is the Katz–Tao sum–difference bound of 2001 (the $\{0, 1, 2, \infty\}$ slope-set bound). We do **not** beat the arithmetic–Kakeya record; the best known sum–difference exponent, $\inf \text{SD}(\mathbb{R}; -1) \leq 1.67513\dots$ (Katz–Tao, restated by Tao 2025), is untouched. Our contributions are (i) the explicit certificate realizing $7/4$ that is *minimal on the 2×2 grid* (proven by exhaustive search) — improving the best small forcing certificate *we had previously constructed*, which scored $S = 9/5$ on a 2×3 grid; and (ii) a precisely-scoped *evidence-for-a-floor* result: a structural argument, exhaustive small-case search, and the Tao (2025) bounded-slope theorem together suggest that the framework’s achievable exponents floor at $\gamma = 1.6751\dots$, the largest real root of $x^3 - 4x + 2$. We are explicit that this is a *characterization supported by evidence, not a theorem*, and we name the universal-reduction step a full proof would require. The practical consequence is that $7/4$ is the exact 2×2 minimum and the second term of the advanced iteration, and reaching even the interval $(3/2, \gamma)$ would require both a new inequality and a new certificate format.

1 Introduction

The **arithmetic Kakeya problem** asks how slowly a sum set can grow relative to its difference set under a Kakeya-type incidence constraint, and is one of the cleaner combinatorial cores underneath the geometric Kakeya conjecture. The relevant quantity is the **sum–difference exponent** $\text{SD}(\mathbb{R}; -1)$, the smallest exponent for which a certain sum-vs-difference inequality survives. The *conjectural* endpoint of the arithmetic Kakeya conjecture (AKC) corresponds to exponent 1; the best known upper bound is $\inf \text{SD}(\mathbb{R}; -1) \leq 1.67513\dots$, due in mechanism to Katz–Tao (2001) [?] and restated in Tao’s November 2025 note [?].

The Epoch **FrontierMath** benchmark includes an “arithmetic–Kakeya” track formulated as a **forcing-certificate** verification problem: a candidate submits a finite combinatorial object (X, d, f, T, R) which a verifier accepts iff a *forcing closure* reaches every vertex of a product grid

and a rational score S falls at or below a threshold τ . The score S equals a certified sum–difference exponent, so a passing certificate at threshold τ is a machine-checkable witness that SD at the certified configuration is at most τ . The benchmark’s warm-up threshold is $\tau = 7/4$; higher tiers push toward the record and ultimately toward the conjectural endpoint.

This note does two things, both stated modestly.

Contribution 1 (an explicit 2×2 -minimal artifact). We give an explicit forcing certificate that passes at $\tau = 7/4$ on the 2×2 grid with no free vertices (§??). It is the **minimum-scoring configuration on the 2×2 grid** (exhaustive search, §??) and improves the best small forcing certificate *we had previously constructed*, which scored $S = 9/5 = 1.8$ on a 2×3 grid and required a free vertex and a telescoping ladder. We prove correctness by exhibiting an explicit rational combination of the certificate’s generators that realizes the target direction $(1, -1)$ exactly at the corner vertex $(2, 2)$ from the empty free set, and then exhibiting one such combination for each of the remaining three vertices. The value $7/4$ is *not new mathematics* (it is the Katz–Tao $\{0, 1, 2, \infty\}$ bound); the contribution is the explicit minimal realization and its end-to-end verification.

Contribution 2 (evidence for a framework floor, scoped honestly). We examine the reach of the certificate framework itself (§??). We give (a) a structural argument that a certificate is a product grid whose axis projections are the Katz–Tao sum–difference scaffold; (b) an **exhaustive** search of the 2×2 grid — with a closed-form configuration count that we verify equals the iterated count — showing minimum score exactly $7/4$, *invariant* under enlarging the slope alphabet from 4 to 7 slopes; and (c) the Tao (2025) bounded-slope theorem, which forces $\text{SD} \rightarrow 2$ when the number of slopes is bounded. Together these give structural and computational evidence — **but not a proof** — that the framework is co-floored by finite sum–difference constructions, whose exponents are bounded below by $\gamma = 1.6751308705666\dots$, the largest real root of $x^3 - 4x + 2$. The honest consequence (§??): within this framework $7/4$ is the exact 2×2 minimum, and even reaching $(3/2, \gamma)$ would require a construction the framework as formatted cannot express.

We do **not** claim a new arithmetic–Takeya bound, and we do **not** claim $7/4$ as a new bound. The record stands. All verifier results are reported against a *reconstruction* of the Epoch verifier (§??), and there is no public submission channel, so we frame $7/4$ as a verified research artifact rather than an official competition solve.

2 The forcing-certificate framework

2.1 The certificate object

A **certificate** is a tuple (X, d, f, T, R) .

- **Slope alphabet** $X \subset \mathbb{Z}^2$. A finite set of slope vectors. It must contain $(0, 0)$, and every other $(a, b) \in X$ must satisfy $a + b \neq 0$. The forbidden direction $a + b = 0$ — represented by the canonical target $(1, -1)$ — is exactly the sum–difference direction one is trying to *synthesize*, so it may never appear as an input slope.
- **Grid** $d = [d_1, \dots, d_k]$. The vertex set is the product

$$V = \{1, \dots, d_1\} \times \dots \times \{1, \dots, d_k\}, \quad n = \prod_{i=1}^k d_i.$$

- **Edge labels** $f = (f_1, \dots, f_k)$. Each f_i assigns a slope in X to a length- i key whose last coordinate lies in $\{1, \dots, d_i - 1\}$; it labels the edges that move along axis i . Omitted keys default to $(0, 0)$.
- **Free set** $T \subseteq V$. Vertices granted “for free” at the start of the closure (they enlarge the denominator of the score and so are penalized).
- **Seed relations** R . A list of **single-support** relations, each a map (one vertex $\mapsto$ one slope in X); each contributes exactly one nonzero entry.

2.2 Score

The cost of the edge labelling is

$$m = \sum_{i=1}^k (\#\{\text{nonzero } f_i\}) \cdot \prod_{j=i+1}^k d_j,$$

each axis- i nonzero edge weighted by the suffix product $d_{i+1} \cdots d_k$. The **score** is

$$S = \frac{m + |R|}{n - |T|}$$

and equals the certified sum-difference exponent $\text{SD}(X; -1)$ for the configuration.

2.3 op1 generators and the module

For each axis- i edge labelled $x \in X$ between two adjacent vertices, the **op1 generator** places x at one endpoint and $-x$ at the other: schematically $\{e_1 : x, e_2 : -x\}$. The seed relations R together with all op1 generators span a rational subspace

$$M = \mathbb{Q}\text{-span}(R \cup \{\text{op1 generators}\}) \subseteq \mathbb{Q}^{2n},$$

where each vertex contributes two coordinates (a 2-vector slope).

2.4 Forcing closure (op2)

For a current free set $F \subseteq V$ and a vertex $e \notin F$, the **antidiagonal direction** at e is a vector $w \in \mathbb{Q}^{2n}$ whose restriction to vertex e is $(s, -s)$ with $s \neq 0$ (i.e. $u + v = 0$, $u \neq 0$ on the two coordinates of e — exactly the forbidden $a + b = 0$ direction). We say e is **forced given** F iff

$$M \text{ (a subspace of } \mathbb{Q}^{2n} \text{) contains } w \text{ with } w|_e = (s, -s), s \neq 0, \text{ and } w|_v = (0, 0) \forall v \notin F \cup \{e\}.$$

Intuitively: the subspace can manufacture the forbidden $(1, -1)$ direction at e while leaving everything outside the already-free region untouched. This is precisely the predicate the verifier computes (its **_op2_can_force**): solve for a generator combination y with (yM) zero on all “outside” coordinates and $(yM)|_e = (1, -1)$ (we may fix $s = 1$, as the line $\{(s, -s) : s \neq 0\}$ is closed under nonzero scaling).

The **closure** starts from $F = T$ and repeatedly adds every forceable vertex until a fixpoint. The certificate **PASSES at threshold** τ iff the closure reaches all n vertices and $S \leq \tau$. A passing certificate at τ certifies $\text{SD}(X; -1) \leq \tau$ for that configuration, with the forbidden $(1, -1)$ direction synthesized from the allowed slopes alone.

3 The 2×2 -minimal $7/4$ certificate

3.1 The certificate

$$\begin{aligned}
X &= \{(0, 0), (1, 0), (1, 1), (1, 2), (0, 1)\} & (\text{slopes } \{0, 1, 2, \infty\}) \\
d &= [2, 2] & (n = 4) \\
f &= [\{(1,) : (1, 0)\}, \{(1, 1) : (1, 1)\}, \{(2, 1) : (1, 1)\}] \\
T &= [] & (|T| = 0) \\
R &= [\{(1, 1) : (1, 2)\}, \{(1, 2) : (0, 1)\}, \{(2, 1) : (0, 1)\}] & (|R| = 3)
\end{aligned}$$

The slope vectors $\{(1, 0), (1, 1), (1, 2), (0, 1)\}$ are exactly the Katz–Tao slopes $\{0, 1, 2, \infty\}$. The grid is the smallest non-trivial one, 2×2 , with no free vertices.

3.2 Arithmetic of the score

The edge cost: f_1 has one nonzero label weighted by suffix product $d_2 = 2$, contributing 2; f_2 has two nonzero labels weighted by suffix product 1, contributing 2. Hence $m = 2 + 2 = 4$, $|R| = 3$, $n = 4$, $|T| = 0$, and

$$S = \frac{m + |R|}{n - |T|} = \frac{4 + 3}{4 - 0} = \frac{7}{4} = 1.75.$$

The reconstructed verifier independently recomputes $m(G) = 4$ and reports $S = (4 + 3)/(4 - 0) = 7/4$.

3.3 Correctness: explicit forcing witnesses

The four vertices are $(1, 1), (1, 2), (2, 1), (2, 2)$. The subspace M is spanned by 7 generators and has rank 7 out of the 8 ambient coordinates (it “wastes no degrees of freedom”): the three seeds and four op1 generators

$$\begin{aligned}
g_0 &= R[(1, 1) : (1, 2)], \quad g_1 = R[(1, 2) : (0, 1)], \quad g_2 = R[(2, 1) : (0, 1)], \\
g_3 &= \text{op1}[(1,) \rightarrow (1, 0)] @ (1, 1), \quad g_4 = \text{op1}[(1,) \rightarrow (1, 0)] @ (1, 2), \\
g_5 &= \text{op1}[(1, 1) \rightarrow (1, 1)] @ (1, 1), \quad g_6 = \text{op1}[(2, 1) \rightarrow (1, 1)] @ (2, 1).
\end{aligned}$$

We exhibit, for each forced vertex, an explicit rational combination realizing the antidiagonal $(1, -1)$ direction there and vanishing outside the permitted region. These are the actual solution vectors, verified by direct substitution — not narrative “telescoping.”

Round 1 — the corner $(2, 2)$ is forced first, and uniquely, from $F = \emptyset$. The combination

$$w_{(2,2)} = g_0 - 2g_1 - g_2 + g_3 - 2g_4 - 2g_5 + g_6$$

evaluates to $w_{(2,2)}|_{(1,1)} = w_{(2,2)}|_{(1,2)} = w_{(2,2)}|_{(2,1)} = (0, 0)$ and $w_{(2,2)}|_{(2,2)} = (1, -1)$. Thus the forbidden target direction is synthesized at $(2, 2)$ alone, from the allowed slopes $\{0, 1, 2, \infty\}$ — the heart of the certificate. A round-by-round driver over the verifier’s `_op2_can_force` predicate confirms $(2, 2)$ is the **unique** vertex forceable in round 1.

Round 2 — the remaining three cascade, each with its own explicit witness. With $(2, 2) \in F$, each remaining e becomes forceable:

$$\begin{aligned}
w_{(1,2)} &= -g_1 + g_4 & \Rightarrow w_{(1,2)}|_{(1,2)} &= (1, -1), \text{ zero at } (1, 1), (2, 1), \quad w_{(1,2)}|_{(2,2)} = (-1, 0), \\
w_{(2,1)} &= -g_0 + 2g_1 - g_2 - g_3 + 2g_4 + 2g_5 & \Rightarrow w_{(2,1)}|_{(2,1)} &= (1, -1), \text{ zero at } (1, 1), (1, 2), \quad w_{(2,1)}|_{(2,2)} = (-2, 0), \\
w_{(1,1)} &= -2g_0 + 3g_1 + 3g_4 + 3g_5 & \Rightarrow w_{(1,1)}|_{(1,1)} &= (1, -1), \text{ zero at } (1, 2), (2, 1), \quad w_{(1,1)}|_{(2,2)} = (-3, 0).
\end{aligned}$$

In each case the antidiagonal $(1, -1)$ appears at the target vertex, the value is zero at the two vertices still outside $F \cup \{e\}$, and the only other nonzero entry sits at the now-free $(2, 2)$ — which is allowed. The driver reports round 1: forced $\{(2, 2)\}$; round 2: forced $\{(1, 1), (1, 2), (2, 1)\}$, reaching all 4 of 4 vertices. Closure succeeds, and since $S = 7/4 \leq 7/4$ the certificate **PASSES** at $\tau = 7/4$. (Each w was verified by direct substitution into M ; see §??.)

3.4 The $9/5 \rightarrow 7/4$ improvement (over our own prior small certificate)

The best small forcing certificate *we had previously constructed* scored $S = 9/5 = 1.8$ on the same slope alphabet but a larger grid:

$$d = [2, 3], \quad n = 6, \quad m = 7, \quad |R| = 2, \quad |T| = 1, \quad S = \frac{7+2}{6-1} = \frac{9}{5}.$$

Here “small” means *the smallest explicit construction we had on hand* ($n \leq 6$), not a global record: the $9/5$ certificate proves only the *upper bound* $SD \leq 9/5$ on $d = [2, 3]$, and indeed our searches on that grid (§??) do not establish $9/5$ as a floor. Re-verification confirms the $9/5$ certificate passes at $\tau = 9/5$, forces all 6 vertices, and matches its arithmetic.

The mechanism of the improvement is structural collapse: the 2×3 ladder spends a free vertex (a $|T| = 1$ penalty) and seven edge labels to walk the forbidden direction down a three-step axis. On the 2×2 grid the corner $(2, 2)$ is forced *directly from the empty free set* — no ladder, no free vertex — so $|T|$ drops from 1 to 0 and m from 7 to 4, and the score falls from $\frac{9}{5}$ to $\frac{7}{4}$. The move is “find the 2×2 optimum, which lies below the previously-exhibited 2×3 construction.”

3.5 Multi-oracle verification and a tamper control

The certificate’s correctness was checked five ways (all re-run 2026-06-14):

1. **Gold verifier (reconstruction).** Returns PASS (exit 0) at $\tau = 7/4$ and REJECT (exit 1) at $\tau = 67/40 = 1.675$, as $7/4 > 67/40$ predicts.
2. **Independent Fraction-RREF oracle.** An exact rational row-reduction reproduces the closure and score; on the published certificate it agrees with the gold verifier (cross-check prints AGREE=True).
3. **Independent two-prime mod- p oracle.** Reducing the generators modulo two distinct primes reproduces the same forcing closure.
4. **Explicit witnesses.** The round-1 and round-2 combinations of §??, each verified by direct substitution.
5. **Tamper / negative control.** Replacing the seed $\{(1, 1) : (1, 2)\}$ with any of the three other allowed slopes breaks the closure: $(1, 1) \mapsto (1, 1)$ reaches $1/4$; $(1, 1) \mapsto (2, 1)$ reaches $0/4$; $(1, 1) \mapsto (1, 0)$ reaches $1/4$. The certificate is **delicately placed**.

Across the recorded audit (300–400 randomly generated small-certificate trials) the three independent oracles registered **zero disagreements** with the gold verifier. “Gold verifier” denotes a *reconstruction* of Epoch’s verifier semantics, not an official Epoch run; all PASS/REJECT statements are against the reconstruction.

4 Evidence for a framework floor

4.1 Katz–Tao sum–difference recursions and the constant γ

The Katz–Tao sum–difference method [?] improves its exponent by iterating a recursion. The **basic** recursion is $\beta \mapsto (4\beta - 1)/(2\beta)$, whose fixed-point equation reduces (verified exactly) to $-2\beta^2 + 4\beta - 1 = 0$, with positive fixed point

$$\beta_{\text{basic}}^* = 1 + \frac{\sqrt{2}}{2} = 1.7071067811865\dots$$

The **advanced** recursion is $\beta \mapsto (3\beta^2 + 2\beta - 2)/(\beta^2 + 3\beta - 2)$. Setting this equal to β and clearing denominators reduces *exactly* (sympy `cancel/fraction/expand`) to

$$-\beta^3 + 4\beta - 2 = 0 \iff \beta^3 - 4\beta + 2 = 0.$$

This cubic has three real roots, $-2.2143\dots$, $0.5392\dots$, and

$$\gamma = 1.6751308705666460709\dots,$$

the largest, the attracting fixed point of the advanced iteration. Substituting γ into $\beta^3 - 4\beta + 2$ gives 0 to machine precision. Comparing to the FrontierMath rejection threshold $67/40 = 1.675$,

$$\gamma - \frac{67}{40} = 1.30871 \times 10^{-4} > 0,$$

so $\gamma > 67/40$ by a small but strictly positive margin.

4.2 The advanced iteration descends to γ from above

Starting the advanced iteration at $\beta = 2$ produces the exact decreasing rational sequence

$$2, \frac{7}{4}, \frac{171}{101}, \frac{101863}{60652}, \dots = 2, 1.75, 1.69307, 1.67947, 1.67618, 1.67539, 1.67519, \dots$$

converging to γ **from above**. The first two terms are exactly 2 and $7/4$. Hence every exponent produced by the advanced iteration is strictly greater than $\gamma > 1.675$, and $7/4$ is the *second* term — the first non-trivial improvement off the base case $\beta = 2$. (This is a statement about the iteration; whether *all* finite sum–difference constructions are iteration outputs is the stronger claim conjectured in §??, not proved here.)

4.3 Exhaustive 2×2 search and slope-count invariance

The load-bearing computational evidence is an **exhaustive** search of the 2×2 grid under the verifier’s forcing semantics, using the fast Fraction-RREF oracle (cross-checked `AGREE=True` against the gold verifier on the published certificate). To make “exhaustive” defensible, we fix a precise enumeration convention and verify the configuration count against a closed form.

Enumeration convention. On $d = [2, 2]$ the vertices are $V = \{(1, 1), (1, 2), (2, 1), (2, 2)\}$ ($n = 4$) and the verifier’s f -domain has exactly $E = 3$ edge slots: the level-1 key $(1,)$ and the level-2 keys $(1, 1), (2, 1)$. With s nonzero slopes, each slot takes a value in $X = \{(0, 0)\} \cup (\text{slopes})$, i.e. $(s + 1)$ choices, so there are $(s + 1)^E = (s + 1)^3$ edge labellings. The free set T ranges over all subsets of V ($2^4 = 16$). The seed list R ranges over all size-0, $\dots$, 3 subsets of the $4s$ vertex–slope pairs (we cap $|R| \leq 3$; the published certificate uses $|R| = 3$). The total count is the closed form

$$N(s) = (s + 1)^3 \cdot 2^4 \cdot \sum_{r=0}^3 \binom{4s}{r}.$$

For each s we verify that the iterated enumeration visits exactly $N(s)$ configurations (the script asserts `MATCH=True`):

slopes s	$(s + 1)^3$	$\sum_{r \leq 3} \binom{4s}{r}$	configs $N(s)$	min score	time
4 $\{0, 1, 2, \infty\}$	125	697	1,394,000	7/4	24.2 s
5	216	1,351	4,669,056	7/4	71.7 s
6	343	2,325	12,759,600	7/4	174.5 s
7	512	3,683	30,171,136	7/4	376.3 s

In every case the closed-form count equals the iterated count, the minimum is exactly 7/4, and the minimizing configuration is exactly the published certificate. **Enlarging the slope alphabet from 4 to 7 slopes adds no forcing power on the 2×2 grid.**

Three scope caveats are required. **(i)** “Exhaustive” is asserted *for this enumeration convention*, made checkable by the closed form $N(s)$; a different convention (counting symmetry orbits, or allowing $|R| > 3$) would give a different total. We report a single reproducible convention with its verified count, and rely on no earlier, unreconciled configuration-count figure. **(ii)** The claim is for $d = [2, 2]$ *only*; it does not prove that no larger grid or alphabet anywhere beats 7/4. **(iii)** The slope-count invariance is *consistent with* — but not a proof of — the Tao (2025) bounded-slope phenomenon (§??).

4.4 Heuristic upper bounds on larger grids (search-found, not floors)

We also ran heuristic (time-budgeted, Dinkelbach-style) searches over a range of grids. These return the **best certificate found** — i.e. *upper bounds* on the achievable score for each grid, not floors — and the searches are non-exhaustive and disagree at the margins. The recorded best-found scores (logs cited in §??) are:

configuration	best certificate found (upper bound)
$d = [2, 3]$	$\leq 9/5$ (explicit cert; the time-budgeted floor probe only reached 2)
$d = [3, 3]$	$\leq 13/7$
$d = [2, 2, 2]$	$\leq 13/7$
$d = [2, 2, 3]$	$\leq 19/10$

Nothing here establishes a *lower bound* (floor) on any grid other than $d = [2, 2]$ (§??). In particular our earlier phrasing “ $d = [2, 3]$ floor = 9/5” was **incorrect** — the time-budgeted floor search on $d = [2, 3]$ in fact returned 2.0, and 9/5 is only the score of an explicit certificate (an upper bound). These multi-grid numbers are **corroborating context, not part of the evidence chain** for the floor claim, and were produced by heuristic search, not re-verified exhaustively. They indicate only that no small grid we searched produced a certificate below 7/4.

4.5 The framework-floor characterization (evidence, not a theorem)

We bring the pieces together into a **characterization** of the framework’s reach, explicit that it is *supported by evidence and not proven*.

Characterization 1 (informal, not a theorem). *Every forcing certificate is plausibly an instance of a finite sum–difference construction; if so, the framework’s achievable exponents are bounded below by $\gamma = 1.6751\dots$ (§??–??), and no certificate of the present format can certify an exponent in $[1, \gamma)$.*

This is the *directional* claim (certificate $\Rightarrow$ finite-SD instance, hence floored by γ); we do not assert the converse (every SD construction is a certificate), which we neither need nor argue. The supporting evidence:

- **(a) Structural.** A certificate is a product grid whose axis projections coincide with the Katz–Tao sum–difference scaffold: each axis carries an edge-labelled difference structure, the op1 generators are the sum–difference pairings $\{x, -x\}$, and the closure synthesizes the forbidden $(1, -1)$ direction exactly as the sum–difference method synthesizes the antidiagonal. The score S is the certified SD exponent by construction. So a certificate is, structurally, an explicit finite instance of a sum–difference construction.
- **(b) Computational.** The exhaustive $d = [2, 2]$ minimum is exactly $7/4$ and is *invariant* to slope-count up to 7 slopes (§??). No configuration in this exhaustive sweep drops below the iteration’s second term $7/4$. (The heuristic multi-grid upper bounds of §?? are consistent with this but, being non-exhaustive, are not part of this evidence list.)
- **(c) Literature — bounded slopes force $SD \rightarrow 2$.** Tao’s November 2025 paper [?] shows that with a *bounded* number of slopes the sum–difference exponents converge to 2. Every certificate uses a finite (hence bounded) slope alphabet, so the bounded-slope regime is exactly the regime certificates inhabit — consistent with the observed slope-count invariance.
- **(d) Literature — no obvious sub-floor lever.** The two natural strengthening routes do not obviously provide a lever below the floor. Pohoata–Zakharov [?] introduce a general way to *strengthen* the Katz–Tao sum–difference inequalities and study a higher-dimensional variant; their abstract describes an iterative upper-bound argument and a higher-dimensional Besicovitch corollary, and we did not find in it a mechanism that lowers the planar sum–difference exponent below γ . Cowen-Breen et al. [?] establish *equivalences* relating arithmetic-Keakeya pattern problems (and prove AKC implies the packing-dimension Keakeya conjecture); it is an equivalence/implication paper rather than a source of a smaller upper-bound exponent. (We verified the titles, authors, and high-level content of [?] and [?]; we did not verify line-by-line that neither contains a sub- γ lever, so we state this as “no obvious lever found,” not a proof.)

What a full proof would require — and what we do not establish. To upgrade this to a theorem one would have to show that **every** forcing certificate — over **all** grids, **all** dimensions, and **all** finite slope alphabets — reduces to a finite sum–difference construction (and hence cannot certify below γ). We do **not** establish this. We therefore do **not** write “we prove the framework equals the sum–difference method.” We write: we give structural, exhaustive-small-case, and literature evidence that the framework’s achievable exponents floor at γ , and we state precisely the universal-reduction step a full proof would require.

5 Consequences

7/4 is the exact 2×2 minimum and the second iteration term. It is the second term of the advanced iteration $(2, 7/4, 171/101, \dots)$ and the exhaustive minimum on the smallest grid (§??), stable to slope-count enrichment up to 7 slopes. The next iteration terms $(171/101 \approx 1.693, \text{ then } 1.679, \dots)$ descend toward γ without ever reaching it.

Sub- γ would require leaving the framework. If the §?? characterization holds, the framework’s exponents floor at $\gamma > 67/40$, so certifying any exponent in $[1, \gamma)$ — including the AKC endpoint 1 (itself conjectural) and even the open band $(3/2, \gamma)$ — cannot be done within the certificate format as defined. The absolute floor of the sum–difference *method* is $3/2$, so $[3/2, \gamma)$ is **not ruled out in principle**, but it is **unreached**: it would require both a new inequality (stronger than sum–difference) *and* a new certificate format able to express it.

The warm-up/full-tier gap. The warm-up threshold $7/4$ is reachable by a tiny, fully explicit, machine-verifiable certificate — the artifact of §??. The next tiers, which approach the record $(1.67513\dots)$ and ultimately the endpoint (1) , do not appear reachable by enlarging the same certificate; the (conjectured) floor at γ would mean incremental grid/alphabet growth asymptotes at γ from above. As an *observation*, the published record $1.67513\dots$ and $\gamma = 1.67513087\dots$ coincide to the digits reported — consistent with the advanced iteration being the operative recursion behind the record — but we do *not* derive an identity “record = γ ” from our framework results; that is a statement about Katz–Tao/Tao’s optimization, not something this paper establishes.

What we did not do. We did not beat the arithmetic–Kakeya record (it stands at $\leq 1.67513\dots$), and $7/4$ is not new mathematics — it is the Katz–Tao $\{0, 1, 2, \infty\}$ bound made into the minimal explicit certificate.

6 Reproducibility

All facts were re-verified on 2026-06-14 with Python 3.12.3 and sympy in a fixed local research environment.

Artifacts (relative research artifact locations). *Verifier and certificates* — `.../dgx_handoff/pipelines/00_gold_verifier_kakeya_verifier.py`; the $7/4$ certificate `SOLVE_kakeya_7over4.cert`; the $9/5$ certificate `best_forcing_cert_9_5.txt`; fast Fraction-RREF oracle `kakeya_fast.py`; explicit witnesses + round-by-round trace + tamper control `round_witnesses.py`, `trace_tamper_witness.py`. *Floor evidence* — `.../dgx_handoff/pipelines/00_kakeya_sub1675/wf_dinkelbach/`: exhaustive 2×2 enumerator (with closed-form count assertion and `AGREE=True` cross-check) `exhaust2x2_clean.py` and log `exhaust2x2_clean.log`; γ closed-form/sequence check `gamma_check.py`; heuristic multi-grid upper-bound searches and logs `full_2x3.py/full_2x3.log`, `sweep_richX.py/sweep_richX.log`, `sweep_3d.py/sweep_3d.log` (the §?? upper bounds; `full_2x3.log` records the $d = [2, 3]$ floor probe returning 2.0).

How to check the certificate.

```
python kakeya_verifier.py SOLVE_kakeya_7over4.cert --threshold 7/4      # PASS, exit 0
python kakeya_verifier.py SOLVE_kakeya_7over4.cert --threshold 67/40   # REJECT, exit 1
```

The CLI takes a **fraction** ($7/4$, $67/40$); a bare decimal *errors* (the parser does `Fraction(int(num), int(den))`, so "1.675" raises). The verifier prints the validated slope set, the grid and $n, m(G)$, the seed/free-set validations, the closure progress (`closure reached 4/4 vertices`), and the score line. The certificate file is six lines (line 1 a human-readable header ignored by the parser; then X, d, f, T, R).

The forcing witnesses of §?? (round-1 combination for $(2, 2)$ and round-2 combinations for the other three vertices), the rank-7 fact, and the tamper control are produced and verified by `round_witnesses.py` / `trace_tamper_witness.py`, which drive the verifier’s internal `build_module_matrix` and `_op2_can_force` routines. The verifier’s `run_closure` returns only the final reached set and emits no per-round output; the round structure is reconstructed by a driver over the `_op2_can_force` predicate, not printed by the verifier. The exhaustive 2×2 search of §?? is driven by `exhaust2x2_clean.py`, which both asserts the iterated configuration count equals the closed form $N(s)$ and prints `AGREE=True` from a direct comparison of the fast oracle and the gold verifier on the published certificate.

Reconstruction caveat (required disclosure). The “gold verifier” `akeya_verifier.py` is a **reconstruction** of Epoch’s FrontierMath verifier semantics, audited against three independent oracles (exact Fraction-RREF, two-prime mod- p , and explicit witnesses) with **zero disagreements** over the recorded 300–400 randomly-generated small-certificate trials. All PASS/REJECT results are **against the reconstructed verifier**, not an official Epoch run. **No public submission channel exists** for the FrontierMath arithmetic-Kekeya track, so the $7/4$ certificate is a verified research artifact, *not* an official competition solve.

Citation status. We verified the arXiv identifier, title, author list, and abstract-level content of each reference (June 2026). We did **not** verify the internal theorem/corollary numbers of [?] against the paper body; those labels are reproduced from the task brief and should be confirmed before being relied upon. Where a specific sub-claim of a cited paper is load-bearing and we could not confirm it from the abstract, we softened the in-text statement accordingly (§??(d)).

References

- [1] N. H. Katz and T. Tao, *New bounds on Kekeya problems*, arXiv:math/0102135 (2001). Sum-difference method; the $\{0, 1, 2, \infty\}$ slope-set bound $SD = 7/4$; the basic iteration $\beta \mapsto (4\beta - 1)/(2\beta)$ and the advanced iteration $\beta \mapsto (3\beta^2 + 2\beta - 2)/(\beta^2 + 3\beta - 2)$. (Specific theorem/corollary numbers not re-verified against the paper body.)
- [2] T. Tao, *Sum-difference exponents for boundedly many slopes, and rational complexity*, arXiv:2511.15135 (19 November 2025). The current record $\inf SD(\mathbb{R}; -1) \leq 1.67513$; with a bounded number of slopes the exponents converge to 2. (Title, author, date, record value, and bounded-slope statement verified from the abstract; no internal theorem number cited.)
- [3] C. Pohoata and D. Zakharov, *Generalized Arithmetic Kekeya*, arXiv:2411.13395 (20 November 2024). A general way to strengthen the Katz–Tao sum-difference inequalities; a higher-dimensional variant with a Besicovitch Minkowski-dimension corollary. (Title, authors, date, and high-level content verified from the abstract; the specific sub-claims “ $\beta_h = \beta$ equivalence” and “ $\beta(\mathbb{R}^d) \geq d$ ” were not confirmed and are not asserted.)
- [4] C. Cowen-Breen, E. Karagozishvili, N. Varadarajan, and T. Wang, *Pattern Problems related to the Arithmetic Kekeya Conjecture*, arXiv:2011.07056 (2020). Equivalences among arithmetic-Kekeya pattern problems; AKC implies the Kekeya conjecture for packing dimension; uses

number-theoretic methods of Green–Ruzsa. (Title, full author list, and content verified from the abstract; the earlier claim that “its only upper-bound tool invokes Katz–Tao sum–difference unmodified” was not supported and has been removed.)