

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

Tejasvi Singh Tomar

23 October 2026

## Abstract

In October 2026 Alman and Vassilevska Williams refuted the 3SUM and APSP hypotheses, and with them the Exact Triangle, 3XOR, weighted  $k$ -Clique and  $(\min, +)$ -convolution hypotheses. Their algorithm counts triangles through prescribed edges in lopsided tripartite graphs whose middle part has  $n^\varepsilon$  vertices for small  $\varepsilon$ , and Majdoddin extended its range to every  $\varepsilon < \ln 4/(5 \ln 5) \approx 0.17227$ . Many conditional lower bounds were derived from the refuted hypotheses. We classify the main ones. Most of them were proved through All-Edges Sparse Triangle or offline Set Disjointness on instances that the new algorithms do not reach, namely balanced graphs and lopsided graphs with a larger middle part. We state the hypotheses that these reductions use, among them a family  $\text{LSTH}_\delta$  for the natural lopsided instances with middle part  $n^\delta$ , a lopsided form for listing a bounded number of triangles, and degree-bounded, all-nodes and batched forms of the hypothesis for balanced graphs. Most set-disjointness, dynamic-graph and triangle bounds follow from these hypotheses. In contrast, a lower bound of Goldstein, Lewenstein and Porat for Set Disjointness over small universes is false, and the trade-offs of Kopelowitz, Pettie and Porat lose their support where they use middle parts below  $n^{0.17227}$ . The quadratic bounds for most 3SUM-hard problems in computational geometry, for jumbled indexing and for local alignment over large alphabets are unsupported. We also observe that 3SUM-Indexing now has a data structure with truly subquadratic preprocessing time and truly sublinear expected query time. Finally, we record why the closed-form analysis of the new method stops at  $\ln 4/(5 \ln 5)$ , and we use border-rank lower bounds to exclude several natural kinds of base identity that might give a larger range.

## 1 Introduction

Alman and Vassilevska Williams [AV26] gave the first polynomially faster algorithms for 3SUM and APSP. Their algorithms are deterministic and run in  $O(n^{1.9992})$  time for 3SUM on integers of polynomial magnitude, in  $O(n^{2.99942})$  time for APSP with polynomially bounded integer weights, and in  $O(n^{2.9983})$  time for Exact Triangle. At their core is an algorithm that computes prescribed entries of a thin matrix product, which is the counting version of All-Edges Sparse Triangle on lopsided graphs. Composed with known reductions, it refutes the Exact Triangle hypothesis, the Zero-, Min- and Max-Weight  $k$ -Clique hypotheses, the  $(\min, +)$ -convolution hypothesis and the Directed Unweighted APSP hypothesis. The same approach gives a truly subquadratic algorithm for 3XOR and Las Vegas algorithms for 3SUM and Exact Triangle over the reals, and the data-structure version refutes three hinted Online Matrix-Vector conjectures in the regime of thin hints. Majdoddin [Maj26] then extended the range of the matrix theorem from  $D \leq N^{0.1204}$  to every  $D \leq N^\varepsilon$  with  $\varepsilon < \ln 4/(5 \ln 5) = 0.17227\dots$ . He also gave reductions that keep its whole saving, with exponents  $1.9945 + o(1)$  for 3SUM and  $2.9945 + o(1)$  for Exact Triangle, the  $(\min, +)$ -product and APSP, and he extended the 3XOR algorithm to 3SUM over  $\mathbb{F}_p^d$  and over abelian groups of

exponent  $p^2$ . Saha and Ye [SY26] gave a simple deterministic  $\tilde{O}(n^{2.9964})$  algorithm for the  $(\min, +)$ -product, Karingula and Lovett [KL26] a simpler proof of a weaker form of the matrix theorem, and Fischer, Polak and Schmidt [FPS26] a reduction from real to integer 3SUM.

For two decades these hypotheses were the main sources of conditional lower bounds in fine-grained complexity, together with SETH and OMv. A lower bound derived from a refuted hypothesis is still a true implication, but its premise is false, so it gives no evidence for the bound. Some of these bounds are now false, many can be derived from hypotheses that no known algorithm contradicts, and some have lost their evidence entirely. In this paper we determine which of these cases applies to each of the main bounds.

**The lopsided regime.** The algorithms of [AV26, Maj26] handle All-Edges Sparse Triangle only when one of the three parts of the tripartite graph is polynomially smaller than the other two, which in the language of sets is offline Set Disjointness over a small universe. Alman and Vassilevska Williams note that the reductions of Kopelowitz, Pettie and Porat [KPP16] and of Vassilevska Williams and Xu [VX20, Corollary 3.12] produce instances of this kind [AV26, Section 3.1]. We parametrize these instances by the exponent  $\delta$  of the universe size. For every  $\delta \in (0, 2)$ , these reductions, and the deterministic one of Fischer, Kaliciak and Polak [FKP24], produce what we call *natural* instances, with universe  $n^\delta$ , sets of size  $n^{\delta/2}$  and  $n^{2-\delta/2}$  queries. The known algorithms reach these instances only for  $\delta < 0.17227$ . Much of the classification then reduces to bookkeeping, namely to finding the shape of the instances behind each downstream bound.

**Relation to [AV26].** Alman and Vassilevska Williams already mark the problems that are reached from 3SUM, APSP or Exact Triangle only by one-way reductions as having lost their evidence [AV26, Figure 1], among them the dynamic problems of [Pät10, AV14, KPP16] and the geometric problems of [GO95]. They also propose the hypothesis that the balanced case of All-Edges Sparse Triangle needs  $m^{4/3-o(1)}$  time [AV26, Section 6]. Our classification is finer. Most of the one-way reductions pass through sparse triangle problems on instances that the new algorithms do not reach. Their bounds therefore survive under hypotheses that nothing contradicts. The hypotheses needed, however, are often stronger than the balanced one, and a few bounds are false.

**Results.** We list the status of the standard hypotheses in Table 1. In Section 3 we state the hypotheses that the downstream reductions use, and in Tables 2 and 3 we give the verdicts for the main conditional lower bounds. Our main findings are as follows.

- Most set-disjointness, dynamic and triangle bounds follow from hypotheses that no known algorithm contradicts. Several of them, however, need more than the balanced hypothesis, namely a degree-bounded form for the distance-oracle bounds of [ABKZ22], a batched form for AE-Mono $\Delta$  [VX20], and a capped-listing form of the lopsided hypothesis for the dynamic bounds of [AV14].
- The lower bound of [GLP19, Theorem 7] for Set Disjointness over a universe of size  $u$ , which asks for  $m^{2-o(1)}$  preprocessing or  $u^{1/2-o(1)}$  query time, fails for  $u = m^r$  with  $0 < r < 0.17227$  (Proposition 4.1).
- The hard instances of [KPP16, Theorem 1.4] with universe below the threshold admit data structures that beat the trade-off they were used to prove (Proposition 4.2). Hence the corresponding corners of the Set Disjointness and dynamic trade-offs of [KPP15, KPP16] are unsupported.

- Several bounds proved by direct reductions from 3SUM are unsupported, namely those for most problems of computational geometry [GO95, Kin04], for jumbled indexing [ACLL14, GKLP16] and for local alignment over large alphabets [AVW14].
- 3SUM-Indexing has a randomized data structure with truly subquadratic preprocessing time and truly sublinear expected query time (Proposition 4.3). Before [AV26], such a data structure would have refuted the 3SUM hypothesis.

**The threshold.** The threshold  $\varepsilon^* = \ln 4 / (5 \ln 5)$  is the exponent of Coppersmith’s rectangular matrix multiplication algorithm from Schönhage’s partial-product tensor [Cop82], which [Maj26] uses. In Section 5 we record why the closed-form analysis of the method stops there, mostly following Majdoddin’s barriers [Maj26, Appendix A]. We also add Koszul-flattening bounds that exclude two natural profiles with two inner outputs. Majdoddin also reports numerically that the method reaches about 0.1784 when the prescribed entries are known in advance [Maj26, Section 8]. Our verdicts use the proved threshold. Those that rest on  $\text{LSTH}_\delta$  with  $0.17227 \leq \delta < 0.1784$  should be read with this in mind.

## 2 Hypotheses after the refutation

In Table 1 we list the standard hypotheses of fine-grained complexity. A hypothesis is *refuted* if a polynomially faster algorithm is known for its problem in the stated form, *refuted in a regime* if this holds for some but not all of its parameters, *open* if the new algorithms do not address it although its evidence came from the refuted hypotheses, and *unaffected* otherwise.

For orientation, the best exponents claimed in research papers are  $1.9945 + o(1)$  for 3SUM and  $2.9945 + o(1)$  for Exact Triangle, the  $(\min, +)$ -product and APSP, all deterministic [Maj26, Corollary 16.1]. The  $o(1)$  hides a factor  $\exp(O(\sqrt{\log n \log \log n}))$ , which makes the gain over [AV26] from the second part of [Maj26] appear only for astronomically large  $n$  [Maj26, Sections 13.7, 15.7]; without that part the bounds are  $O(n^{1.99863})$  for 3SUM,  $O(n^{2.9973})$  for Exact Triangle and  $\tilde{O}(n^{2.99909})$  for the  $(\min, +)$ -product and APSP [Maj26, Corollary 16.1]. Saha and Ye give  $\tilde{O}(n^{2.9964})$  for the  $(\min, +)$ -product and APSP [SY26], and an unrefereed repository [Shr26] claims  $O(n^{1.9931055})$  for 3SUM and  $O(n^{2.9926862626})$  for APSP.

## 3 Sparse triangle hypotheses

Most downstream bounds were derived from the refuted hypotheses through All-Edges Sparse Triangle or offline Set Disjointness, on instances of a particular shape. In this section we state the forms of these problems that the downstream reductions use, determine which of them the new algorithms reach, and name the rest. Throughout,  $\varepsilon^* := \ln 4 / (5 \ln 5) = 0.17227 \dots$

### 3.1 Lopsided instances

**Definition 3.1** (Lopsided All-Edges Sparse Triangle). An instance of  $\text{LST}(n, u, Q)$  is a tripartite graph with parts  $A$  and  $B$  of  $n$  vertices each and a middle part  $M$  of  $u$  vertices, arbitrary edges in  $A \times M$  and  $M \times B$ , and a set  $W \subseteq A \times B$  of  $Q$  query pairs. The task is to decide, for every  $(a, b) \in W$ , whether  $a$  and  $b$  have a common neighbour in  $M$ . In the counting version the task is to output the number of common neighbours.

| Hypothesis                                                                                  | Status              | Reference                                                                                                          |
|---------------------------------------------------------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------|
| 3SUM, integers of polynomial magnitude;<br>Strong 3SUM (integers in $[-n^2, n^2]$ )         | refuted             | [AV26, Maj26]                                                                                                      |
| 3SUM and Exact Triangle over the reals                                                      | refuted             | real RAM, Las Vegas [AV26, Theorem 35], through [CVX22]; for 3SUM also [FPS26]                                     |
| APSP and its subcubic equivalence class [VW18, AGV15]                                       | refuted             | [AV26, Maj26, SY26]                                                                                                |
| Exact Triangle; Strong Exact Triangle (weights in $[-n, n]$ ) [CX24]                        | refuted             | [AV26, Theorem 19]                                                                                                 |
| 3XOR; 3SUM over $\mathbb{F}_p^d$ and over abelian groups of exponent $p^2$                  | refuted             | [AV26, Remark 23], [Maj26, Theorem E, Section 19]                                                                  |
| Zero-, Min- and Max-Weight $k$ -Clique                                                      | refuted             | [AV26, Corollary 39]                                                                                               |
| (min, +)-convolution [KPS17, CMWW19]                                                        | refuted             | [AV26], through [BCD+14, BIS17, CMWW19]                                                                            |
| Directed Unweighted APSP                                                                    | refuted             | [AV26, Theorem 34]                                                                                                 |
| Hinted OMv [BNS19, Conjectures 5.2, 5.7, 5.12]                                              | refuted in a regime | thin hints: $\tau < 0.1204$ [AV26, Corollary 40], $\tau < \varepsilon^*$ [Maj26, Corollary 10.1, Proposition 10.3] |
| Lopsided Sparse Triangle LSTH $_\delta$                                                     | refuted in a regime | $\delta < \varepsilon^*$ (Theorem 3.4)                                                                             |
| SETH, Orthogonal Vectors                                                                    | unaffected          |                                                                                                                    |
| OMv without hints [HKNS15]                                                                  | unaffected          |                                                                                                                    |
| $k$ -SUM and $k$ -XOR for $k \geq 4$                                                        | unaffected          | [AV26, Section 1.2]                                                                                                |
| 3SUM-Indexing (no truly subquadratic space with $\tilde{O}(1)$ query time)                  | unaffected          | [AV26, Section 1.2]; the strong form fell in 2019 [KP19, GGHPV20]                                                  |
| Multiphase conjecture [Pät10]                                                               | open                | its 3SUM-based case follows from LST(0.4, 0.2, 1.4) (Section 4)                                                    |
| BST, its degree-bounded, all-nodes and batched forms, TLH, LST( $r, \sigma, \kappa$ ), LSTL | open                | Section 3                                                                                                          |
| Combinatorial BMM                                                                           | unaffected          |                                                                                                                    |
| Combinatorial versions of 3SUM and APSP                                                     | open                | the new algorithms are algebraic [AV26, Section 1.2]                                                               |

Table 1: Status of fine-grained hypotheses. Integer inputs have polynomial magnitude and the integer algorithms are deterministic.  $\varepsilon^* = \ln 4 / (5 \ln 5) = 0.17227 \dots$

In the language of sets,  $\text{LST}(n, u, Q)$  is offline Set Disjointness with  $n$  sets on each side over a universe of size  $u$  and  $Q$  queries, and the counting version asks for the sizes of the intersections. Answering each query by merging the two neighbourhoods takes  $O(Q\Delta)$  time, where  $\Delta$  is the largest degree into  $M$ . Computing the whole Boolean product of the two biadjacency matrices takes  $n^{\omega(1, \log_n u, 1) + o(1)}$  time, which is  $n^{2+o(1)}$  when  $u \leq n^{0.321}$  [VXXZ24].

**Definition 3.2** (Natural instances). For  $0 < \delta < 2$ , a *natural* instance of  $\text{LST}_\delta(n)$  is an instance of  $\text{LST}(n, n^\delta, Q)$  in which every vertex of  $A \cup B$  has at most  $n^{\delta/2}$  neighbours in  $M$  and  $Q \leq n^{2-\delta/2}$ .

Merging solves natural instances in  $O(n^2)$  time for every  $\delta$ .

**Hypothesis 3.3** ( $\text{LSTH}_\delta$ ). For every  $\varepsilon > 0$ , natural instances of  $\text{LST}_\delta(n)$  cannot be solved in  $O(n^{2-\varepsilon})$  time. The counting form  $\#\text{LSTH}_\delta$  is the same statement for the counting version.

The reductions of [KPP16, Theorem 1.4] from 3SUM, of [FKP24, Theorem 1.4] from deterministic 3SUM and of [VX20, Corollary 3.12] from Exact Triangle produce, for a parameter  $0 < \gamma < 1$ , offline Set Disjointness with  $\tilde{O}(N)$  sets of size  $O(N^{1-\gamma})$  over a universe of size  $O(N^{2-2\gamma})$  and  $\tilde{O}(N^{1+\gamma})$  queries, where  $N$  is the number of sets up to logarithmic factors. After padding the instance with isolated vertices to absorb the logarithmic factors, these are natural instances of  $\text{LST}_{2-2\gamma}$ . As noted in the introduction, this is the regime pointed out in [AV26, Section 3.1]. Consequently 3SUM and Exact Triangle each imply  $\text{LSTH}_\delta$  for every  $0 < \delta < 2$ , under randomized reductions (deterministic ones for 3SUM, by [FKP24]). The refutation, however, reaches only small  $\delta$ .

**Theorem 3.4** ([AV26, Maj26]). *For every  $\delta < \varepsilon^*$  and every  $\mu > 0$  there is  $\gamma > 0$  such that the counting version of  $\text{LST}(n, n^\delta, n^{2-\mu\delta})$  can be solved deterministically in  $O(n^{2-\gamma\delta})$  time. In particular  $\text{LSTH}_\delta$  and  $\#\text{LSTH}_\delta$  are false for every  $\delta < \varepsilon^*$ .*

*Proof.* Majdoddin’s Theorem 7.4, which is part of his Theorem A [Maj26], computes the entries of a product  $XY$  with  $X \in \mathbb{Z}^{N \times D}$  at any  $N^2/D^\mu$  prescribed positions in  $O(N^2/D^\gamma)$  time when  $D \leq N^\varepsilon$ ,  $\varepsilon < \varepsilon^*$ , and for  $\varepsilon < \ln 4/(5 \ln 10) = 0.1204$  this is [AV26, Theorem 25]. We apply it to the two biadjacency matrices, with  $N = n$ ,  $D = n^\delta$  and the query pairs as positions.  $\square$

Theorem 3.4 covers only the proved range. Majdoddin’s numerical range of about 0.1784, mentioned in the introduction, assumes that  $W$  is known in advance and that the leaves of the recursion are counted exactly [Maj26, Section 8]. If this numerical range is confirmed,  $\text{LSTH}_\delta$  fails for  $\delta < 0.1784$  as well. For larger  $\delta$  no algorithm faster than  $O(n^2)$  is known for natural instances.

Some conditional lower bounds use lopsided instances that do not saturate these bounds, with smaller sets or fewer queries. Merging solves them in less than  $n^2$  time, so  $\text{LSTH}_\delta$  says nothing about them. The following general form covers them.

**Hypothesis 3.5** ( $\text{LST}(r, \sigma, \kappa)$ ). Let  $0 < \sigma \leq r < 2$  with  $\min(r, 2\sigma) \geq 1/5$ , and let  $\kappa + \sigma \leq 2$ . For every  $\varepsilon > 0$ , instances of  $\text{LST}(n, n^r, n^\kappa)$  in which every vertex of  $A \cup B$  has at most  $n^\sigma$  neighbours in  $M$  cannot be solved in  $O(n^{\kappa+\sigma-\varepsilon})$  time.

Natural instances are the case  $\sigma = r/2$ ,  $\kappa = 2 - r/2$ . The condition  $\kappa + \sigma \leq 2$  keeps the merging time below the cost  $n^{2+o(1)}$  of the whole product for small  $r$ . Since every set has at most  $n^\sigma$  elements, hashing the universe into  $O(n^{2\sigma})$  buckets, with  $O(\log n)$  independent hash functions, preserves the answers with high probability [Pät10, GLP19]. Hence only  $\min(r, 2\sigma)$  matters. The restriction  $\min(r, 2\sigma) \geq 1/5$  keeps the hypothesis clear of both the proved and the numerical range, and every use below satisfies it.

## 3.2 Listing

The dynamic bounds of [AV14] and the triangle-listing bound of [Pät10, Theorem 5] do not follow from the decision version of LST. They use a capped listing problem, in which the task is to list all triangles that contain a query pair or to report that there are more than a given number of them [AV14, Theorem 10.1]. A fast algorithm for the decision version does not obviously give a fast algorithm for capped listing, so the corresponding hypothesis has to be stated separately.

**Hypothesis 3.6** ( $\text{LSTL}(r, \sigma, \kappa, \lambda)$ ). Let  $r, \sigma, \kappa$  be as in  $\text{LST}(r, \sigma, \kappa)$  and  $0 \leq \lambda \leq \kappa + \sigma$ . For every  $\varepsilon > 0$ , no algorithm, given an instance of  $\text{LST}(n, n^r, n^\kappa)$  in which every vertex of  $A \cup B$  has at most  $n^\sigma$  neighbours in  $M$ , lists in  $O(n^{\kappa+\sigma-\varepsilon})$  time all triangles that contain a query pair whenever there are at most  $n^\lambda$  of them.

For  $\lambda < \kappa$ ,  $\text{LSTL}(r, \sigma, \kappa, \lambda)$  implies  $\text{LST}(r, \sigma, \kappa)$ , because a decision algorithm finds the query pairs that lie in triangles, and if there are at most  $n^\lambda$  of them, merging their neighbourhoods lists the triangles in  $O(n^{\lambda+\sigma})$  further time. For balanced graphs the listing hypothesis is the following, in the general form that matches the listing algorithm of Björklund, Pagh, Vassilevska Williams and Zwick [BPWZ14] when  $\omega = 2$ .

**Hypothesis 3.7** (TLH). For every  $\varepsilon > 0$  and  $m \leq t \leq m^{3/2}$ , listing  $t$  triangles of a graph with  $m$  edges, or all of them if there are fewer, cannot be done in  $O(m^{1-\varepsilon}t^{1/3})$  time.

At  $t = m$  the bound is  $m^{4/3}$ . The algorithms of [AV26, Maj26] count triangles without listing them, and they apply only to lopsided graphs whose middle part has fewer than  $n^{\varepsilon^*}$  vertices. They touch none of the listing hypotheses.

### 3.3 Balanced instances

At  $\delta = 1$  the three parts of a natural instance have the same size, and natural instances are All-Edges Sparse Triangle on graphs with  $n$  vertices of degree about  $\sqrt{n}$  and  $m = \Theta(n^{3/2})$  edges, where merging takes  $n^2 = m^{4/3}$  time. As recalled in the introduction, Alman and Vassilevska Williams propose that the balanced case needs  $m^{4/3-o(1)}$  time [AV26, Section 6]. We use this hypothesis in the form for all graphs with  $m$  edges.

**Hypothesis 3.8** (BST). For every  $\varepsilon > 0$ , All-Edges Sparse Triangle on graphs with  $m$  edges cannot be solved in  $O(m^{4/3-\varepsilon})$  time.

Several downstream reductions need more than BST, because they start from graphs of bounded degree, solve the all-nodes version, or solve many independent instances at once. We name the corresponding forms as follows, with  $\varepsilon > 0$  arbitrary in each.

- *Degree-bounded form*  $\text{BST}_\beta$ , for  $0 < \beta \leq 1/2$ : All-Edges Sparse Triangle on  $n$ -vertex graphs of maximum degree at most  $n^\beta$  cannot be solved in  $O(n^{1+2\beta-\varepsilon})$  time. This is the conclusion of [VX20, Corollary 3.11], and  $\text{BST}_{1/2}$  is the base of [ABKZ22, Theorem 4.1]. Since such graphs have  $m \leq n^{3/2}$  edges,  $\text{BST}_{1/2}$  implies BST; the converse is not known [ABKZ22, Section 7].
- *All-nodes form*: All-Nodes Sparse Triangle (decide for every vertex whether it lies in a triangle) on graphs with  $m$  edges cannot be solved in  $O(m^{5/4-\varepsilon})$  time. This is the statement of [CX24, Theorem 4.6].
- *Batched form*:  $m^{2/3}$  independent instances of All-Edges Sparse Triangle with  $m$  edges each cannot be solved in  $O(m^{2-\varepsilon})$  total time. This is what the proof of [VX20, Corollary 3.10] uses.

### 3.4 Relations between different $\delta$

Hardness of natural instances is not obviously monotone in  $\delta$ , and padding transfers it only with a loss.

**Lemma 3.9** (padding). *Let  $0 < \delta < \delta' < 2$  and  $n' := \lceil n^{(2-\delta/2)/(2-\delta'/2)} \rceil$ . Every natural instance of  $\text{LST}_\delta(n)$  is a natural instance of  $\text{LST}_{\delta'}(n')$  after adding isolated vertices to all three parts. Consequently, if natural instances of  $\text{LST}_{\delta'}$  can be solved in  $O(n^{2-\varepsilon})$  time for some  $\varepsilon > (\delta' - \delta)/(2 - \delta/2)$ , then  $\text{LSTH}_\delta$  is false.*

*Proof.* After padding, the outer parts have  $n'$  vertices and the middle part  $n'^{\delta'} \geq n^\delta$  vertices. Degrees are at most  $n^{\delta/2} \leq n'^{\delta'/2}$ , and  $Q \leq n^{2-\delta/2} \leq n'^{2-\delta'/2}$  by the choice of  $n'$ . Isolated vertices need not be written down, so an algorithm for  $\text{LST}_{\delta'}$  running in time  $O(n'^{2-\varepsilon})$  solves the original instance in time  $O(n^{(2-\varepsilon)(2-\delta/2)/(2-\delta'/2)} + n^{2-\delta/2} + n^{1+\delta/2})$ , where the last two terms read the input. The first exponent is below 2 exactly when  $\varepsilon(2-\delta/2) > \delta' - \delta$ .  $\square$

Read in the other direction, Lemma 3.9 says that  $\text{LSTH}_\delta$  implies the lower bound

$$n^{2 - \frac{\delta' - \delta}{2 - \delta/2} - o(1)}$$

for natural instances of  $\text{LST}_{\delta'}$ , which approaches  $n^2$  as  $\delta' \downarrow \delta$ .

## 4 Which lower bounds survive

In this section we classify the main conditional lower bounds that were derived from 3SUM, APSP, Exact Triangle or weighted  $k$ -Clique. For each of them we record the shape of its hard instances and one of the following verdicts.

**Faster.** The problem is equivalent to a refuted one and now has a polynomially faster algorithm.

**Refuted.** The stated lower bound is false, because a known algorithm beats it. (The original theorem is an implication from a hypothesis that is now false, so it remains true, and what fails is its conclusion.)

**Implied by  $H$ .** The bound follows from a hypothesis  $H$  that no known algorithm contradicts, by the original derivation started from  $H$  instead of the refuted hypothesis.

**Unaffected.** The bound rests on a hypothesis that was never derived from the refuted ones and is not refuted.

**Unsupported.** The base hypothesis is refuted and no surviving hypothesis is known to imply the bound. The problem need not have become easier, but the bound has lost its evidence.

Tables 2 and 3 summarize the classification, and in the rest of the section we give the reasons where they are not immediate. As explained in Section 3, the reductions of [KPP16, Theorem 1.4], [FKP24, Theorem 1.4] and [VX20, Corollary 3.12] produce natural instances of  $\text{LST}_{2-2\gamma}$ , so the value of  $\gamma$  that a downstream bound uses decides its fate. The bound is refuted or unsupported when  $2-2\gamma < \varepsilon^*$ , that is, when  $\gamma > 1-\varepsilon^*/2 = 0.91386\dots$ , and it is implied by  $\text{LSTH}_{2-2\gamma}$  otherwise. Where the references list a full version of a paper, its theorem numbers refer to that version.

### 4.1 Set problems and data structures

We begin with a published lower bound that the data structure of [AV26, Maj26] falsifies as stated. Goldstein, Lewenstein and Porat [GLP19, Theorem 7] show from 3SUM that every data structure for Set Disjointness on  $m$  sets  $S_1, \dots, S_m \subseteq [u]$  of total size  $N$ , for every  $u \in [N^\delta, N]$  with  $\delta > 0$ , needs  $\Omega(m^{2-o(1)})$  preprocessing time or  $\Omega(u^{1/2-o(1)})$  query time.

**Proposition 4.1.** *For every  $0 < r < \varepsilon^*$  and every  $q > 0$  there is  $c > 0$  such that Set Disjointness on  $m$  sets over a universe of size  $u = m^r$  has a deterministic data structure with preprocessing time  $O(m^{2-c})$  that returns the size of any intersection in  $O(u^q)$  time. In particular the conclusion of [GLP19, Theorem 7] fails for  $u = m^r$ .*

| Lower bound                                                                                                                                                | Instances                                                | Verdict                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Static set problems and data structures</i>                                                                                                             |                                                          |                                                                                                                                                            |
| Offline Set Disjointness [KPP16, Thm 1.4], [FKP24, Thm 1.4], [VX20, Cor 3.12], parameter $\gamma$                                                          | $LST_{2-2\gamma}$                                        | refuted for $\gamma > 1 - \varepsilon^*/2$ ; implied by $LSTH_{2-2\gamma}$ otherwise                                                                       |
| Set Disjointness over a universe $[u]$ , $m$ sets: $m^{2-o(1)}$ preprocessing or $u^{1/2-o(1)}$ query [GLP19, Thm 7]                                       | $LST_r$ , $u = m^r$                                      | refuted for $r < \varepsilon^*$ (Proposition 4.1); implied by $LSTH$ near $r$ for $r > \varepsilon^*$                                                      |
| Set Disjointness trade-off at $\gamma$ [KPP16, Thm 1.7]; $p + 2q \geq 2$ [KPP16, Cor 1.8]                                                                  | $LST_{2-2\gamma}$ , $\gamma \approx 2 - 2/p$             | unsupported for $\gamma > 1 - \varepsilon^*/2$ , i.e. $p > 4/(2 + \varepsilon^*) = 1.8413\dots$ (Proposition 4.2); implied by $LSTH_{2-2\gamma}$ otherwise |
| Range mode, stretch- $(< 2)$ distance oracles, 3SUM-Indexing over small universes [GLP19, Cor 9, Thm 12, 14]                                               | via [GLP19, Thm 7]                                       | unsupported where they use $r < \varepsilon^*$ ; implied by $LSTH$ near $r$ for $r > \varepsilon^*$                                                        |
| Set Intersection with reporting, $p + q \geq 2$ [KPP16, Cor 5.2]                                                                                           | [KPP16, Thm 1.5]                                         | unsupported                                                                                                                                                |
| Consequences of Strong 3SUM-Indexing [GKLP17, Thm 1, 2, 7, 11, 12]                                                                                         |                                                          | unsupported since 2019 [KP19, GGHPV20]                                                                                                                     |
| Set Disjointness, universe $\log^c m$ , constant query time [GKLP17, Conj 1]                                                                               | polylog universe                                         | unaffected                                                                                                                                                 |
| Multiphase problem from 3SUM [Pät10, Thm 3]                                                                                                                | $(r, \sigma, \kappa) = (0.4, 0.2, 1.4)$                  | implied by $LST(0.4, 0.2, 1.4)$                                                                                                                            |
| <i>Dynamic problems</i>                                                                                                                                    |                                                          |                                                                                                                                                            |
| Multiphase conjecture $\Rightarrow$ dynamic bounds [Pät10, Thm 2]                                                                                          |                                                          | implied by the multiphase conjecture                                                                                                                       |
| st-SubConn, st-Reach, BPMatch, SC, $\emptyset$ -PP from 3SUM [AV14, Thm 5.2]; st-reachability from Exact Triangle [VX20, Cor 1.6]; $\alpha \in [1/6, 1/3]$ | capped listing, $r = \frac{2}{3(1-\alpha)} \in [0.8, 1]$ | implied by $LSTL(r, \sigma, \kappa, \lambda)$ (Section 4.2)                                                                                                |
| Decremental/incremental st-SP, BWMatch from APSP [AV14, Thm 5.4]                                                                                           | one-way APSP-hard                                        | unsupported                                                                                                                                                |
| Fully dynamic MCM [KPP16, Thm 1.11]; incremental set intersection emptiness [KPP15, Thm 9]                                                                 | $LST_{2-2\gamma}$                                        | constant-update (resp. constant-query) corner unsupported; near-linear-preprocessing corner implied by $LSTH_\delta$ , $\delta \rightarrow 2$              |
| Incremental MCM, $\Omega(\hat{n}^{0.3903})$ [KPP16, Thm 1.12]                                                                                              | $LST_{1.1231}$                                           | implied by $LSTH_{1.1231}$                                                                                                                                 |
| $d$ -failure connectivity, deletion time $d^a$ [KPP16, Thm 5.3]                                                                                            | $LST_{1/a}$                                              | implied by $LSTH_{1/a}$ for $a \leq 1/\varepsilon^* = 5.80\dots$ ; unsupported for $a > 1/\varepsilon^*$                                                   |
| Dynamic matrix inverse and relatives [BNS19], partially dynamic distance oracles [HLS24], dynamic attention [BSZ24]                                        | hinted OMv                                               | thin-hint end unsupported; main bounds implied by the hinted OMv conjectures where they are not refuted [AV26, Section 5.4]                                |

Table 2: Set problems and dynamic problems after the refutation.  $\varepsilon^* = \ln 4 / (5 \ln 5)$ ;  $LST_\delta$ : natural lopsided instances;  $LST$  and  $LSTL$ : general-parameter and capped-listing forms (Section 3).

| Lower bound                                                                                                                | Instances                                     | Verdict                                                                              |
|----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|--------------------------------------------------------------------------------------|
| <i>Triangles, cycles and distance oracles</i>                                                                              |                                               |                                                                                      |
| All-Edges Sparse Triangle, $m^{4/3}$ [VX20, Cor 1.1], [CX24, Thm 3.1, 4.3, A.3]                                            | $m$ -edge graphs                              | this is BST                                                                          |
| Reporting $m$ triangles, $m^{4/3}$ [Pät10, Thm 5]; enumerating $t$ triangles [KPP16, Thm 1.10]                             | balanced; dense, $t \leq m^{3/2}$             | implied by TLH                                                                       |
| AE-Sparse Triangle, maximum degree $n^\beta$ [VX20, Cor 3.11]                                                              | equal parts                                   | this is $\text{BST}_\beta$                                                           |
| AE-Mono $\Delta$ , $n^{5/2}$ [VX20, Cor 3.10]                                                                              | $\tilde{O}(n)$ independent instances          | implied by the batched form of BST                                                   |
| Distance oracles, dynamic APSP, cycle enumeration [ABKZ22, Cor 1.7–1.9]                                                    | maximum degree $\sqrt{n}$                     | implied by $\text{BST}_{1/2}$ ; for odd $k$ in Cor 1.9 also TLH                      |
| 4-cycle and related bounds [ABKZ22, Cor 1.11]                                                                              | triangle detection, maximum degree $\sqrt{n}$ | unaffected (Section 4.3)                                                             |
| Graphs with few short cycles [ABF23, Thm 1.1–1.5], [JX23, Thm 1.5, 1.7]                                                    | from 3SUM                                     | sharper exponents unsupported; weaker forms from $\text{BST}_{1/2}$ through [ABKZ22] |
| Triangle detection, maximum degree $n^{1/4}$ , $n^{1.5}$ [JX23, Thm 1.14]                                                  | Strong 3SUM                                   | unsupported                                                                          |
| All-Nodes Sparse Triangle, $m^{5/4}$ [CX24, Thm 4.6]                                                                       | $m = N^{4/3}$                                 | this is the all-nodes form                                                           |
| Sparse Triangle Detection, $m^{9/7}$ [CX24, Thm 4.7]                                                                       | Strong Exact Triangle                         | unsupported                                                                          |
| <i>Strings, geometry, convolutions</i>                                                                                     |                                               |                                                                                      |
| Jumbled (histogram) indexing [ACLL14, Thm 2, 3, 5], [GKLP16, Thm 6, 7]                                                     | (strong) Convolution-3SUM, 3SUM               | unsupported                                                                          |
| Local alignment, alphabet $n^{1-\delta}$ [AVW14, Thm 1]; alphabet $\sqrt{n}$ [AVW14, Cor 1]                                | 3SUM; Max-Weight 4-Clique                     | unsupported                                                                          |
| Local alignment, binary alphabet [AVW14, Thm 2]                                                                            | SETH                                          | unaffected                                                                           |
| GeomBase, 3SUM' [GO95]                                                                                                     | equivalent to 3SUM                            | faster                                                                               |
| Three collinear points, point covering, motion planning [GO95]; minimum-area triangle [Kin04]; polygon containment [BHP01] | one-way from 3SUM                             | unsupported; for three collinear points a restricted-model bound remains [ES95]      |
| Problems as hard as sorting $X + Y$ , $o(n^2 \log n)$ level [Bar96]                                                        |                                               | unaffected                                                                           |
| (min, +)-convolution class [CMWW19, Thm 2]                                                                                 | reduces to 3SUM, APSP                         | faster                                                                               |
| MonoConvolution [LPV20, Thm 7, 8]                                                                                          | equivalent to 3SUM                            | faster                                                                               |
| <i>Other bases</i>                                                                                                         |                                               |                                                                                      |
| APSP-equivalent problems [VW18, AGV15]                                                                                     |                                               | faster [AV26]                                                                        |
| One-way APSP-hard problems, e.g. [AV14, Thm 5.4], local alignment at $n^{1.5}$ [AVW14]                                     |                                               | unsupported                                                                          |
| Average-case Zero- $k$ -Clique [LLV19, AHY25]                                                                              |                                               | refuted in every form with polynomially bounded weights [AV26, Cor 39]               |

Table 3: Triangle, string, geometric and other bounds after the refutation.  $\text{BST}_\beta$ , the all-nodes and batched forms, and TLH are defined in Section 3.

*Proof.* Apply the data-structure statement of [Maj26, Theorem 7.4] to the  $m \times u$  incidence matrix  $X$  and to  $Y = X^\top$ , with  $N = m$  and  $D = u = N^r$ . After preprocessing in  $O(m^2/u^\gamma) = O(m^{2-r\gamma})$  time, any entry of  $XY$ , which is the size of an intersection, is computed in  $O(u^q)$  time. Since  $N \leq mu$ , the universe size  $u = m^r$  lies in  $[N^\delta, N]$  for  $\delta = r/(1+r)$ . Hence  $q < 1/2$  contradicts [GLP19, Theorem 7].  $\square$

The hard instances of [GLP19, Theorem 7] have  $m$  sets of size  $O(\sqrt{u})$  and  $\tilde{O}(m^2/\sqrt{u})$  queries, up to factors  $u^{O(\varepsilon)}$  for an arbitrarily small  $\varepsilon > 0$ . Writing  $u = m^r$ , they are therefore natural instances of  $\text{LST}_{r'}$  for  $r'$  close to  $r$ . For  $r > \varepsilon^*$  the bound follows from these hypotheses, since a data structure with preprocessing  $m^{2-\varepsilon}$  and query time  $u^{1/2-\varepsilon}$  would solve them in  $O(m^{2-\varepsilon} + m^{2-r\varepsilon})$  time. The bounds that [GLP19] derive from its Theorem 7 by further reductions are unsupported where they use it with  $r < \varepsilon^*$  and implied by  $\text{LSTH}_r$  otherwise.

The trade-offs of Kopelowitz, Pettie and Porat for Set Disjointness are not refuted, because they concern data structures for all instances, but the instances used to prove them lose their force for  $\gamma > 1 - \varepsilon^*/2$ .

**Proposition 4.2.** *For every  $\gamma > 1 - \varepsilon^*/2$  there is  $c > 0$  such that the instances of [KPP16, Theorem 1.4] with parameter  $\gamma$  admit a deterministic data structure with*

$$t_p + N^{\frac{1+\gamma}{2-\gamma}} t_q = O(N^{\frac{2}{2-\gamma}-c}),$$

where  $N$  is the total size of the sets,  $t_p$  the preprocessing time and  $t_q$  the query time. The trade-off of [KPP16, Theorem 1.7] asserts that no data structure for Set Disjointness achieves this.

*Proof.* The instances have  $\Theta(n \log n)$  sets of size  $O(n^{1-\gamma})$  over a universe of size  $D = \Theta(n^{2-2\gamma})$ , so  $N = \Theta(n^{2-\gamma} \log n)$ ,  $N^{(1+\gamma)/(2-\gamma)} = \tilde{\Theta}(n^{1+\gamma})$  and  $N^{2/(2-\gamma)} = \tilde{\Theta}(n^2)$ . Moreover,  $D \leq (n \log n)^\varepsilon$  for some  $\varepsilon < \varepsilon^*$  because  $2 - 2\gamma < \varepsilon^*$ . Fix  $q' < 1/2$ . By the data-structure statement of [Maj26, Theorem 7.4] there is  $\gamma' > 0$  with  $t_p = O(n^2/D^{\gamma'})$  and  $t_q = O(D^{q'})$ , so  $t_p + n^{1+\gamma} t_q = \tilde{O}(n^{2-(2-2\gamma)\gamma'} + n^{1+\gamma+(2-2\gamma)q'})$ , and both exponents are below 2 because  $q' < 1/2$ .  $\square$

Corollary 1.8 of [KPP16] ( $p + 2q \geq 2$  for preprocessing  $N^p$  and query time  $N^q$ ) is proved by applying Theorem 1.7 with  $\gamma = 2 - 2/p + \varepsilon'$  for a small  $\varepsilon' > 0$ . For  $p > 4/(2 + \varepsilon^*) = 1.8413\dots$  this  $\gamma$  exceeds  $1 - \varepsilon^*/2$ , and Proposition 4.2 shows that the instances no longer support the bound. The corollary may still be true, but it has lost its evidence there. For  $p < 4/(2 + \varepsilon^*)$  it follows from  $\text{LSTH}_{4/p-2-2\varepsilon'}$ . The same applies to the constant-update corner of the fully dynamic trade-off [KPP16, Theorem 1.11] and the constant-query corner of [KPP15, Theorem 9], which use  $\gamma \rightarrow 1$ . At the other end, with near-linear preprocessing, these bounds use  $\gamma \rightarrow 0$  and follow from  $\text{LSTH}_\delta$  with  $\delta$  close to 2.

The reporting bound [KPP16, Corollary 5.2] uses the instances of [KPP16, Theorem 1.5], which must enumerate up to  $n^{2-\delta}$  elements of the intersections with  $\delta \rightarrow 0$ , about as many as the merging algorithm reads. It is thus a bound on the overhead of reporting. None of the hypotheses of Section 3 covers it, and the counting algorithms of [AV26, Maj26] do not report elements.

The multiphase problem of Pătraşcu [Păt10] gives one set in its second phase, so the data structure of [AV26], which preprocesses both families of sets, does not apply to it. No known algorithm contradicts the multiphase conjecture. Its evidence, however, came from 3SUM. The derivation [Păt10, Theorem 3] queries the intersections of  $k = R\sqrt{n}$  sets of size  $n/R$  over a universe of size  $\Theta((n/R)^2)$  with one set per vertex of the other side, for  $nR$  query pairs, with  $R = n^{0.75}$ . In the notation of  $\text{LST}(r, \sigma, \kappa)$ , with  $n^{1.25}$  vertices on each outer side, these instances have  $(r, \sigma, \kappa) = (0.4, 0.2, 1.4)$ . A multiphase data structure that beats [Păt10, Theorem 3] would answer all these

queries in  $n^{2-\Omega(1)}$  time, so  $\text{LST}(0.4, 0.2, 1.4)$  implies the theorem. The triangles found are listed only afterwards, and their number matters only for that step.

Finally, the new algorithms change what is possible for 3SUM-Indexing, where the arrays  $A$  and  $B$  are preprocessed and a query asks whether  $a + b = z$  for some  $a \in A$  and  $b \in B$ . A data structure with truly subquadratic preprocessing time and truly sublinear query time would have refuted the 3SUM hypothesis [DG25], and such a data structure now exists.

**Proposition 4.3.** *3SUM-Indexing on  $n$  integers of polynomial magnitude has a randomized data structure with preprocessing time  $O(n^{2-c})$  and expected query time  $O(n^{1-c})$  for some  $c > 0$ .*

*Proof.* Goldstein, Lewenstein and Porat [GLP19, Lemma 6] reduce a 3SUM-Indexing instance, for any  $X \in [n^\delta, n]$  and  $\varepsilon > 0$ , in truly subquadratic time to  $O(\varepsilon \log X)$  instances of Set Disjointness, the  $i$ -th with  $m_i = n\sqrt{X/u_i}$  sets of size  $O(\sqrt{u_i})$  over a universe of size  $u_i = X^{1+\varepsilon}/2^{i-1} \in [X^{1-\varepsilon}, X^{1+\varepsilon}]$ , so that each 3SUM-Indexing query is answered by  $O(n/\sqrt{X})$  queries to each instance and truly sublinear expected additional time. Take  $X = n^{1/10}$  and  $\varepsilon$  small. Then  $u_i \leq m_i^{0.11}$ , and the data structure of [Maj26, Theorem 7.4], used as in Proposition 4.1 with  $D = u_i \leq m_i^{0.11}$ , preprocesses the  $i$ -th instance in  $O(m_i^2/u_i^\gamma) \leq O(n^2 X^{1-(1-\varepsilon)(1+\gamma)})$  time and answers its queries in  $O(u_i^q) \leq O(X^{q(1+\varepsilon)})$  time. For  $\varepsilon < \gamma/3$  and  $q < 1/4$  the preprocessing is  $O(n^{2-c})$  and a 3SUM-Indexing query takes  $O(nX^{q(1+\varepsilon)-1/2} \log n) = O(n^{1-c})$  time.  $\square$

The 3SUM-Indexing conjecture itself, which concerns truly subquadratic space with  $\tilde{O}(1)$  query time, is not affected [AV26, Section 1.2].

## 4.2 Dynamic problems

The 3SUM-based bounds of Abboud and Vassilevska Williams [AV14, Theorem 5.2] use the triangle instances of [Pät10, Lemma 11] and need their capped listing version [AV14, Theorem 10.1], which lists all triangles or reports that there are more than  $\Delta = O(n^2/R)$ . The graphs have two parts of  $R\sqrt{n}$  vertices, a third part of  $n$  vertices, degrees  $O(n/R)$  into the third part and  $O(nR)$  query edges, with  $R \approx n^{1-3\alpha/2}$  for the trade-off parameter  $\alpha \in [1/6, 1/3]$ . With  $N := R\sqrt{n} = n^{3(1-\alpha)/2}$  vertices on each outer side, they are instances of  $\text{LST}(N, N^r, N^\kappa)$  with degrees at most  $N^\sigma$  and at most  $N^\lambda$  triangles, where

$$r = \frac{2}{3(1-\alpha)}, \quad \sigma = \frac{\alpha}{1-\alpha}, \quad \kappa = \frac{4-3\alpha}{3(1-\alpha)}, \quad \lambda = \frac{2+3\alpha}{3(1-\alpha)},$$

so that  $r \in [0.8, 1]$  and  $N^{\kappa+\sigma} = N^{2r} = n^2$ . The parameters saturate the bounds of natural instances only at  $\alpha = 1/3$ , and  $\lambda < \kappa$  for  $\alpha < 1/3$ . These bounds therefore follow from  $\text{LSTL}(r, \sigma, \kappa, \lambda)$ , far from the range of the known algorithms. Pagh's problem is not derived from these instances but through the multiphase route [Pät10]. The trade-off for dynamic st-reachability stated from Exact Triangle in [VX20, Corollary 1.6] is the same as [AV14, Theorem 5.2] and has the same verdict. The APSP-based bounds [AV14, Theorem 5.4] reduce from Min-Weight Triangle with polynomial weights, which is equivalent to APSP and now faster, to the dynamic problems. Since this reduction goes only one way, these bounds are unsupported.

The bound for incremental maximum cardinality matching [KPP16, Theorem 1.12] takes  $\gamma = (1+\alpha)/(2+3\alpha) = 0.43845\dots$  with  $\alpha = (\sqrt{17}-1)/8$ , hence  $\delta = 2-2\gamma = 1.1231\dots$ . The bound for  $d$ -failure connectivity [KPP16, Theorem 5.3] uses  $1/2 \leq \gamma < 1$  and forces deletion time  $d^{1/(2-2\gamma)}$  unless preprocessing or queries are slow. A lower bound of  $d^a$  on the deletion time therefore needs  $\delta = 1/a$ , and it is supported by  $\text{LSTH}_{1/a}$  exactly when  $1/a \geq \varepsilon^*$ . The data structure of Duan

and Pettie discussed in [KPP16, Section 5] has deletion time  $d^{2c+4}$ , which for  $c = 1$  lies in the unsupported range  $a > 1/\varepsilon^*$ .

The hinted OMv conjectures of van den Brand, Nanongkai and Saranurak [BNS19] are refuted for thin hints, and the lower bounds that used that end of their trade-offs (for dynamic matrix inverse with fast queries, partially dynamic distance oracles [HLS24] and dynamic attention [BSZ24]) have lost their support [AV26, Section 5.4], although no faster algorithm for these dynamic problems follows. The main bounds of [BNS19], at hint dimension  $n^{0.53}$ , use the conjectures far from the refuted range.

### 4.3 Triangles, cycles and distance oracles

Most triangle bounds derived from Exact Triangle, 3SUM or APSP use graphs whose parts have comparable sizes. The precise hypothesis they need varies, however, and several need more than BST. The All-Edges Sparse Triangle bounds of [VX20, Corollary 1.1] and [CX24, Theorems 3.1, 4.3, A.3] are statements about all graphs with  $m$  edges, which is BST itself. Corollary 3.11 of [VX20] is the degree-bounded form. The bounds of Abboud, Bringmann, Khourey and Zamir for distance oracles, dynamic APSP and cycle enumeration [ABKZ22, Corollaries 1.7–1.9] start from all-edges triangle detection in  $n$ -vertex graphs of maximum degree  $\sqrt{n}$  [ABKZ22, Theorem 4.1], which is  $\text{BST}_{1/2}$ . The authors of [ABKZ22] do not know how to reduce the general case to it. The bound  $n^{5/2}$  for AE-Mono $\Delta$  [VX20, Corollary 3.10] packs  $m^{1/3}$  independent sparse instances into one instance of AE-Mono $\Delta$ . It therefore needs the batched form, since a single balanced instance embedded in AE-Mono $\Delta$  gives only an  $n^2$  bound. The all-nodes bound [CX24, Theorem 4.6] uses graphs with  $N$  vertices per part,  $m = N^{4/3}$  edges and maximum degree  $N^{1/3}$ .

The refinements of Abboud, Bringmann and Fischer and of Jin and Xu for graphs with few short cycles [ABF23, JX23] assume 3SUM, so their sharper exponents are unsupported. Weaker versions, however, follow from  $\text{BST}_{1/2}$  through the short-cycle removal of [ABKZ22]. The bound of [JX23, Theorem 1.14] assumes the Strong 3SUM hypothesis (integers in  $[-n^2, n^2]$ ), and the detection bound of [CX24, Theorem 4.7] assumes the Strong Exact Triangle hypothesis (weights in  $[-n, n]$ ). The algorithms of [AV26] cover all polynomially bounded inputs and refute both hypotheses, so these bounds are unsupported.

The 4-cycle bounds of [ABKZ22, Corollary 1.11] assume that detecting one triangle in a graph of maximum degree  $\sqrt{n}$  needs  $n^{2-o(1)}$  time. This hypothesis was never derived from 3SUM or APSP [ABKZ22], and the new algorithms do not touch it. The evidence offered for it in [ABKZ22, Section 7] is a reduction (Theorem 7.6) from Zero-Triangle with parts of sizes  $n$ ,  $\sqrt{n}$ ,  $\sqrt{n}$  and weights up to  $\sqrt{n}$ , under a Strong Zero-Triangle Conjecture stated for all parameters (Conjecture 7.5). That conjecture is refuted in its symmetric case, with parts of size  $n$  and weights up to  $n$ , by the Exact Triangle algorithm of [AV26], but the regime used by Theorem 7.6 is not known to be affected.

### 4.4 Strings, geometry and convolutions

The hardness of jumbled indexing was derived from Convolution-3SUM over  $\{-n^2, \dots, n^2\}$  and from its strong form over  $\{-n, \dots, n\}$  [ACLL14], and from 3SUM [GKLP16]. The hardness of local alignment over large alphabets was derived from 3SUM and from Max-Weight 4-Clique [AVW14]. All of these hypotheses are refuted, and none of the bounds is known to follow from a surviving hypothesis. The quadratic bound for binary local alignment rests on SETH [AVW14, Theorem 2] and is unaffected.

In computational geometry, GeomBase is equivalent to 3SUM [GO95] and now has a polynomially faster algorithm. Three collinear points, point covering and planar motion planning [GO95], minimum-area triangle [Kin04] and polygon containment under translation [BHP01] were shown hard by reductions from 3SUM in one direction only, so their quadratic bounds are now unsupported. Lower bounds in restricted decision-tree models remain for 3SUM itself [Eri99] and for three collinear points [ES95]. Problems shown as hard as sorting  $X + Y$  [Bar96], which concern an  $o(n^2 \log n)$  algorithm, have a base that the new algorithms do not affect. The problems equivalent to  $(\min, +)$ -convolution [CMWW19, Theorem 2] are faster, because  $(\min, +)$ -convolution reduces to 3SUM and to the  $(\min, +)$ -product [BCD+14, BIS17, CMWW19]. MonoConvolution, for its part, is equivalent to 3SUM at its  $n^{3/2}$  baseline [LPV20, Theorems 7, 8].

## 5 Why the analysis stops at $\ln 4/(5 \ln 5)$

The hypotheses of Section 3 are stated beyond the range of the known algorithms, and it is fair to ask whether a better base identity would move that range. In this section we collect what is known about this question. Most of it is due to Majdoddin [Maj26, Appendix A]. We add the enumeration of the shapes that could improve the closed-form range and Koszul-flattening bounds for identities with two inner outputs.

A *base identity* is a decomposition  $\sum_{\lambda=1}^R \varphi_{\lambda} \psi_{\lambda} \chi_{\lambda}$ , exact up to an error that the method tolerates or in the border sense, of the trilinear form  $\sum_{i \leq k, j \leq n} x_i y_j z_{ij} + F z_0$ , where  $F$  is a bilinear form whose rank  $q_0$  is the *inner length*. The method takes the  $L$ -th tensor power, keeps the output strings that contain  $z_0$  at exactly  $m = L/c$  positions, and computes only the leaf products that the wanted outputs need [AV26, Section 2]. A term is *private* if its output form is a multiple of  $z_0$ . A variable of  $F$  that also occurs in the outer product is *shared*. For one shared variable,  $S(k, n, q')$  denotes the tensor with inner form  $\sum_{s \leq q'} p_s q_s + c y_1$ , of effective inner length  $q' + 1$  [Maj26, Appendix A.5]. The five-term partial-product tensor of [Cop82, Maj26] is  $S(2, 2, 1)$ .

**The closed-form range.** The analysis of [Maj26] saves a polynomial factor exactly for  $\varepsilon < \varepsilon_5^* = \ln 4/(5 \ln 5)$  with the five-term tensor. That of [AV26] does so exactly for  $\varepsilon < \varepsilon_{10}^* = \ln 4/(5 \ln 10) = 0.1204$  with Schönhage’s ten-term identity. Both limits have the form

$$\varepsilon_R^* = \frac{2 \ln q_0}{R \ln R} \quad (1)$$

for an identity with  $R$  terms and inner length  $q_0$  per half of the squared run ( $R = 5, q_0 = 2$  and  $R = 10, q_0 = 4$ ). The reason is that in both cases the leaf count needs  $c > R$ , and at  $c = R$  the condition that a tile fits coincides with the cost of the encodings [Maj26, Proposition 8.1]. The derivation uses only these features, so we use (1) to compare identities with one private term in which every outer output lies in exactly one term. This is a limit of the closed-form analysis rather than of the method. With  $W$  known in advance and the exact leaf count, Majdoddin’s numerical range is about 0.1784 for the five-term tensor (see Section 3) and 0.1292 for Schönhage’s identity [Maj26, Section 8]. As noted in the introduction,  $\varepsilon_5^*$  is the exponent of Coppersmith’s rectangular matrix multiplication algorithm from the five-term tensor [Cop82].

**Barriers.** Majdoddin proves the following barriers.

- For  $k, n \geq 2$ , an exact identity whose inner variables are disjoint from the outer ones, with a harmless error, has  $q_0 \leq R - k - n$  [Maj26, Proposition A.2].

- A border identity has  $q_0 \leq R - k - n + 1$  with disjoint inner variables and effective inner length at most  $R - k - n + 2$  with one shared variable [Maj26, Propositions A.4, A.5].
- A run in which every output is a correct single output of a distinct entry and block cannot use two shared variables on the same side, nor, if the blocks that compute each entry are disjoint, one on each side [Maj26, Theorem A.9].
- $R_0$  private terms give a saving only if  $R_0^2 R_1^{2(c-1)} < q_0$ , where every outer output lies in at least  $R_1$  terms that cover no other outer output [Maj26, Lemma A.1].

The exact identity obtained from Schönage's by  $z_{11} \mapsto z_{11} + z_0$  reaches inner length  $R - k - n + 1$  but shares  $x_1$  and  $y_1$ , which Theorem A.9 rules out [Maj26, Appendix A.1]. Full products as base identities are excluded for many shapes [Maj26, Theorem A.14, Corollary A.18].

**One private term.** Combining the cap with (1) leaves only a few shapes.

**Proposition 5.1.** *Let a border identity (for instance an exact identity without error) have one private term, outer product  $\langle k, 1, n \rangle$  with  $k \leq n$ ,  $R = kn + 1$  terms, at most one shared variable and inner length  $q_0$ . If  $\varepsilon_R^* \geq \varepsilon^*$ , then  $(k, n, q_0)$  is  $(2, 2, 2)$ , the shape of the five-term tensor, for which  $\varepsilon_R^* = \varepsilon^*$ , or one of  $(1, 1, 2)$ ,  $(1, 2, 2)$ ,  $(1, 3, 2)$ ,  $(1, 4, 2)$ ,  $(2, 2, 3)$ ,  $(2, 3, 4)$ . No identity of the last six shapes exists.*

*Proof.* Since  $q_0 < R$ ,  $\varepsilon_R^* < 2/R \leq 1/6 < \varepsilon^*$  for  $R \geq 12$ , so  $kn \leq 10$ . Evaluating (1) for these shapes with  $q_0 \leq R - k - n + 2$  gives the list, with equality  $\varepsilon_R^* = 2 \ln 2 / (5 \ln 5) = \varepsilon^*$  for  $(1, 4, 2)$  and  $(2, 2, 2)$  (computed<sup>†</sup>). Each of the last six shapes has  $q_0 = R - k - n + 2$ , so it needs one shared variable, and up to exchanging left and right its tensor is  $S(k, n, q_0 - 1)$  or  $S(n, k, q_0 - 1)$ . For  $(1, n, 2)$ ,  $S(n, 1, 1)$  is concise of left dimension  $n + 2 > R$ , and  $S(1, n, 1)$  has border rank at least  $n + 2$  for  $n \leq 6$  [Maj26, Remark A.6]. For  $(2, 2, 3)$  and  $(2, 3, 4)$ ,  $S(2, 2, 2)$ ,  $S(2, 3, 3)$  and  $S(3, 2, 3)$  have border rank at least 6, 8 and 8 [Maj26, Remark A.7], more than  $R = 5$  and  $R = 7$ . Table 4 reproduces these bounds. Identities that are exact only up to a tolerated error are covered by [Maj26, Proposition A.2] when  $k, n \geq 2$  and their inner variables are disjoint from the outer ones, and not otherwise.  $\square$

With  $R_0 \geq 2$  private terms, (1) is not the range of the method. Adding  $R_0 - 1$  private terms  $p_s q_s z_0$  to the five-term tensor gives border identities with  $R = R_0 + 4$  terms and  $q_0 = R_0 + 1$ , for which (1) exceeds  $\varepsilon^*$  when  $R_0 \leq 5$ . By [Maj26, Lemma A.1], however, they give no saving at all, since  $R_0^2 \geq q_0$ . In general, when every outer output lies in a term that covers no other outer output, Lemma A.1 requires  $R_0^2 < q_0 \leq R - k - n + 2 = (k - 1)(n - 1) + R_0 + 1$ . For  $k = 1$  and for  $(k, n) = (2, 2)$  this is impossible, and for  $(k, n) = (2, 3)$  it forces  $R_0 = 2$  and  $q_0 = 5$ . No border identity attains this, because  $S(2, 3, 4)$  and  $S(3, 2, 4)$  have border rank at least  $9 > 8$  (Table 4). For larger shapes with  $R_0 \geq 2$  we do not know the range.

**Several inner outputs.** The barriers of [Maj26] assume one inner output per level. A natural way around them is to let the outer terms serve two inner products at once. The private-leaf count suggests that this would help. In the leaf model of [Maj26] with pruned encodings, which gives  $\varepsilon\gamma = 0.0061$  for the five-term tensor (against 0.005825 in [Maj26]), the profile with outer product  $\langle 2, 1, 2 \rangle$ , two private terms and four outer terms that each serve two inner outputs of length 2 scores 0.0102, and the analogous profile with outer product  $\langle 3, 1, 3 \rangle$  and two inner outputs of length 4 scores 0.0074 (computed<sup>†</sup>). Proposition 5.2 shows that neither profile exists when the inner products are disjoint.

**Proposition 5.2.** *The tensor  $\langle 2, 1, 2 \rangle \oplus \langle 1, 2, 1 \rangle \oplus \langle 1, 2, 1 \rangle$  has border rank at least 7, and  $\langle 3, 1, 3 \rangle \oplus \langle 1, 4, 1 \rangle \oplus \langle 1, 4, 1 \rangle$  has border rank at least 13, while the two profiles need 6 and 11 terms.*

*Proof.* The bounds come from Koszul flattenings [LO13], and the ranks are listed in Table 4.  $\square$

| Tensor                                                              | flattening            | rank        | bound            | terms   |
|---------------------------------------------------------------------|-----------------------|-------------|------------------|---------|
| $S(1, n, 1), n = 1, 2, 3, 4$                                        | $B (n = 1), A; p = 1$ | 3, 7, 9, 11 | 3, 3.5, 4.5, 5.5 | $n + 1$ |
| $S(2, 2, 2)$                                                        | $C, p = 1$            | 22          | 5.5              | 5       |
| $S(2, 3, 3)$                                                        | $C, p = 3$            | 141         | 7.05             | 7       |
| $S(3, 2, 3)$                                                        | $C, p = 1$            | 45          | 7.5              | 7       |
| $S(2, 3, 4); S(3, 2, 4)$                                            | $C, p = 3; p = 1$     | 161; 51     | 8.05; 8.5        | 8       |
| $\langle 2, 1, 2 \rangle \oplus \langle 1, 2, 1 \rangle^{\oplus 2}$ | $C, p = 2$            | 66          | 6.6              | 6       |
| $\langle 3, 1, 3 \rangle \oplus \langle 1, 4, 1 \rangle^{\oplus 2}$ | $C, p = 5$            | 3102        | 12.31            | 11      |

Table 4: Koszul-flattening lower bounds on border rank. The flattening is  $\Lambda^p V \otimes U^* \rightarrow \Lambda^{p+1} V \otimes W$  for a factor  $V$  ( $C$ : the output factor) and the bound is the rank divided by  $\binom{\dim V - 1}{p}$ ; “terms” is the number of terms the identity or profile would have. Ranks are exact over  $\mathbb{Q}$  and agree with ranks modulo  $2^{31} - 1$ .<sup>†</sup>

## 6 Open problems

We close with some questions that the classification leaves open.

1. Is  $\text{LSTH}_\delta$  true for some  $\delta$ ? Is hardness monotone in  $\delta$ , so that there is a threshold  $\delta^* \in [\varepsilon^*, 1]$  below which natural instances are easy and above which they are hard? Lemma 3.9 gives only a weak form of monotonicity. Can the numerical range 0.1784 of [Maj26, Section 8] be proved, and how far does the method reach with the exact leaf count?
2. Is there a base identity outside the classes of Section 5 whose range exceeds Coppersmith’s exponent? The thin full products that [Maj26, Appendix A.9] does not exclude, the smallest being  $\langle 34, 2, 35 \rangle$ , are candidates, and so are identities with several private terms, whose range is not known.
3. Can capped listing, triangle listing or Set Intersection with reporting be solved faster on lopsided instances, as counting can? A positive answer would bear on the listing hypotheses and on the status of [KPP16, Corollary 5.2].
4. Do the general-parameter and listing hypotheses  $\text{LST}(r, \sigma, \kappa)$  and  $\text{LSTL}(r, \sigma, \kappa, \lambda)$  follow from the natural ones, or from BST? Does BST imply  $\text{BST}_{1/2}$  [ABKZ22]?
5. Can the quadratic lower bounds for three collinear points and the other 3SUM-hard problems of computational geometry be derived from a surviving hypothesis, such as  $k$ -SUM for some  $k \geq 4$ , BST, or SETH? No polylogarithmic improvement is known for three collinear points [AV26].
6. Triangle detection in tripartite graphs with parts of sizes  $n, n$  and  $\sqrt{n}$  and  $O(n^{3/2})$  edges would have consequences for girth approximation if it could be solved in  $O(n^{2-\varepsilon})$  time [RV12, AV26]. These instances lie outside the range of Theorem 3.4. Can they be solved faster?

**Computations.** The values marked  $\dagger$  were computed with short Python scripts, the ranks in exact rational arithmetic with a cross-check modulo  $2^{31} - 1$ , and the ranges and leaf-model scores in floating point. The rank bounds in Table 4, except those for  $S(1, n, 1)$ , are also proved in Lean 4, where the flattenings are built from the tensors and the kernel checks the bounds over  $\mathbb{Q}$  using only the standard axioms. The Landsberg–Ottaviani inequality enters these proofs as an explicit hypothesis. The scripts, their outputs and the Lean files are in the supplement to this paper, <https://doi.org/10.5281/zenodo.23602531>.

## References

- [ABF23] A. Abboud, K. Bringmann, and N. Fischer. Stronger 3-SUM lower bounds for approximate distance oracles via additive combinatorics. In *STOC 2023*, pages 391–404. Full version: arXiv:2211.07058.
- [ABKZ22] A. Abboud, K. Bringmann, S. Khoury, and O. Zamir. Hardness of approximation in P via short cycle removal: cycle detection, distance oracles, and beyond. In *STOC 2022*, pages 1487–1500. Full version: arXiv:2204.10465.
- [ACLL14] A. Amir, T. M. Chan, M. Lewenstein, and N. Lewenstein. On hardness of jumbled indexing. In *ICALP 2014*, LNCS 8572, pages 114–125. Full version: arXiv:1405.0189.
- [AGV15] A. Abboud, F. Grandoni, and V. Vassilevska Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *SODA 2015*, pages 1681–1697.
- [AHY25] J. Alman, Y. Huang, and K. Yeo. Fine-grained complexity in a world without cryptography. In *EUROCRYPT 2025, Part VII*, LNCS 15607, pages 375–405.
- [AV14] A. Abboud and V. Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *FOCS 2014*, pages 434–443. Full version: arXiv:1402.0054.
- [AV26] J. Alman and V. Vassilevska Williams. Truly subquadratic 3SUM and truly subcubic APSP via triangles in sparse lopsided graphs. arXiv:2610.06783, 2026.
- [AVW14] A. Abboud, V. Vassilevska Williams, and O. Weimann. Consequences of faster alignment of sequences. In *ICALP 2014*, LNCS 8572, pages 39–51.
- [Bar96] A. Hernández Barrera. Finding an  $o(n^2 \log n)$  algorithm is sometimes hard. In *CCCG 1996*, pages 289–294.
- [BCD+14] D. Bremner, T. M. Chan, E. D. Demaine, J. Erickson, F. Hurtado, J. Iacono, S. Langerman, M. Pătraşcu, and P. Taslakian. Necklaces, convolutions, and  $X + Y$ . *Algorithmica* 69(2):294–314, 2014.
- [BHP01] G. Barequet and S. Har-Peled. Polygon containment and translational min-Hausdorff-distance between segment sets are 3SUM-hard. *Int. J. Comput. Geom. Appl.* 11(4):465–474, 2001.
- [BIS17] A. Backurs, P. Indyk, and L. Schmidt. Better approximations for tree sparsity in nearly-linear time. In *SODA 2017*, pages 2215–2229.
- [BNS19] J. van den Brand, D. Nanongkai, and T. Saranurak. Dynamic matrix inverse: improved algorithms and matching conditional lower bounds. In *FOCS 2019*, pages 456–480. Full version: arXiv:1905.05067.
- [BPWZ14] A. Björklund, R. Pagh, V. Vassilevska Williams, and U. Zwick. Listing triangles. In *ICALP 2014*, LNCS 8572, pages 223–234.

- [BSZ24] J. van den Brand, Z. Song, and T. Zhou. Algorithm and hardness for dynamic attention maintenance in large language models. In *ICML 2024*, PMLR 235, pages 49008–49028.
- [CMWW19] M. Cygan, M. Mucha, K. Węgrzycki, and M. Włodarczyk. On problems equivalent to  $(\min, +)$ -convolution. *ACM Trans. Algorithms* 15(1):14:1–14:25, 2019. Full version: arXiv:1702.07669.
- [Cop82] D. Coppersmith. Rapid multiplication of rectangular matrices. *SIAM J. Comput.* 11(3):467–471, 1982.
- [CVX22] T. M. Chan, V. Vassilevska Williams, and Y. Xu. Hardness for triangle problems under even more believable hypotheses: reductions from real APSP, real 3SUM, and OV. In *STOC 2022*, pages 1501–1514.
- [CX24] T. M. Chan and Y. Xu. Simpler reductions from exact triangle. In *SOSA 2024*, pages 28–38. Full version: arXiv:2310.11575.
- [DG25] I. Dinur and A. Golovnev. Improved time-space tradeoffs for 3SUM-Indexing. arXiv:2512.04258, 2025.
- [Eri99] J. Erickson. Lower bounds for linear satisfiability problems. *Chicago J. Theoret. Comput. Sci.* 1999(8).
- [ES95] J. Erickson and R. Seidel. Better lower bounds on detecting affine and spherical degeneracies. *Discrete Comput. Geom.* 13:41–57, 1995.
- [FKP24] N. Fischer, P. Kaliciak, and A. Polak. Deterministic 3SUM-hardness. In *ITCS 2024*, LIPIcs 287, 49:1–49:24. Full version: arXiv:2310.12913.
- [FPS26] N. Fischer, A. Polak, and J. Schmidt. 3SUM is really hard: a real-to-integer reduction. In *FOCS 2026*. arXiv:2610.10153.
- [GGHPV20] A. Golovnev, S. Guo, T. Horel, S. Park, and V. Vaikuntanathan. Data structures meet cryptography: 3SUM with preprocessing. In *STOC 2020*, pages 294–307.
- [GKLP16] I. Goldstein, T. Kopelowitz, M. Lewenstein, and E. Porat. How hard is it to find (honest) witnesses? In *ESA 2016*, LIPIcs 57, 45:1–45:16. Full version: arXiv:1706.05815.
- [GKLP17] I. Goldstein, T. Kopelowitz, M. Lewenstein, and E. Porat. Conditional lower bounds for space/time tradeoffs. In *WADS 2017*, LNCS 10389, pages 421–436. Full version: arXiv:1706.05847.
- [GLP19] I. Goldstein, M. Lewenstein, and E. Porat. On the hardness of set disjointness and set intersection with bounded universe. In *ISAAC 2019*, LIPIcs 149, 7:1–7:22. Full version: arXiv:1910.00831.
- [GO95] A. Gajentaan and M. H. Overmars. On a class of  $O(n^2)$  problems in computational geometry. *Comput. Geom.* 5(3):165–185, 1995.
- [HKNS15] M. Henzinger, S. Krinninger, D. Nanongkai, and T. Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *STOC 2015*, pages 21–30.
- [HLS24] B. Haeupler, Y. Long, and T. Saranurak. Dynamic deterministic constant-approximate distance oracles with  $n^\epsilon$  worst-case update time. In *FOCS 2024*, pages 2033–2044.
- [JX23] C. Jin and Y. Xu. Removing additive structure in 3SUM-based reductions. In *STOC 2023*, pages 405–418. Full version: arXiv:2211.07048.
- [Kin04] J. King. A survey of 3SUM-hard problems. Manuscript, 2004.

- [KL26] S. R. Karingula and S. Lovett. A simplified proof of polynomial improvements for 3SUM and APSP. arXiv:2610.15663, 2026.
- [KP19] T. Kopelowitz and E. Porat. The strong 3SUM-INDEXING conjecture is false. arXiv:1907.11206, 2019.
- [KPP15] T. Kopelowitz, S. Pettie, and E. Porat. Dynamic set intersection. In *WADS 2015*, LNCS 9214, pages 470–481. Full version: arXiv:1407.6755.
- [KPP16] T. Kopelowitz, S. Pettie, and E. Porat. Higher lower bounds from the 3SUM conjecture. In *SODA 2016*, pages 1272–1287. Full version: arXiv:1407.6756v3.
- [KPS17] M. Künnemann, R. Paturi, and S. Schneider. On the fine-grained complexity of one-dimensional dynamic programming. In *ICALP 2017*, LIPIcs 80, 21:1–21:15.
- [LLV19] R. LaVigne, A. Lincoln, and V. Vassilevska Williams. Public-key cryptography in the fine-grained setting. In *CRYPTO 2019, Part III*, LNCS 11694, pages 605–635.
- [LO13] J. M. Landsberg and G. Ottaviani. Equations for secant varieties of Veronese and other varieties. *Ann. Mat. Pura Appl.* 192(4):569–606, 2013.
- [LPV20] A. Lincoln, A. Polak, and V. Vassilevska Williams. Monochromatic triangles, intermediate matrix products, and convolutions. In *ITCS 2020*, LIPIcs 151, 53:1–53:18. Full version: arXiv:2009.14479.
- [Maj26] R. Majdoddin. Thin matrix products up to  $D \leq N^{0.172}$  and lossless reductions for 3SUM, Exact Triangle and APSP, with 3SUM over abelian groups and a dichotomy for coordinatewise problems. arXiv:2610.17420, 2026.
- [Pät10] M. Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *STOC 2010*, pages 603–610.
- [RV12] L. Roditty and V. Vassilevska Williams. Subquadratic time approximation algorithms for the girth. In *SODA 2012*, pages 833–845.
- [Shr26] shreshth-s. Thin-matrix results. Unrefereed repository, <https://github.com/shreshth-s/thin-matrix-results>, accessed 23 October 2026.
- [SY26] B. Saha and C. Ye. Simple, faster (min, +) matrix multiplication and APSP. arXiv:2610.20297, 2026.
- [VW18] V. Vassilevska Williams and R. Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM* 65(5):27:1–27:38, 2018.
- [VX20] V. Vassilevska Williams and Y. Xu. Monochromatic triangles, triangle listing and APSP. In *FOCS 2020*, pages 786–797. Full version: arXiv:2007.09318.
- [VXXZ24] V. Vassilevska Williams, Y. Xu, Z. Xu, and R. Zhou. New bounds for matrix multiplication: from alpha to omega. In *SODA 2024*, pages 3792–3835.

TEJASVI SINGH TOMAR

ORCID: <https://orcid.org/0000-0003-4668-3639>