

A polynomial versus superpolynomial gap in vertex-distribution-free graph testing

Alec Kriebel

Independent researcher

<https://orcid.org/0009-0001-9320-500X>

17 October 2026 Version 1.0.0

Abstract

We give a negative answer to the comparison question of Gishboliner and Shapira in Oberwolfach Report 19/2019, Problem 5, and their later Problem 5.6. The fixed property consisting of K_4 -free graphs that become 3-colorable after deleting at most one vertex has a polynomial classical canonical tester, for every input size. In the vertex-distribution-free model, every adaptive two-sided tester has superpolynomial total sample-plus-adjacency cost as the proximity parameter tends to zero. The lower bound follows by an elementary tripartite-cover and heavy-apex reduction from the established two-sided triangle-testing lower bound. All distances allow arbitrary edge additions and deletions. The reduction remains valid when the tester is given the vertex set, its size, the full positive rational distribution, and the heavy vertex. A self-contained classical sampling argument and exact finite verification materials accompany the note.

1 The question and the fixed property

All graphs are finite, simple and undirected. Sampling inputs have a nonempty vertex set. For a probability distribution D on that set, use the unordered-pair normalization

$$d_D(G, F) = \sum_{\{u,v\} \in E(G) \triangle E(F)} D(u)D(v), \quad \text{dist}_D(G, \mathcal{R}) = \min_{F \in \mathcal{R}, V(F)=V(G)} d_D(G, F).$$

Thus uniform distance is the number of changed edges divided by n^2 . We call distance at least ε ε -far. Strict-threshold conventions give the same conclusions after a fixed proximity rescaling.

A tester accepts members and rejects ε -far inputs with probability at least $2/3$. The VDF oracle supplies independent samples from the unknown vertex distribution and adjacency answers. We count their *total* number, as in the original workshop model [1, p. 1138]. The comparison question asks whether an extendable hereditary property has the same, or nearly the same, testing complexity in the VDF and uniform dense models [1, Problem 5], [2, Problem 5.6]. Hereditary means closed under induced subgraphs; extendable means that every member has a one-vertex extension which is also a member.

Let $\mathcal{C}_3$ be 3-colorability, let $\mathcal{A}_3$ mean that deletion of at most one vertex leaves a 3-colorable graph, and fix

$$\mathcal{P} = \{K_4\text{-free graphs}\} \cap \mathcal{A}_3.$$

This property is hereditary, closed under deleting edges, and extendable by adding an isolated vertex. Moreover, $\mathcal{C}_3 \subseteq \mathcal{P}$.

Theorem 1. *For this single fixed property $\mathcal{P}$:*

1. *There is a one-sided classical canonical tester, without knowing n , using $O(\varepsilon^{-6} \log(1/\varepsilon))$ independent uniform samples and $O(\varepsilon^{-12} \log^2(1/\varepsilon))$ adjacency queries as $\varepsilon \rightarrow 0$.*
2. *Along a sequence of positive rational parameters $\varepsilon \rightarrow 0$, every two-sided adaptive VDF tester has worst-case total oracle cost $\exp(\Omega(\log^2(1/\varepsilon)))$. This holds even with the vertex set, its size, the complete positive rational weight vector and the distinguished heavy vertex supplied.*

Consequently VDF complexity admits no fixed polynomial bound in the classical canonical complexity and $1/\varepsilon$, even after any fixed positive constant rescaling of classical proximity.

The lower-bound ingredients are established results: polynomial colorability sampling [4, 5], the hard-triangle constructions and their two-sided extensions [6, 7, 8], and the VDF framework [3, 2]. The reduction below transfers triangle hardness to a property with an easy classical canonical tester. It supplies the full negative comparison, rather than a new triangle lower bound. Fei’s recent edge-distribution work [9] also uses classical parity constructions; its default sample oracle and arbitrary pair weights differ from the present adjacency oracle and vertex-product distance. The current priority audit, including primary theorem bodies and version histories, identified no earlier full comparison answer. This is a bounded literature finding: the 2021 journal version-of-record body of [2] was unavailable, although its arXiv v5, ECCC revision 2 and author manuscripts were inspected; forward-citation retrieval was incomplete. Detailed coverage and limitations accompany the package.

2 A classical tester at every input size

We include a conservative form of the standard branching-list colorability argument [4, 5], so no large-input qualification is needed.

Lemma 2. *For $0 < \varepsilon \leq 1/8$, put*

$$h = \lfloor 6/\varepsilon \rfloor + 1, \quad \ell = \left\lceil (2/\varepsilon) \log(12 \cdot 3^h) \right\rceil, \quad r = h\ell.$$

If a nonempty graph is ε -far from $\mathcal{C}_3$, then r independent uniform samples induce a graph outside $\mathcal{C}_3$ with probability at least $11/12$. Repeated labels are identified.

Proof. Fix a proper partial coloring on S . For $v \notin S$, let $L(v)$ be the colors absent from its neighbors in S ; let U comprise the empty-list vertices. For $v \notin S \cup U$, define

$$\delta(v) = \min_{i \in L(v)} |\{u \notin S \cup U : uv \in E(G), i \in L(u)\}|.$$

Color every such v by a minimizing color and color U arbitrarily. Deleting monochromatic edges gives a 3-colorable graph with at most $n|U| + \sum_v \delta(v)$ changes. If $W = \{v \notin S \cup U : \delta(v) \geq \varepsilon n/2\}$, fairness implies

$$\varepsilon n^2 \leq n|U| + \sum_v \delta(v) \leq n(|U| + |W|) + \varepsilon n^2/2.$$

Hence $|U \cup W| \geq \varepsilon n/2$. The list potential $\Phi = \sum_{v \notin S} |L(v)|$ starts at $3n$. Coloring a vertex of W with any feasible color reduces it by at least $\varepsilon n/2$; a vertex in U has no feasible color.

Partition the samples into h independent blocks of ℓ samples. At each node of a ternary tree of proper partial colorings, choose the first sample in that level’s block lying in its $U \cup W$, and

branch on its feasible colors. A node fails if its block misses $U \cup W$. Conditional on all previous blocks, this probability is at most $\exp(-\varepsilon\ell/2) \leq (12 \cdot 3^h)^{-1}$. There are fewer than 3^h possible node addresses at depths $0, \dots, h-1$. Summing these conditional bounds shows that no node fails with probability at least $11/12$; nodes sharing a block need not be independent.

On this event no feasible branch survives $h > 6/\varepsilon$ levels, since Φ stays nonnegative. A proper coloring of the sampled induced graph would follow a surviving branch, a contradiction. Each chosen vertex lies outside its path's previously colored set, so repeated samples cause no difficulty. $\square$

Set $N = 12r^2$ and $M = \lceil N \log(3N) \rceil$. Sample M vertices uniformly with replacement, query every pair of distinct sampled labels, and accept exactly when the induced graph belongs to $\mathcal{P}$. Membership can be decided by finite exhaustive search; the assertion concerns oracle complexity, with no polynomial running-time claim. Completeness is one-sided by heredity.

If G is ε -far from $\mathcal{P}$, it is also ε -far from $\mathcal{C}_3$. If $n \geq N$, the first two r -blocks both induce non-3-colorable graphs and have disjoint label sets with probability at least

$$1 - \frac{1}{12} - \frac{1}{12} - \frac{r^2}{n} \geq \frac{3}{4}.$$

Deleting one vertex cannot destroy both disjoint witnesses, so the whole sample lies outside $\mathcal{A}_3$, hence outside $\mathcal{P}$. If $1 < n < N$, the probability of missing any vertex is at most

$$n(1 - 1/n)^M \leq ne^{-M/n} \leq Ne^{-M/N} \leq 1/3.$$

The case $n = 1$ is immediate. Seeing all vertices rejects every far input. The algorithm need not know which case holds; $M \geq 2r$. Since $h = O(\varepsilon^{-1})$, $\ell = O(\varepsilon^{-2})$ and $r = O(\varepsilon^{-3})$, this proves part 1 of Theorem 1. For $\varepsilon > 1/8$, use parameter $1/8$. The empty graph belongs to $\mathcal{P}$ and can be accepted if supplied as an explicit empty input; it carries no sampling distribution.

3 Transferring triangle hardness

Write $\mathcal{T}$ for triangle-freeness.

Lemma 3 (Heavy apex). *If H is 3-colorable on $n \geq 1$ vertices, add a universal vertex a to form $J(H)$ and set $D(a) = 1/2$, $D(v) = 1/(2n)$ on the tail. Then*

$$\text{dist}_D(J(H), \mathcal{P}) = \frac{1}{4} \text{dist}_{\text{unif}}(H, \mathcal{T}).$$

Proof. Delete a minimum triangle-hitting set of tail edges. The resulting tail remains 3-colorable and triangle-free, so its join with a is K_4 -free and belongs to $\mathcal{P}$. Each tail edge costs $1/(4n^2)$, giving the upper bound.

For any repair $F \in \mathcal{P}$, let S be the tail vertices whose edge to a was deleted. Its apex-edge cost is $\delta_a = |S|/(4n)$; denote its tail-edit cost by δ_t , including additions. The restriction $F[V(H) \setminus S]$ is triangle-free, since a triangle together with a would give a K_4 . Use that restriction and isolate S to repair H . This changes at most $4n^2\delta_t + n|S|$ edges. Therefore

$$\text{dist}_{\text{unif}}(H, \mathcal{T}) \leq 4\delta_t + |S|/n = 4(\delta_t + \delta_a).$$

Minimizing over all repairs proves the reverse inequality. $\square$

Lemma 4 (Tripartite cover). *For any graph H on $n \geq 1$ vertices, let $L(H)$ have vertex set $\{1, 2, 3\} \times V(H)$, with $(i, u)(j, v)$ an edge exactly when $i \neq j$ and $uv \in E(H)$. Then $L(H)$ is 3-colorable, $H \in \mathcal{T}$ implies $L(H) \in \mathcal{T}$, and*

$$\text{dist}_{\text{unif}}(L(H), \mathcal{T}) \geq \frac{1}{9} \text{dist}_{\text{unif}}(H, \mathcal{T}).$$

Proof. The three parts color $L(H)$. A triangle projects to three distinct pairwise adjacent vertices of H . Triangle-free repairs can be restricted to deletions: intersecting a repair's edge set with the original graph preserves triangle-freeness and never increases cost. Take a minimum deletion repair of $L(H)$, with t deleted edges. Delete uv from H if any of its six cross-copy edges was deleted. This deletes at most t original edges. Every surviving triangle would lift, one vertex in each part, to a surviving triangle in the cover, a contradiction. Dividing the two deletion counts by n^2 and $(3n)^2$ gives the claim. $\square$

Combining the lemmas gives, for every H ,

$$\text{dist}_D(J(L(H)), \mathcal{P}) \geq \frac{1}{36} \text{dist}_{\text{unif}}(H, \mathcal{T}), \quad (1)$$

where $D(a) = 1/2$ and every other vertex has weight $1/(6n)$. Triangle-free H gives a member of $\mathcal{P}$.

To use a VDF tester at proximity ε , simulate it on $(J(L(H)), D)$. A sample is generated locally from this explicit distribution. Self-queries return zero, distinct apex–tail pairs are adjacent, and same-part pairs or copies of the same original vertex are nonadjacent. Every remaining adjacency uses at most one query to H . All promised extra information is supplied explicitly. Consequently a VDF tester with total budget Q gives a two-sided adaptive tester for triangle-freeness at proximity 36ε , using at most Q adjacency queries. Samples require no graph queries. If the VDF execution uses s samples and e adjacency calls, the sampled original indices and query endpoints involve at most $s + 2e \leq 2Q$ source vertices. This simulation also works for randomized, history-dependent choices of queries and for early stopping.

The established two-sided triangle-testing lower bound is $\exp(\Omega(\log^2(1/\delta)))$ in the uniform dense graph model: see [7, Theorem 4(ii) and §7], including its explicit undirected consequence. That work gives the vertex count and permits adaptive adjacency queries; its inspected-vertex convention, with all induced pairs available, changes the present query count by at most a constant factor. Fixed changes in the unordered/ordered distance normalization also leave this bound unchanged. Applying it at $\delta = 36\varepsilon$ to the simulation proves part 2. It suffices to retain a rational sequence of parameters; the weights above are already strictly positive and rational. Combining this lower bound with the polynomial classical bound proves the final assertion of Theorem 1.

4 Scope and verification

The separation concerns total VDF oracle cost. A sample-only lower bound with unlimited free arbitrary-name adjacency queries is not asserted. The original proof preserved in the supporting package also gives an explicit rational hard family and a direct parity/exposure proof, with all its combinatorial and probabilistic estimates. The shorter proof here imports the established triangle lower bound and gives the complete reduction instead.

Independent programs test arbitrary weighted repairs on small graphs, the cover reduction, coloring lists and finite adaptive queries. These finite checks test the arguments; the universal claims rest on the proofs and the cited triangle theorem. The original programs and their hashes accompany the reports. The public checkers raise errors explicitly, and run their checks with Python optimization enabled as well as in ordinary mode.

AI use and review status. AI tools were used extensively in developing and checking the arguments, computations, literature review and writing. The note and supplements are unrefereed. Independent adversarial reviews were AI-assisted; no external human peer review, journal acceptance or formal proof-assistant certification is claimed.

References

- [1] L. Gishboliner, joint work with A. Shapira, contribution to *Combinatorics, Probability and Computing*, Oberwolfach Report 19/2019, Problem 5, p. 1139. Publisher report.
- [2] L. Gishboliner and A. Shapira, *Testing graphs against an unknown distribution*, Israel J. Math. **245** (2021), 787–837, doi:10.1007/s11856-021-2228-8. Inspected manuscript: arXiv:1905.09903v5, §5.3.
- [3] O. Goldreich, *Testing graphs in vertex-distribution-free models*, STOC 2019, 527–534. Inspected extended author manuscript, 4 October 2021. <https://www.wisdom.weizmann.ac.il/~oded/R3/vdf.pdf>.
- [4] O. Goldreich, S. Goldwasser and D. Ron, *Property testing and its connection to learning and approximation*, J. ACM **45** (1998), 653–750.
- [5] N. Alon and M. Krivelevich, *Testing k -colorability*, SIAM J. Discrete Math. **15** (2002), 211–227.
- [6] N. Alon, *Testing subgraphs in large graphs*, Random Structures Algorithms **21** (2002), 359–370.
- [7] N. Alon and A. Shapira, *Testing subgraphs in directed graphs*, J. Comput. System Sci. **69** (2004), 354–382, doi:10.1016/j.jcss.2004.04.008. Inspected author manuscript: <https://www.math.tau.ac.il/~asafico/subdir.pdf>.
- [8] N. Alon, T. Kaufman, M. Krivelevich and D. Ron, *Testing triangle-freeness in general graphs*, SIAM J. Discrete Math. **22** (2008), 786–819.
- [9] Y. Fei, *Testing properties of edge distributions*, arXiv:2603.22702v2, 12 August 2026, especially §7.1.