Inflatable graph properties and naturalization of property tests through strong canonicality

Eldar Fischer* Eyal Rozenberg*

May 21, 2011

Abstract

We consider *natural* graph property tests, which act entirely independently of the size of the graph being tested (not just having a number of queries independent of the size). We introduce the notion of graph properties being *inflatable* — closed under taking (balanced) blowups — and show that the query complexity of natural tests for a property are related to the degree to which it is approximately hereditary and approximately inflatable. Specifically, we show that for properties which are almost hereditary and almost inflatable, any test can be made natural, with a polynomial increase in the number of queries. The naturalization is carried out as a sort of extension of the canonicalization due to Goldreich and Trevisan in [15], so that natural canonical tests can be described as *strongly canonical*. In the reverse direction, we show that properties admitting natural tests are approximately inflatable and approximately hereditary, with these parameters depending on the test's number of queries.

Using the technique for naturalization, we restore in part the claim in [15] (which was qualified in the errata [16]) regarding testing hereditary properties by ensuring that a small random subgraph itself satisfies the tested property. This restoration allows us to generalize the result of Alon and Shapira in [5], regarding the lower bound on triangle-freeness testing: Any lower bound, not only the currently established quasi-polynomial one, on one-sided testing for triangle freeness holds essentially for two-sided testing as well. We also explore the relations of the notion of inflatability and other already-studied features of properties and property tests in the dense graph model such as one-sidedness, heredity, and proximity-oblivion.

To Do (not necessarily before finishing my Ph.D. or resubmitting if we don't get accepted to RANDOM):

- Try to salvage the lemma saying that $s' < s$ natural test for a hereditary-downto-s property implies a one-sided s-test.
- Try to construct a hereditary property, which is not 1/10 - approximately-inflatable, even with a threshold / on the average.
- Try to characterize the properties which are both proximity-oblivious-testable and hereditary (but not necessarily inflatable).
- Can we say something interesting about estimation / tolerant testing?

*Department of Computer Science, Technion, Haifa, Israel. Email: {eldar|eyalroz}@cs.technion.ac.il

1 Introduction

[I separated the discussion of triangle testing. Need to decide whether we integrate the concept of 'strong canonicality' into the introduction as well.]

In Property Testing, one is allowed oracle access to some combinatorial object, and must distinguish with high probability between the case of this object satisfying a certain property, and the case of the object being far from satisfying it by some distance measure. The study of property testing in the context of graphs began with the work of Goldreich, Goldwasser and Ron in [11]. Their paper introduced the *dense model* for graph testing; in this model, graphs on n vertices are close to each other if one needs to add and/or remove an ε-fraction of all possible $\binom{n}{2}$ edges from one graph to convert it into the other; sparse graphs are all close to being empty, hence the model's name. Beginning with [11], much work has been done to characterize which graph properties admit tests in this model in certain query complexity classes with respect to n and ε, and with certain features of the tests.

While [11] and some later contributions deal with bounds on the number of queries in terms of n (specifically, [12]; see discussion below), most of the study of testing in the dense model has focused on testing properties using a number of queries independent of the size of the input graph, depending only on the distance parameter ε (called simply "testable" properties). [11] established a large class of properties as testable, and posed the characterization of the class of properties testable in the dense model as an open problem. In the following decade, a series of results gradually progressed towards this goal, and a characterization was achieved by Alon, Fischer, Newman and Shapira in [4], and independently by [7] (in terms of graph limits); while the techniques of [4] rely heavily on the use of Szemerédi's Regularity Lemma, some of the ideas used there have bearing on this paper, as will be described below.

A notable result on the way to characterizing the testable properties is the work of Alon and Shapira in [6], proving all *hereditary* graph properties to be testable. In fact, this result establishes testability with one-sided error: Tests which have zero probability of rejecting a graph in the property. This kind of link between a feature of a class of properties and a feature of the test is a main focus of this paper.

Beyond the fundamental testability of properties, much study has focused on properties' *query complexity*: The dependence of the number of queries necessary for testing a property on the distance parameter ε; it is also a measure of practicality for 'real life' application of property testing. Most general testability results in the dense model are based on the use of Szemerédi's regularity lemma, incurring a prohibitive dependence on ε, while the class of generalized partition properties proved testable in [11] have a mere polynomial dependence on ε in the number of queries.

Links between features of properties and features of tests have been studied also in this context of bounding query complexity: Specifically, the adaptivity of property tests — the possibility of making different queries based on previous query results — has been shown to offer a definite advantage in query complexity. Goldreich and Ron showed in [14] (following the earlier work of Gonen and Ron in [17]) that some testable graph properties, exhibit a polynomial gap between an upper bound on the query complexity of adaptive tests, and a lower bound on the query complexity non-adaptive tests. In other testing models such gaps can be even exponential.

Goldreich and Trevisan's [15] includes two results tying several features of properties and tests together. Their paper defined the 'canonicality' feature of property tests: A *canonical* test uniformly samples a small induced subgraph, queries it entirely, and makes a deterministic decision based on this subgraph. [15] proved that any test can be made canonical with at most about a squaring of its number of queries, immediately implying that the gap between adaptive and non-adaptive query complexity is at most quadratic. [15] also included a proposition communicated by Noga Alon: Testable hereditary properties can be tested by merely ensuring that most small induced subgraphs themselves satisfy the property (with the same mild increase in the number of queries for the switch to a canonical test). Unfortunately, it later turned out that this second result only holds for tests which are *natural*: Tests acting independently of the size of the input graph. This qualification, and the definition of a natural test, appear in the errata [16].

It seems odd that properties which in many cases are highly 'local' in their definition should have tests significantly benefitting somehow from basing their action on the order of the entire input graph. If we constrain ourselves to properties with features preventing blatant 'pathologies' which preclude natural tests (e.g. the property of graphs having an odd number of vertices) — one tends to believe that property tests are 'essentially natural', so that perhaps one can 'smooth out' any non-natural artificial dependence of tests on n.

With regards to the relevant feature constraints, we have already mentioned a concept relating a property to itself at different orders — heredity; but this only covers one direction of change in n : As the input order increases, the set of forbidden subgraphs gains more and more relevant elements, so in a sense one expects the the set of graphs accepted by a tester to shrink gradually. The idea for the converse feature, implying a set of acceptable graphs shrinking as n *decreases*, came to us from earlier works on lower bound results. The concept which springs to mind is graph blowups: The dense model hierarchy theorem in [12] relies on the use of blowups; and Alon's triangle testing lower bound construction in [2] involves a blowup (more on this below). The feature we define is being *inflatable* — being closed to blowups. This takes care, specifically, of the pathology of graphs going from satisfying a property at order n to being very far from it by merely adding a vertex.

With regards to the idea of 'smoothing out' non-naturality, a typical example would be a test which arbitrarily rejects some specific queried subgraph at, say, even orders, and accepts it at odd ones. If this subgraph is very unlikely to appear in graphs in the property, a natural test could be 'spoiled' by adding this behavior to it, while still remaining a valid test. However, this can only be done for a single possible queried subgraph, or a few of them — such behavior is impossible with all acceptable graphs, nor in fact with any subset of them which has an overall high probability of being sampled. This leads one to recall that, in [4], the characterization of testability uses the set of all subgraphs of a fixed order accepted by a canonical test. Even more relevant is Fischer and Newman's [9] (proving that testable properties are also estimable, a key result necessary for the characterization in [4]), where it is observed that if one has an estimate of the subgraph distribution, one knows in particular whether a test querying subgraphs of this order accepts with high probability or not. In fact, disregarding the heavy use of Szemerédi's regularity lemma in [9], its result is based on estimating the subgraph distribution up to a small variation distance.

Relating to all of the above, in this paper we prove that, indeed, tests can be made natural, with a polynomial penalty in the number of queries, under the constrains of heredity and inflatability (even with a little relaxation), with the analysis focusing on the distribution of subgraphs of a fixed order and its behavior in subgraphs and blowups.

An immediate noteworthy applications of our naturalization technique regards lower bounds on the query complexity of testing triangle-freeness.

The question of whether a graphs contains a triangles (or more generally, whether it is free of certain induced subgraphs) is a very basic question to ask when considering the queries one can make in the dense model, and is perhaps the most studied property in this model. While it is known to be testable, there is a vast gap between the lower and upper bounds for it. The best upper bound until recently was obtained by Alon, applying Szemerédi's regularity lemma [1]; a proof sketch appears in Fischer's survey [8], and a more general treatment covering any induced subgraph is found in [3]. This construction yields a query complexity equal to a (single) tower function of height polynomial in $1/\varepsilon$; recently, Fox has proven in [10] a tower function upper bound for forbidden induced subgraphs, whose height is only logarithmic in $1/\varepsilon$, by a technique similar to the one used for proving Szemerédi's Regularity Lemma itself, customized to the problem of subgraph-freeness.

It is the establishment of lower bounds for testing triangle-freeness, however, that has seen use of relations between feature of properties and features of property tests, and which allows us to apply our result regarding natural tests, as triangle-freeness is both a hereditary and an inflatable property.

The standard approach for proving lower bounds on a property's query complexity is a principle observed by Yao in [19]: The average probability of success of a random algorithm on an input from a certain distribution is no lower than the best probability of success of a deterministic algorithm on the same input distribution. For a query complexity lower bound, one generally devises two different distributions of input graphs, one supported on graphs satisfying a property with high probability, and the other supported on graphs far from satisfying the property with high probability: If these can be shown not to be distinguishable from each other using a certain number of queries, then this number is insufficient for testing the property, since a test will either accept far graphs, or reject satisfying graphs, with an overly high probability. The requirements from the 'adversarial' distributions becomes more complex if the test is allowed to be adaptive.

There is challenge, however, even in the construction of a single graph, rather than a distribution, which is hard to distinguish as being far from a property — as such a construction can yield lower bound results. This is the case for one-sided testing of triangle-freeness: A test querying a subgraph which in itself contains no triangles would have to accept, as it is possible that there are no edges in the graph except the queried ones. A bound therefore requires only constructing a single graph (for every order n) which has very few triangles, but no small set of edges intersecting all of them. Indeed, such a construction by Alon in [2] established the best such bound known, a bound quasi-polynomial in $1/\varepsilon$.

A one-sided query complexity lower bound can be converted into a two-sided lower bound, provided one can show that any test for the property can be made one-sided. Indeed, this was to be possible using Alon's proposition appearing in [15] — but as mentioned above, the proof there relied on the test being natural. Alon and Shapira worked in [5] around this issue, by proving the same quasi-polynomial lower bound of the one-sided case for any triangle freeness test — directly, using Yao's method. Their proof, however, is tailored to the specific construction, and may not be able to convert possible stronger one-sided lower bounds on triangle testing to general two-sided bounds.

Using our naturalization technique, in this paper we restore in part the proposition testing hereditary properties, thus proving that the query complexity of testing for triangle-freeness is entirely determined by the one-sided-error query complexity, in general rather for the best lower bound currently known.

The rest of this paper is organized as follows. Our main result regarding naturalization of tests, as well as its significant implications, are stated in Section 3. This follows Section 2, which includes the formal definitions of the testing models and of the various features of properties and of tests, as well as a few technical lemmata regarding graph blowups and their effects on distances and subgraph distributions. Proof of our main Theorem 1 regarding the naturalization of tests constitutes Section 4. Following is Section 5, containing further discussion of inflatability and natural testability, and proofs of additional results, including the weak reverse of Theorem 1. We conclude by putting forward some open questions in Section 6.

2 Preliminaries

2.1 Graph properties and the dense model for property testing

Definition 2.1. The *absolute distance* between two graphs G, H of order n is the number of edges one has to add and/or remove in G to make it into an isomorphic copy of H; in other words, it is $\binom{n}{2}^{-1}$ times the minimum over all bijections $\phi : V(G) \to V(H)$ of the number of edge discrepancies — the number of elements in the symmetric difference between $E(H)$ and $\{\{\phi(u'), \phi(v')\} \mid \{u', v'\} \in E(G)\}$. The *distance* $\text{dist}\big(G, H\big)$ between G and H is the absolute distance between them normalized by a factor of $\binom{n}{2}^{-1}$.

Two graphs are said to be ε-*far* if their distance is at least ε.

Definition 2.2. A *property* of graphs is a set $\Pi = \bigcup_{n=1}^{\infty} \Pi_n$ of graphs, closed under graph isomorphism, where Π_n is supported on graphs of order n.

A graph is set to *satisfy* a property Π if it is an element of the set Π; a graph G of order n is said to be ε-*far* from satisfying a property Π if G is ε-far from every graph $H \in \Pi_n$.

Definition 2.3. A *property test* for a graph property Π is a probabilistic oracle machine which, given the values (n, ε), as well oracle access to a graph G of order n, makes a certain number of edge queries ("is there an edge between the vertices u and v?"), and distinguishes with probability at least $2/3$ between the case of G being in Π and the case of G being ε-far from Π. The (possibly adaptive) number and and choice of queries, as well as the rest of the algorithm, may in general depend on the value of n, as can the decision to accept or reject.

Note. Many results regard tests for specific values of ε, rather than tests receiving ε as a parameter (the difference between these settings regards the question of computability as a function of ε). The results in this paper hold for both settings.

The above, traditional, definition of a property test in the dense model includes an artificial dependence of the query model on the value of n: Without utilizing this value it is not possible to make any samples. The results and observations in [15, Section 4] emphasize the artifice of this particular dependence, and lead to an alternative definition of a test avoiding it:

Definition 2.4 (Alternative to Definition 2.3)**.** A *property test* for a graph property Π graphs is a probabilistic oracle machine which is given the values (n, ε), as well access to a graph G of order n, through an oracle which takes two types of requests: A request to uniformly sample an additional vertex out of the remaining vertices of G, and an edge query within the subgraph induced by the sampled vertices ("is there an edge between the i^{th} and j^{th} sampled vertices?"). The machine makes a sequence of requests to the oracle, and distinguishes with probability at least $2/3$ between the case of G being in Π and the case of G being ε-far from Π. If the test has sampled the entire input graph, additional requests to sample an additional vertex will indicate that there are none left.

Definition 2.3 and Definition 2.4 are not equivalent as computational models in general. But in the context of graph properties — which are closed under isomorphism — they are equivalent. This is established, for all intents and purposes, in [15], albeit not formally stated there. We shall not go into the details in this paper.

2.2 Features of property tests

Definition 2.5. A property test is said to be *one-sided* (or said to have *one-sided error*) if it accepts all graphs in Π with probability 1.

Definition 2.6. A property test is said to be *adaptive* if the queries it makes to the input graph may depend in some way on the results of previous queries. If no query made by the test depends on previous query results, the test is said to be *non-adaptive*.

Definition 2.7. A test for a graph property Π is said to be *canonical* if, for some function $s : \mathbb{N} \to \mathbb{N}$ and some sequence of properties $\left(\Pi^{(i)}\right)_{i=1}^{\infty}$, the test operates as follows: on input n and oracle access to an n-vertex graph G, the test samples uniformly a set of $s(n, \varepsilon)$ distinct vertices of G, queries the entire corresponding induced subgraph and accepts if and only if this subgraph is in $\Pi^{(n)}$. If the graph has fewer than $s(n, \varepsilon)$ vertices, the test queries the entire graph and accepts if it is in Π.

Theorem ([15, Theorem 2])**.** *If a graph property has a test with queries involving at most $s(\varepsilon)$ vertices, independently of the size of the input graph, then it has a canonical test with queried subgraph order at most $9s(\varepsilon)$. If the original test is one sided, this canonical test's queried subgraph order is $s(\varepsilon)$ and it is also one-sided.*

Note. The theorem is not phrased in terms of the number of sampled vertices, but this is evident from the proof of theorem: The original test is repeated 9 times and the majority-vote is used, to amplify the probability of success from $1/3$ to $1/6$; see also [16, Page 2, Footnote 1]. If one wishes the canonical test to succeed with higher probability, this can be achieved by repeating the original pre-canonized test additional times (and using a majority vote) before applying canonization; the penalty is a constant-factor increase in the final order of the queried subgraph.

A canonical test, which accepts a graph G when the queried subgraph on its sampled vertices is G', is said to *accept G by sample G'*.

Definition 2.8 (as appearing in [16])**.** A graph property test is said to be *natural* if its query complexity is independent of the size of the tested graphs, and on input (n, ε) and oracle access to a graph of order n, the test's output is based solely on the sequence of oracle answers it receives [not adding "and not on n"; while it makes things extra clear, it is implicit, plus, that's how the definition appears originally.] (while possibly using more random bits, provided that their number is also independent of n).

If our graph property tests are as defined traditionally (Definition 2.3), the above definition of a natural test is flawed, and no test which makes any queries can be natural: A test cannot make $q(\varepsilon)$ queries to an input graph with less than $\sqrt{q(\varepsilon)}$ vertices (this point is also mentioned in [6]). Instead of amending the definition of naturality to avoid this semantic issue, it seems more reasonable to use the alternative definition for the dense graph model, Definition 2.4, in which the artificial dependence on n is removed. In this case, Definition 2.8 is valid: If the test attempts to sample too many vertices, the oracle indicates its failure to do so and the test proceeds accordingly. In fact, in this paper we will assume implicitly that whenever a test attempts to sample more vertices than the input graph has, it proceeds to query the entire graph and accept it deterministically if it satisfies the property being tested.

In this paper we will be dealing mostly with tests which combine both the above features, or rather, we will focus on making canonical tests natural as well. In the context of a canonical test, the naturalness means that the 'internal' property, the one for which the sampled subgraph is checked for, does not depend on the size order of the input graph. This observation leads us to use naturality to define several 'levels' of canonicality for a property test:

Definition 2.9. Consider a canonical test for graph property Π, with $\left(\Pi^{(i)}\right)_{i=1}^{\infty}$ being the sequence of properties the satisfaction of which the test checks for its sampled order-s subgraph. The test is said to be

perfectly canonical	when $\Pi^{(i)} = \Pi$	The test does nothing but ensure that a small random subgraph satisfies the same property the larger input graph is being tested for.
strongly canonical	when $\Pi^{(i)} = \Pi'$	The test ensures a small sampled subgraph satisfies some fixed property, the same for any order of the input graph, but not necessarily Π itself.
weakly canonical	for any $\left(\Pi^{(i)}\right)_{i=1}^{\infty}$	It may be the case that $\Pi^{(i)}$ is different for different values of i.

Notes.

- Indeed, a test is strongly canonical if and only if it is both canonical and natural.
- In [6], the term *oblivious* is used for what we have defined as a strongly canonical test.
- There is only one perfectly canonical test for any queried subgraph order; of course, for many properties this will not constitute a test, as it does not distinguish satisfying graphs from far graphs with sufficient probability.

2.3 Features of graph properties

In this subsection we define strict and approximate notions of graph properties being hereditary and inflatable.

Definition 2.10. A graph $G' = (V', E')$ is a (balanced) *blowup* of a graph $G = (V, E)$ if V' can be partitioned into $|V|$ clusters of vertices, all of the same size up to a difference of at most 1, each corresponding to a vertex in V, where the edges in E' between these clusters correspond to the edges of E. In other words, if $(u, v) \in E$ then the bipartite graph between the clusters corresponding to u and v is complete, and if $(u, v) \notin E$ then this bipartite graph is empty. There are no edges inside each partition set.

Definition 2.11. If $G' = (V', E')$ is a blowup of G, and the clusters in V' (corresponding to the vertices of G) all have exactly the same size (and, in particular, $|V|$ divides $|V'|$), then G' is said to be an *exactly-balanced blowup*.

We will sometimes refer to a "random" or a "uniformly sampled" blowup of a graph from order n to some order n'; this will mean that the $n' \pmod{n}$ vertices which have the larger clusters in the blowup (clusters of size $\lceil n'/n \rceil$ rather than $\lfloor n'/n \rfloor$) are chosen at random.

Lemma 2.12. *Let $G \neq H$ be graphs of order n, let $n' > n$, and let $\phi : V(G) \to V(H)$ be a bijection achieving $\mathrm{dist}\big(G, H\big)$, i.e. exhibiting $\mathrm{dist}\big(G, H\big) \cdot \binom{n}{2}$ discrepancies. If one uniformly samples a blowup G' of G to order n', and applies the same blowup to H — in the sense that for every $v \in G$, the size of v's cluster in G' is the same as the size of $\phi(v)$'s cluster in the blowup H' of H — then the expected distance between between the two blowups is strictly lower than $\mathrm{dist}\big(G, H\big)$.*

Proof. We show that the expected number of discrepancies under a bijection mapping each vertex v's cluster to a vertex in the cluster of $\phi(v)$ is less than $\mathrm{dist}\big(G, H\big)\binom{n'}{2}$, implying the claim. By the linearity of expectation, it suffices to show that for every pair of vertices u, v

which exhibit a discrepancy under ϕ before the blowup, the expected number of discrepancies of the two corresponding clusters in G' and H' is under $(n'/n)^2$.

Now, let $k = n' \pmod{n}$ and $m = \lfloor n'/n \rfloor$. The number of discrepancies due to $\{u, v\}$ is the product of the sizes of u and v's clusters (denote their sizes $\mathrm{cs}(u)$, $\mathrm{cs}(v)$). Each of these clusters has size either m or $m + 1$; thus

$$\begin{aligned}
\mathbf{Ex}[\mathrm{cs}(u) \cdot \mathrm{cs}(v)] &= 1 \cdot (m \cdot m) + \mathbf{Pr}[\mathrm{cs}(u) = m + 1] \cdot (1 \cdot m) \\
&\quad + \mathbf{Pr}[\mathrm{cs}(v) = m + 1] \cdot (m \cdot 1) \\
&\quad + \mathbf{Pr}[\mathrm{cs}(u) = \mathrm{cs}(v) = m + 1] \cdot (1 \cdot 1) \\
&= m^2 + 2 \cdot m \cdot \frac{k}{n} + \mathbf{Pr}[\mathrm{cs}(u) = \mathrm{cs}(v) = m + 1] \\
&= m^2 + 2 \cdot m \cdot \frac{k}{n} + \left(\frac{k}{n} \cdot \frac{k-1}{n-1}\right) < \left(m + \frac{k}{n}\right)^2 = \left(\frac{n'}{n}\right)^2
\end{aligned}$$

This completes the proof. □

Incidentally, Pikhurko has shown in [18, Lemma 14] that the distance between blowups can't be very far below the distance between the original graphs: $\mathrm{dist}(G', H') \geq \frac{1}{3} \mathrm{dist}(G, H)$, for exactly-balanced blowups; this non-trivial direction of the distance bound, however, is not used in our paper.

Definition 2.13. A property Π is said to be *inflatable* if it is closed under blowups, i.e. if G satisfies Π, then so does any blowup of G.

Definition 2.14. A graph property Π is said to be (s, δ)-*inflatable* if for any graph G satisfying Π, of order at least s, all blowups of G are δ-close to satisfying Π. A property Π is said to be (s, δ)-*inflatable on the average* if for any graph G satisfying Π, of order at least s, the expected distance from Π of blowups of G to any fixed order (a uniform sampling out of all possible blowups to that order) is less than δ.

As noted above, blowups do not affect graph distances overmuch. This implies that taking a blowup cannot drive you too far away from an inflatable property:

Proposition 2.15. *Let property Π be (s, δ)-inflatable on the average, let G be a graph of order $n \geq s$, and let $n' > n$. The expected distance of a uniformly-sampled blowup of G to order n from Π is less than $\mathrm{dist}(G, \Pi) + \delta$.*

Proof. Let $H \in \Pi$ be a graph of the same order as G such that $\mathrm{dist}(G, \Pi) = \mathrm{dist}(G, H)$. Let G' and H' be corresponding random blowups of G and H respectively, as defined in the statement of the Lemma 2.12. The lemma gives $\mathbf{Ex}_{G'}[\mathrm{dist}(G', H')] < \mathrm{dist}(G, H)$; also, since Π is (s, δ)-inflatable on the average, and since H if of order at least s, and since H' is a also random blowup, its own expected distance from Π is less than δ. We can now use the triangle inequality to conclude that:

$$\begin{aligned}
\mathop{\mathbf{Ex}}_{G'}[\mathrm{dist}(G', \Pi)] &\leq \mathop{\mathbf{Ex}}_{G'}[\mathrm{dist}(G', H') + \mathrm{dist}(H', \Pi)] \\
&= \mathop{\mathbf{Ex}}_{G'}[\mathrm{dist}(G', H')] + \mathop{\mathbf{Ex}}_{G'}[\mathrm{dist}(H', \Pi)] \\
&< \mathrm{dist}(G, H) + \delta = \mathrm{dist}(G, \Pi) + \delta
\end{aligned}$$

as claimed. □

Definition 2.16. A graph property is said to be *hereditary* if it is closed under the taking of induced subgraphs. A property is said to be *hereditary down to order* n_0 if it is closed under the taking of induced subgraphs of order no less than n_0.

Observation 2.17. Hereditary properties can be characterized by a (possibly infinite) set of forbidden induced subgraphs — a graph satisfies a hereditary property Π if and only if it has no induced subgraph from the forbidden set $\mathcal{F}_\Pi$.

Definition 2.18. A property Π is said to be (s, δ)*-hereditary* if, for every graph in Π, all of its induced subgraphs of order at least s are δ-close to Π. A property Π is said to be (s, δ)*-hereditary on the average* if, for every graph in Π, the expected distance from Π of a uniformly-sampled subgraph of any fixed order $s' \geq s$ is less than δ.

2.4 Fixed-order subgraph distributions of graphs

Definition 2.19. Given a graph G, consider the graph induced by a uniformly sampled subset of s vertices. We denote the distribution of this induced subgraph by D^s_G, the *order-s subgraph distribution* of G; $\mathrm{D}^s_G(G')$ is the relative frequency of a subgraph G' of order s in G.

Note. In [9], this distribution is called the graph's *q-statistic*.

Our analysis of property tests in this paper focuses on such distributions and their distances.

Definition 2.20. Let $\mathcal{G}^s$ denote all graphs of order s. The *distance between two distributions* D, D' over graphs of order s, denoted $\mathrm{dist}\big(D, D'\big)$, is the variation distance between them, i.e.

$$\mathrm{dist}\big(D, D'\big) = \frac{1}{2} \sum_{G \in \mathcal{G}^s} \big|D(G) - D'(G)\big|$$

The distance between two graphs' order-s subgraph distributions cannot exceed the two graphs' relative distance overmuch:

Lemma 2.21. *If two graphs G, H (of order $n \geq s$) are $\delta\binom{s}{2}^{-1}$-close, then their order-s subgraph distributions are δ-close, i.e.* $\mathrm{dist}\big(\mathrm{D}^s_G, \mathrm{D}^s_H\big) \leq \delta$.

Proof. Let $\phi : V(G) \rightarrow V(H)$ be a bijection achieving the minimum of the number of edge discrepancies. The graphs' being $\delta\binom{s}{2}^{-1}$-close means that there are at most $\delta\binom{s}{2}^{-1} \cdot \binom{n}{2}$ such discrepancies. Now consider a uniformly-sampled set of s vertices in $V(G)$, and the subgraph they induce in G and (through ϕ) in H. Every pair of vertices in the subgraph is uniformly distributed among the pairs of vertices of G or of H, so the probability of having any discrepant edges between these two subgraphs under ϕ is at most δ. When we condition on the sample not containing any vertex pair discrepant under ϕ, the distributions of such an order-s subgraph of G and of H becomes identical; the variation distance between the unconditioned distributions cannot, therefore, exceed δ. □

Another feature of the the order-s subgraph distribution is that it does not change overmuch when taking the blowup of a graph.

Lemma 2.22. *Let $\delta > 0$, let G be a graph of order $n \geq \frac{2}{\delta}\binom{s}{2}$, let G' be a random blowup of G to order $n' > n$, and let $\mathcal{H} \subseteq \mathcal{G}^s$. Then*

$$\left| \mathop{\mathbf{Ex}}_{G'} \left[\mathop{\mathbf{Pr}}_{H \sim \mathrm{D}^s_{G'}} [H \in \mathcal{H}] \right] - \mathop{\mathbf{Pr}}_{H \sim \mathrm{D}^s_G} [H \in \mathcal{H}] \right| < \delta$$

Proof. Let $\tilde{\mathrm{D}}^s_{G'}$ denote the order-s subgraph distribution of G', conditioned on the event that every vertex of the subgraph is in the cluster of a different vertex of G. For any fixed G', we have

$$\left| \mathop{\mathbf{Pr}}_{H\sim \mathrm{D}^s_{G'}}[H \in \mathcal{H}] - \mathop{\mathbf{Pr}}_{H\sim \tilde{\mathrm{D}}^s_{G'}}[H \in \mathcal{H}] \right| \le \mathrm{dist}\big(\mathrm{D}^s_{G'}, \tilde{\mathrm{D}}^s_{G'}\big)$$

This variation distance is bounded by the probability p that multiple vertices in H sampled uniformly from G' are in the same cluster of vertex of G. For a given pair of vertices of H, the probability of their being in the same cluster is at most the relative size of a large cluster, which is bounded by $2/n$; union-bounding over all pairs, we have, irrespective of G',

$$p < \binom{s}{2} \cdot \frac{2}{n} \le \binom{s}{2} \cdot \frac{2}{\frac{2}{\delta}\binom{s}{2}} = \delta$$

The claim can now be established if we show that

$$\mathop{\mathbf{Ex}}_{G'}\left[\mathop{\mathbf{Pr}}_{H\sim \tilde{\mathrm{D}}^s_{G'}}[H \in \mathcal{H}] \right] = \mathop{\mathbf{Pr}}_{H\sim \mathrm{D}^s_{G}}[H \in \mathcal{H}]$$

For this purpose, let us analyze separately the various sets of s vertices in G (corresponding to sets of s clusters in G'): The probability of sampling H in $\mathcal{H}$ is the probability of sampling a set S of s vertices, such that the induced graph $H = H_S$ on these vertices is in $\mathcal{H}$; in G', it is the probability of sampling vertices from the appropriate sets of s clusters. Let $\mathcal{S}_{\mathcal{H}}$ be the family of s-vertex sets S with $H_S \in \mathcal{H}$. Denote by $p_S(G')$ the probability that a set S' of s vertices, each from a different cluster of a G vertex, equals S. Now, by the linearity of expectation,

$$\mathop{\mathbf{Ex}}_{G'}\left[\mathop{\mathbf{Pr}}_{H\sim \tilde{\mathrm{D}}^s_{G'}}[H \in \mathcal{H}] \right] = \mathop{\mathbf{Ex}}_{G'}\left[\sum_{S\in\mathcal{S}_{\mathcal{H}}} p_S(G') \right] = \sum_{S\in\mathcal{S}_{\mathcal{H}}} \mathop{\mathbf{Ex}}_{G'}\big[p_S(G')\big]$$

The expectation $\mathbf{Ex}_{G'}[p_S(G')]$ is the same, by symmetry, for all s-subsets S, as the blowup G' is sampled uniformly. It must therefore be equal to the inverse of the number of sets S, i.e. $\binom{n}{s}^{-1}$. Thus

$$\mathop{\mathbf{Ex}}_{G'}\left[\mathop{\mathbf{Pr}}_{H\sim \tilde{\mathrm{D}}^s_{G'}}[H \in \mathcal{H}] \right] = \sum_{S\in\mathcal{S}_{\mathcal{H}}} \mathop{\mathbf{Ex}}_{G'}\big[p_S(G')\big] = \sum_{S\in\mathcal{S}_{\mathcal{H}}} \binom{n}{s}^{-1} = \mathop{\mathbf{Pr}}_{H\sim \mathrm{D}^s_{G}}[H \in \mathcal{H}]$$

as claimed. □

[following is a proposition we don't use.]

Note that while any single event (or single order-s subgraph or set of s clusters) has the same expected probability, in specific blowups this may very well not be the case, even for $n \gg s$, as one may choose to have, say, the higher-degree vertices have bigger clusters, and the lower-degree vertices in smaller clusters. The following proposition gives a deterministic bound on the distance between the subgraph distributions using both the order of the pre-blowup graph n and the 'imbalance' of the blowup:

Proposition 2.23. *Let G be a graph of order $n \ge s$ and G' a blowup of G to order $n' \ge n$, and let $k = n' \pmod{n}$. Then $\mathrm{dist}\big(\mathrm{D}^s_{G'}, \mathrm{D}^s_{G}\big) < \binom{s}{2} \cdot \frac{1}{n} + s \cdot \frac{\min\{k, n-k\}}{n'} < \binom{s}{2} \cdot \frac{1}{n} + s \cdot \frac{n}{n'}$, and $\mathrm{dist}\big(\mathrm{D}^s_{G'}, \mathrm{D}^s_{G}\big) < \binom{s}{2} \cdot \frac{1}{n}$ if n divides n'.*

Proof. Let us first analyze the case of the blowup G' being exactly-balanced, i.e. $n' = n \cdot k$ for some $k \in \mathbb{N}$. Consider a sample of an s-vertex subgraph of G'. Conditioning on the event of every vertex being sampled from the cluster of a different vertex of G, the distribution of order-s subgraphs of G' is exactly D^s_G. Thus the unconditioned distance $\mathrm{dist}\big(\mathrm{D}^s_{G'}, \mathrm{D}^s_G\big)$ is at most the probability of sampling at least two of the s vertices from the same cluster. Since G' is an exactly-balanced blowup, this probability is less than $1/n$ for a single pair of vertices. Applying a union bound over the $\binom{s}{2}$ pairs of vertices yields $\mathrm{dist}\big(\mathrm{D}^s_{G'}, \mathrm{D}^s_G\big) < \frac{1}{n}\binom{s}{2}$.

In the general case, G' is not necessarily exactly-balanced. However, let us choose one vertex from each of the $n' \pmod n$ larger clusters to form a set U. the subgraph of G' induced by $V(G') \setminus U$ is an exactly-balanced blowup of G; and with probability at least $1 - s \cdot \frac{k}{n'}$, a sample of s vertices from $V(G)$ is in fact sampled from $V(G') \setminus U$ only, conditioning on which event the above distance bound holds. Alternatively, think of an exactly-balanced blowup G'' of G, to order $n' + n - k$. The exactly-balanced distance holds for G'', but when conditioning on the event of no vertices being sampled out of the $n - k$ additional vertices in G'', it has the same order-s subgraph distribution as G'; this event's probability is at least $1 - s \cdot \frac{n-k}{n'}$.

In the general case, therefore, we have $\mathrm{dist}\big(\mathrm{D}^s_{G'}, \mathrm{D}^s_G\big) < \min\big\{\frac{1}{n}\binom{s}{2} + \frac{k}{n'}, \frac{1}{n}\binom{s}{2} + \frac{n-k}{n'}\big\}$ as claimed. $\square$

3 Our results

We first state our our main result in a simplified manner, for motivation and clarity:

Theorem 1. *If a hereditary, inflatable graph property has a test making $q(\varepsilon)$ queries, regardless of the size of the input graph, then it has a strongly canonical test — specifically, a natural test — making $O\big(q(\varepsilon)^4\big)$ queries.*

We will in fact prove a mildly stronger version, with the above being a special case:

Theorem 1 (exact version)**.** *Let Π be a graph property has a test with queries involving at most $s(\varepsilon)$ distinct vertices, regardless of the size of the input graph, and let $s_1 = 12\binom{31s}{2}$. If Π is $\big(s_1, \frac{1}{6}\binom{s_1}{2}^{-1}\big)$-hereditary on the average and $\big(s_1, {s_1}^{-1}\big)$-inflatable on the average, then it has a strongly canonical test whose queried subgraph order is $s_1 = O\big(s(\varepsilon)^2\big)$.*

Note. This theorem should also hold also for properties with weaker inflatability — a higher threshold value than stated above for ε-inflatability on the average — with some modifications of our proof, and with a worse dependence of the queried subgraph order on s.

We also prove a weak inverse of Theorem 1:

Theorem 2. *If a graph property Π has a natural (not necessarily canonical) test with queries involving $s(\varepsilon)$ distinct vertices, then for every $\varepsilon' > \varepsilon$, Π is $\big(s_h, \varepsilon'\big)$-hereditary on the average and $\big(s_i, \varepsilon'\big)$-inflatable on the average, for $s_h = O\big(s \cdot \log(\frac{1}{\varepsilon' - \varepsilon})\big)$ and $s_i = O\big(s^2 \cdot (\varepsilon' - \varepsilon)^{-1} \log(\frac{1}{\varepsilon' - \varepsilon})\big)$ respectively.*

Let us now recall the proposition from Goldreich and Trevisan discussed in the introduction: [I'm not numbering this one. I'm only numbering propositions and theorems which we're proving in this paper; this one already has a number - D.2 ...]

Proposition ([15, proposition D.2], corrected as per [16])**.** *Let Π be a hereditary graph property, with a <u>natural</u> test making $q(\varepsilon)$ queries. Then Π has a perfectly canonical test with queried subgraph order $O\big(q(\varepsilon)\big)$.*

Originally, this proposition was stated without requiring that the test be natural (merely that the number of queries be independent of the order of the input graph).

Combining this corrected, qualified version above with Theorem 1, one obtains:

Corollary 3. *If a hereditary inflatable property* Π *has a test making* $q(\varepsilon)$ *queries, it has a perfectly canonical test with queried subgraph order* $\text{poly}(q(\varepsilon))$*.*

The converse of this corollary, useful for proving lower bounds, is that if a hereditary inflatable property has no perfectly canonical test with queried subgraph order $\text{poly}(q(\varepsilon))$, then it has no test whatsoever whose number of queries is polynomial in ε (natural or otherwise, with one-sided or two-sided error). For triangle-freeness specifically, combining Corollary 3 with [2, Lemma 3.1], and the fact that the property of triangle-freeness is inflatable, one obtains an alternative, 'generic' proof of the following:

Corollary 4 (first proven as [5, Theorem 1])**.** *The query complexity of any* ε*-test — natural or otherwise, with one-sided or two-sided error — for the property of being triangle-free makes is at least* $(c/\varepsilon)^{c\cdot\log(c/\varepsilon)}$*, for some global constant* c*.*

As mentioned in the introduction, the proof in [5] uses a construction specific to the details of the $(c/\varepsilon)^{c\cdot\log(c/\varepsilon)}$ one-sided lower bound. The construction we will be utilizing in the proof of Theorem 1 applies to any test for triangle-freeness, so a proposition similar to the above corollary would hold for any possible lower bound on the query complexity of testing triangle-freeness. It will similarly hold for the property of being free of any single non-bipartite graph which is not a blowup of a smaller graph.

Returning to [15, proposition D.2], while for hereditary inflatable properties we have established it holds with a power-of-four penalty on the number of queries, for properties with one-sided tests it can be shown to hold as stated:

Proposition 3.1. *If a hereditary inflatable property* Π *has a one-sided (not necessarily natural) test making* $q(\varepsilon)$ *queries, then* Π *has a perfectly canonical test with queried subgraph order at most* $2q$*.*

Finally, putting inflatability in the context of proximity-oblivious testing, we prove the following partial characterization:

Proposition 3.2. *Let* Π *be an inflatable hereditary graph property.* Π *has a constant-query proximity-oblivious test if and only if there exists a constant* s *such that, for* $n \geq s$, Π_n *consists exactly of those graphs of order* n*, which are free of order-*s *graphs outside of* Π_s*.*

4 Naturalizing tests — proof of Theorem 1

Let Π be a property meeting the conditions of Theorem 1. As Π has a test with queries involving at most $s(\varepsilon)$ vertices (independently of n), by [15, Theorem 2] it has a canonical test, querying a uniformly sampled subgraph of order at most $9s$, in its entirety. As discussed in Subsection 2.2, we may assume that the canonical test's probability of error is at most $\frac{1}{36}$ rather than $\frac{1}{3}$, at the cost of increasing the queried subgraph order to $s_0 = 31s$.

One may think of the existence of such a canonical test as meaning that the membership of a graph in Π is essentially determined by its distribution of (induced) subgraphs of order s_0. This being the case, let us consider a (canonical) 'meta-test' for Π, which estimates whether the subgraph distribution leads to acceptance: This meta-test is listed as Algorithm 1.

Note. The order s_1 of the larger subgraph used for this estimate is chosen so as to ensure the stability of the distribution under blowups — a consideration which will become relevant later in this section. On the other hand, s_1 is not high enough to properly *estimate the distribution*, i.e. estimate the frequency of specific order-s_0 subgraphs (there are $\exp(\Omega(s_0{}^2))$ of them) in G.

Algorithm 1 A Meta-Test for Π

1: Uniformly query a subgraph G_{sample} of order $s_1 = 12\binom{s_0}{2} = 12\binom{31s(\varepsilon)}{2}$.
2: If at least a $\frac{1}{6}$-fraction of the order-s_0 subgraphs G' of G_{sample} are such that the (canonical) s_0-test accepts G by sample G', accept. Otherwise reject.

Lemma 4.1. *Algorithm 1 is a valid test for property* Π*, with probability of failure at most* $1/6$*.*

Proof. Suppose the input graph G either satisfies Π or is ε-far from satisfying Π. Let G' be one of the $\binom{s_1}{s_0}$ order-s_0 subgraphs of G_{sample}. Let $X_{G'}$ be the indicator for the s_0-test erring (that is, rejecting G in case G satisfies Π, or accepting G in case G is far from Π) by sample G'. Every order-s_0 subgraph of G_{sample} is in fact uniformly sampled from the input graph, thus $\mathbf{Ex}[X_{G'}]$ is the probability of the s_0-test erring — at most $\frac{1}{36}$. The expected fraction of order-s subgraphs of G_{sample} by which the s_0-test errs is therefore also at most $\frac{1}{36}$. Considering the meta-test's behavior again, it only errs if at least a $\frac{1}{6}$-fraction of the subgraphs of G_{sample} cause the s_0-test to err. by Markov's inequality the probability of this occurring is at most $\frac{1}{36}/\frac{1}{6} = \frac{1}{6}$. □

Let us now modify Algorithm 1 to reject samples which are themselves not in the property at order s_1; the result is listed as Algorithm 2.

Algorithm 2 Modified Meta-Test for Π

1: Uniformly query a subgraph G_{sample} of order $s_1 = 12\binom{s_0}{2} = 12\binom{31s(\varepsilon)}{2}$.
2: If G_{sample} is not in Π, reject.
3: If at least a $\frac{1}{6}$-fraction of the order-s_0 subgraphs G' of G_{sample} are such that the s_0-test accepts G by sample G', then accept. Otherwise reject.

Lemma 4.2. *Algorithm 2 is a valid test for property* Π*.*

Proof. The additional check only increases the probability of rejection of any input graph, so it does not adversely affect the soundness of the modified test (that is, a graph ε-far from Π is still rejected by Algorithm 2 with probability at least $\frac{5}{6} \geq \frac{2}{3}$).

As for the modified test's completeness, we recall that Π is $\left(s_1, \frac{1}{6}\binom{s_1}{2}^{-1}\right)$-hereditary on the average. This implies that, for an input graph in Π, the average distance of subgraphs of order s_1 from Π is $\frac{1}{6}\binom{s_1}{2}^{-1}$; as each order-$s_1$ subgraph not in Π is at least $\binom{s_1}{2}^{-1}$-far from Π, the fraction of order-s_1 subgraphs of G which aren't in Π is at most $\frac{1}{6}$. Regardless of these, at most a $\frac{1}{6}$-fraction of the order-s_1 subgraphs of a satisfying graph cause Algorithm 1 to reject. Union bounding over these two sets of subgraphs causing rejection we find that the probability of the modified meta-test rejecting a graph in Π is less than $2 \cdot \frac{1}{6} = \frac{1}{3}$. □

If Algorithm 2 were somehow also natural, this would complete the proof of Theorem 1, as it otherwise meets its requirements. Interestingly enough, our modification has indeed made it natural, as implied by the following final element of our proof:

Lemma 4.3. *Let* H *be a graph of order* s_1 *by which sample Algorithm 2 accepts for at least some input graph order* n*. Algorithm 2 cannot reject for any input graph order* $n' \geq s_1$ *by sample* H*.*

Proof. Assume on the contrary that Algorithm 2 rejects by sample H for some $n' \geq s_1$. We first note that Algorithm 2 does not reject by H at order n' on account of H not being in Π (as samples which aren't in Π are rejected at all input orders). We will show that this invariably implies that the original test is unsound.

Let $\Pi'_{n'}$ denote the set of order-s_0 subgraphs by which sample the s_0-test accepts an input graph G at order n'. Our assumption is that the probability of the s_0-test accepting a subgraph of H is less than $\frac{1}{6}$, or in terms of the subgraph distribution, $\mathbf{Pr}_{H_\mathrm{s} \sim \mathrm{D}_H^{s_0}}\left[\Pi'_{n'}\right] < \frac{1}{6}$. [I don't like the notation H_s; but H_sample would be too long; and I can't really start using a capital letter other than G and H.]

Now, consider a random blowup H' of H to order n'. Π is $\left(s_1, \frac{1}{12}\binom{s_0}{2}^{-1}\right)$-inflatable on the average, and H is in Π, so

$$\mathop{\mathbf{Ex}}_{G'}\left[\mathrm{dist}(H', \Pi)\right] < \frac{1}{12}\binom{s_0}{2}^{-1}$$

and by Markov's inequality,

$$\mathop{\mathbf{Pr}}_{H'}\left[\mathrm{dist}(H', \Pi) \geq \frac{1}{6}\binom{s_0}{2}^{-1}\right] < \frac{1}{2}$$

Also, let $\delta = \frac{1}{6}$. Since $s_1 \geq \frac{2}{\delta}\binom{s_0}{2}$, we may apply Lemma 2.22 (substituting H and H' for G and G', s_0 for s, s_1 for n) for the event of the s_0-test accepting at order n':

$$\begin{aligned}\mathop{\mathbf{Ex}}_{H'}\left[\mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H'}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right]\right] &\leq \mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] + \left|\mathop{\mathbf{Ex}}_{H'}\left[\mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H'}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right]\right] - \mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right]\right| \\ &< \mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] + \delta < \frac{1}{6} + \frac{1}{6} = \frac{1}{3}\end{aligned}$$

and again by Markov's inequality

$$\mathop{\mathbf{Pr}}_{H'}\left[\mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H'}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] \geq \frac{2}{3}\right] < \frac{1}{2}$$

Combining these two facts, we conclude that with positive probability, H' is a graph which is both very close to Π and is accepted by the s_0-test with probability at most $\frac{2}{3}$.

Now, let $\widetilde{H'}$ be a graph in Π at distance at most $\frac{1}{6}\binom{s_0}{2}^{-1}$ from H'. By Lemma 2.21, these two graphs' order-s_0 subgraph distributions are $\frac{1}{6}$-close, implying that

$$\left|\mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{\widetilde{H'}}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] - \mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H'}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right]\right| < \frac{1}{6}$$

We now use the triangle inequality to bound the probability of the s_0-test accepting $\widetilde{H'}$:

$$\begin{aligned}\mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{\widetilde{H'}}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] &\leq \mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H'}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] + \left|\mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{H'}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right] - \mathop{\mathbf{Pr}}_{H_\mathrm{s} \sim \mathrm{D}_{\widetilde{H'}}^{s_0}}\left[H_\mathrm{s} \in \Pi'_{n'}\right]\right| \\ &< \frac{2}{3} + \frac{1}{6} = \frac{5}{6}\end{aligned}$$

This contradicts the original test's probability of error — it must accept $\widetilde{H'}$, a graph in Π, with probability at least $1 - \frac{1}{36} > \frac{5}{6}$. It can therefore not be the case that Algorithm 2 rejects H at order n'. □

Proof of Theorem 1. Given a property Π satisfying the conditions, we have devised Algorithm 2: This is a canonical test for Π, with queried subgraph order $s_1 = 12\binom{31s}{2}$; by Lemma 4.3, it accepts and rejects the same set of queried subgraphs for all graph orders $n \geq s_1$ — that is, it is a natural test. □

5 Natural testability and inflatability — further discussion

5.1 One-sidedness and Proposition 3.1

Observation 5.1. If a hereditary property has a strongly canonical test, this test must be one-sided.

Proof. If the test for the hereditary property Π (deterministically) rejects any sampled subgraph G' of a graph $G \in \Pi$, the test also rejects G' when it is the entire graph. But when G' is the entire graph, it will always be the sampled subgraph, i.e. the test rejects G' with probability 1. G' can therefore not be in Π — a contradiction to Π being hereditary. □

The implication in Observation 5.1 can be reversed, in a way — weak approximate heredity as a consequence of one-sided testability:

Lemma 5.2. *If a property Π has a one-sided strongly canonical test with queried subgraph order $s(\varepsilon)$ for some ε, then Π is $\big(s(\varepsilon), \varepsilon\big)$-hereditary.*

Proof. Let $G \in \Pi_n$ for $n \geq s(\varepsilon)$, and let G' be a subgraph of G of order at least $s(\varepsilon)$. If G' is ε-far from Π, then it must have an order-s subgraph G'' by which sample the test rejects G'. But the test also rejects G by sample G'', in contradiction to its one-sidedness. □

Note. This lemma is somewhat similar to the second direction of [6, Theorem 2], in which the existence of a one-sided natural test is shown to imply 'semi-heredity'.

One would hope to somehow get rid of the dependence on ε and find conditions under which the property is hereditary, at least down to some n_0; this becomes possible if the test is proximity-oblivious, but note that if a property Π has a natural proximity-oblivious test, then Π is simply the property of being free of subgraphs by which this test rejects (at least for $n \geq s$; see discussion in Subsection 5.3, and specifically Proposition 3.2).

In the proof of Lemma 5.2, we used the one-sidedness of the test to obtain deterministic approximate heredity; Subsection 5.2 below deals with the general, two-sided case, and establishes approximate heredity only on the average. Deterministic approximate heredity may indeed require the test to be one-sided. For example, the property Π_{half}, containing those graphs with at most $\frac{1}{2}\binom{n}{2}$ edges, is $\big(O(\frac{1}{\delta}), \delta\big)$-hereditary on the average, has a two-sided natural test (in fact, its query complexity can be shown to be $O\big(1/\varepsilon\big)$), but it is not $\big(s, \frac{1}{2} - \delta\big)$-hereditary for any s and $\delta > 0$ (there are satisfying graphs with arbitrarily large complete subgraphs).

Returning again to the direction of Theorem 1, let us follow an alternate line of argumentation than the one used to prove Theorem 1, this time for the case of one-sided tests.

Lemma 5.3. *Let Π be an inflatable property. A one-sided canonical test for Π can only reject an input graph when it samples a subgraph which is not itself in Π.*

Proof. Suppose for some input graph G of order n the test samples a subgraph $G' \in \Pi$. Since Π is inflatable, the exists a blowup G'' of G' to order n such that $G'' \in \Pi$. Now, G' is an induced subgraph of G'', so it is possible for the test to sample G' when G'' is the input graph. Since the test is one-sided, it can not, therefore, reject an input graph of order n with G' as the sample. □

Proof of Proposition 3.1. By [15, Theorem 2], Π has a canonical one-sided test with queried subgraph order $s(\varepsilon) \leq 2q(\varepsilon)$, which is also one-sided. By Lemma 5.3, this test only rejects sampled subgraphs which aren't themselves in Π. Now suppose we modify the test so as to reject all sampled subgraphs not in Π. As we are only rejecting additional subgraphs, the test's soundness can only improve. As for its completeness, we note that since Π is hereditary, no graph in Π has any subgraphs outside of Π, so the test still accepts graphs in Π with probability 1. The resulting test is indeed perfectly canonical. □

5.2 Reversing the direction of Theorem 1

Lemma 5.4. *If a property Π has a strong canonical test with queried subgraph order $s(\varepsilon)$, with probability of error $\delta \leq \frac{1}{3}$, then Π is $\left(\frac{2}{\delta}\binom{s}{2}, \varepsilon + 3\delta\right)$-inflatable on the average.*

Proof. Let G be a graph in Π_n for $n \geq \frac{2}{\delta}\binom{s}{2}$ and let G' be a uniformly sampled blowup of G to some higher order. Let Π' be as in Definition 2.9 (the set of order-s subgraphs by which sample the test accepts an input graph). By Lemma 2.22, we have

$$\left| \underset{G'}{\mathbf{Ex}}\left[\underset{H \sim \mathrm{D}^s_{G'}}{\mathbf{Pr}}\left[H \notin \Pi'\right] \right] - \underset{H \sim \mathrm{D}^s_{G}}{\mathbf{Pr}}\left[H \notin \Pi'\right] \right| < \delta$$

so $\mathbf{Ex}_{G'}\left[\mathbf{Pr}_{H \sim \mathrm{D}^s_{G'}}[H \notin \Pi']\right] < 2\delta$. By Markov's inequality

$$\underset{G'}{\mathbf{Pr}}\left[\underset{H \sim \mathrm{D}^s_{G'}}{\mathbf{Pr}}\left[H \notin \Pi'\right] > 1 - \delta \right] \leq \frac{2\delta}{1-\delta} \leq 3\delta$$

Now, if G' is rejected by the test with probability at most $1-\delta$, it cannot be ε-far from Π; if it is rejected with higher probability, we can't make any assumptions regarding its distance. Thus

$$\mathbf{Ex}\left[\mathrm{dist}(G', \Pi)\right] < \underset{G'}{\mathbf{Pr}}\left[\underset{H \sim \mathrm{D}^s_{G'}}{\mathbf{Pr}}\left[H \notin \Pi'\right] \leq 1 - \delta \right] \cdot \varepsilon + \underset{G'}{\mathbf{Pr}}\left[\underset{H \sim \mathrm{D}^s_{G'}}{\mathbf{Pr}}\left[H \notin \Pi'\right] > 1 - \delta \right] \cdot 1 \leq \varepsilon + 3\delta$$

□

Lemma 5.5. *If a property Π has a strongly canonical test, with queried subgraph order $s(\varepsilon)$, with probability of error $\delta \leq \frac{1}{3}$, then Π is $\left(s, \varepsilon + \frac{3}{2}\delta\right)$-hereditary on the average.*

Proof. Let G be a graph in Π of order at least s, let G' a uniformly-sampled subgraph of G of order $s' \geq s$, and let $p_{G'}$ denote the probability of the test rejecting with G' rather than G as its input graph. The expectation of $p_{G'}$ is exactly δ, the probability of the test rejecting G — as the process of sampling an order-s' subgraph, then sampling an order-s subgraph out of it, is the same as just sampling an order-s subgraph of G. We can apply Markov's inequality and bound the probability of $p_{G'}$ being too high: $\mathbf{Pr}_{G'}[p_{G'} \geq 1-\delta] \leq \frac{\delta}{1-\delta}$. Since the test is sound, we know that if $p_{G'}$ is lower than $1-\delta$, G' cannot be ε-far from Π; if $p_{G'}$ is higher, we don't assume anything about G''s distance from Π. Thus

$$\begin{aligned}\underset{G'}{\mathbf{Ex}}\left[\mathrm{dist}(G', \Pi)\right] &\leq \underset{G'}{\mathbf{Pr}}[p_{G'} < 1-\delta] \cdot \varepsilon + \underset{G'}{\mathbf{Pr}}[p_{G'} \geq 1-\delta] \cdot 1 \\ &\leq 1 \cdot \varepsilon + \frac{\delta}{1-\delta} \cdot 1 = \varepsilon + \frac{\delta}{1-\delta} \leq \varepsilon + \frac{3}{2}\delta\end{aligned}$$

□

Proof of Theorem 2. Let $\delta = \frac{1}{3}(\varepsilon' - \varepsilon)$. Our first step is the same as in the proof of Theorem 1 — pre-amplifying the probability of success of the natural test and canonicalizing it. Our modified test remains natural (thus being strongly canonical), with probability of failure at most δ, and its queried subgraph size is $s_{\mathrm{h}} = O\big(s \cdot \log\big(\delta^{-1}\big)\big)$, as per the discussion of canonization in Subsection 2.2. Now, by Lemma 5.5, Π is $\big(s_{\mathrm{h}}, \varepsilon + \frac{3}{2}\delta\big)$-hereditary on the average, and by Lemma 5.4, Π is $\big(\frac{2}{\delta}\binom{s_{\mathrm{h}}}{2}, \varepsilon + 3\delta\big)$-inflatable on the average. This meets the claim. □

5.3 Natural testability and proximity-oblivious testing

In most works regarding property testing, tests are devised based on a foreknowledge of the proximity parameter ε: Either the test is given ε as input, or ε is fixed globally. Goldreich and Ron explore an alternative approach in [13], studying one-sided tests which act irrespectively of ε. Characterizing the probability of rejection ρ as a function of an input graph's distance from the property, one then obtains an ε-test by invoking the *proximity-oblivious* test $\Theta\big(1/\rho(\varepsilon)\big)$ times. A proximity oblivious test has query complexity which may depend on n, but not on ε.

Lemma 5.6. *If a hereditary, inflatable graph property has a proximity-oblivious test with making c queries (using s sampled vertices), then it has a perfectly canonical proximity-oblivious test with queried subgraph order s (making at most $\binom{c}{2}$ queries).*

The proof of this lemma is exactly the proof of Proposition 3.1, which applies regardless of whether the test depends on the proximity parameter ε or not.

The general results of [13] regarding the dense graph model include a characterization of the properties admitting a constant-query proximity-oblivious test:

Theorem ([13, Theorem 4.7])**.** *A property Π has a constant-query proximity-oblivious test if and only if there exists a constant c and a finite sequence $\overline{\mathcal{F}} = (\mathcal{F}_n)_{n\in\mathbb{N}}$ of sets of graphs, such that each $\mathcal{F}_n$ contains graphs of size at most c, and Π_n is the set of order-n $\mathcal{F}_n$-free graphs.*

When limiting our focus to properties which we know to be naturally testable, we can tighten the characterization:

Proof of Proposition 3.2. If Π is the property of being $\mathcal{F}$-free for $\mathcal{F} = (\Pi_s)^c$, Π is proximity-oblivious testable with a constant number of queries: By [3], either the graph is close to being $\mathcal{F}$-free, or it has $\delta(\varepsilon) \cdot n^s$ induced copies of this forbidden subgraph (with δ being a double-tower function of $(1/\varepsilon)$, as this fact is established using a version of Szemerédi's regularity lemma). In this direction, the argument is the same as for the general characterization theorem of proximity-oblivious-testable properties [13, Theorem 4.7].

The other direction follows from Lemma 5.6: The existence of a proximity-oblivious test implies the existence of a perfectly canonical test querying a subgraph of order s and rejecting if it isn't in Π_s. This test accepts, with probability 1, exactly those graphs which are free of induced subgraphs outside Π_s; as it is one-sided, this implies that Π, at order s and above, is the set of $(\Pi_s)^c$-free graphs. □

6 Open questions

Naturalization without canonization This paper can in fact be said not to be concerned with natural tests per se, but rather with a further development of the concept of canonical tests, with naturalization being a 'strengthening' of the degree of canonicality. What can be said regarding the naturality of non-canonical, or even non-adaptive, tests? Can such tests be made natural without incurring the penalty for making them non-adaptive and canonical?

Testing a large graph by testing small subgraphs Goldreich and Trevisan posed in [16] the question of whether any test for a hereditary properties can be replaced with merely ensuring that a random *small* induced subgraph (not much larger than the subgraph queried by the original test) has the property — as was originally claimed in [15, Proposition D.2]. We've shown that being hereditary and inflatable, or one-sided-testable, is a sufficient condition for this to hold. Are these conditions, or similar ones, also necessary?

The benefit of non-natural testing Some testable properties have a non-constant-factor gap in query complexity between their adaptive and non-adaptive tests; Such a gap may also exist between natural and n-dependent tests. As with adaptivity, it will be bounded by the penalty of naturalizing the test when at all possible. Can one find specific properties exhibiting such a gap, or 'non-contrived' properties for which there is no gap (similarly to Goldreich and Ron's [14])?

A more appropriate notion of inflatability Our choices for the definition of a blowup and of (perfect) inflatability are somewhat arbitrary. For example, the property of being the empty graph is inflatable, but the property of being the complete graph is not — since the clusters in a blowup are empty rather than, say, supporting a clique. Also, the property of being H-free, when H itself is a blowup of a smaller graph, is not inflatable, and not even (s, δ)-inflatable, for any s. However, these properties are inflatable on the average (even though for the case of subgraph freeness the known threshold s is exceedingly high). Can one devise a more appropriate, perhaps more relaxed notion of inflatability, which covers such properties as well, while still allowing for naturalization with the same polynomial penalty as in Theorem 1? We are uncertain whether one could devise a useful notion of graph blowup under which all such properties would be considered 'perfectly' inflatable. Of course, this is not much of an issue with regard to (s, δ)-inflatability, as at high orders the edges within the clusters are immaterial to the distance.

References

[1] Noga Alon. Private communication, 1999.

[2] Noga Alon. Testing subgraphs in large graphs. *Random Structures and Algorithms*, 21(3-4):359–370, 2002.

[3] Noga Alon, Eldar Fischer, Michael Krivelevich, and Mario Szegedy. Efficient testing of large graphs. *Combinatorica*, 20:451–476, 2000.

[4] Noga Alon, Eldar Fischer, Ilan Newman, and Asaf Shapira. A combinatorial characterization of the testable graph properties: It's all about regularity. *SIAM Journal on Computing*, 39(1):143–167, 2009. An earlier version appeared in the proceedings of the 38^{th} STOC, 2006.

[5] Noga Alon and Asaf Shapira. A characterization of easily testable induced subgraphs. *Combinatorics, Probability and Computing*, 15(6):791–805, 2006. An earlier version appeared in the proceedings of the 15^{th} SODA, 2004.

[6] Noga Alon and Asaf Shapira. A characterization of the (natural) graph properties testable with one-sided error. *SIAM Journal on Computing*, 37(6):1703–1727, 2008. An earlier version appeared in the proceedings of the 46^{th} FOCS, 2005.

[7] Christian Borgs, Jennifer Chayes, László Lovász, Vera T. Sós, Balázs Szegedy, and Katalin Vesztergombi. Graph limits and parameter testing. In *Proceedings of the 38th STOC*, pages 261–270, New York, NY, USA, 2006. ACM Press.

[8] Eldar Fischer. The art of uninformed decisions: A primer to property testing. In G. Paun, G. Rozenberg, and A. Salomaa, editors, *Current Trends in Theoretical Computer Science: The Challenge of the New Century*, volume 1, pages 229–264. World Scientific Publishing, 2004.

[9] Eldar Fischer and Ilan Newman. Testing versus estimation of graph properties. *SIAM Journal on Computing*, 37(2):482–501, 2007.

[10] Jacob Fox. A new proof of the graph removal lemma. *Annals of Mathematics*, to appear. available at the following URL: http://math.mit.edu/~fox/paper-removal.pdf.

[11] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.

[12] Oded Goldreich, Michael Krivelevich, Ilan Newman, and Eyal Rozenberg. Hierarchy theorems for property testing. In *Proceedings of RANDOM 2009*, pages 504–519, Berlin, Heidelberg, 2009. Springer-Verlag.

[13] Oded Goldreich and Dan Ron. On proximity oblivious testing. In *Proceedings of the 41st STOC*, pages 141–150, New York, NY, USA, 2009. ACM.

[14] Oded Goldreich and Dana Ron. Algorithmic aspects of property testing in the dense graphs model. In Oded Goldreich, editor, *Property Testing*, volume 6390 of *Lecture Notes in Computer Science*, pages 295–305. Springer, 2010.

[15] Oded Goldreich and Luca Trevisan. Three theorems regarding testing graph properties. *Random Structures and Algorithms*, 23(1):23–57, 2003.

[16] Oded Goldreich and Luca Trevisan. Errata for [15], 2005. available at the following URL: http://www.wisdom.weizmann.ac.il/~oded/PS/tt-err.ps.

[17] Mira Gonen and Dana Ron. On the benefits of adaptivity in property testing of dense graphs. In *Proceedings of RANDOM 2007*, pages 525–539, Berlin, Heidelberg, 2007. Springer-Verlag.

[18] Oleg Pikhurko. An analytic approach to stability. *Discrete Mathematics*, 310(21):2951–2964, 2010. Manuscript available on ArXiv since 2008, at http://arxiv.org/abs/0812.0214.

[19] Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th FOCS*, pages 222–227, 1977.