

A Characterization of Constant-Sample Testable Properties

Eric Blais*

David R. Cheriton School of Computer Science
University of Waterloo
eric.blais@uwaterloo.ca

Yuichi Yoshida†

National Institute of Informatics *and*
Preferred Infrastructure, Inc.
yyoshida@nii.ac.jp

December 16, 2016

Abstract

We characterize the set of properties of Boolean-valued functions on a finite domain \mathcal{X} that are testable with a constant number of samples. Specifically, we show that a property \mathcal{P} is testable with a constant number of samples if and only if it is (essentially) a k -part symmetric property for some constant k , where a property is k -part symmetric if there is a partition S_1, \dots, S_k of \mathcal{X} such that whether $f : \mathcal{X} \rightarrow \{0, 1\}$ satisfies the property is determined solely by the densities of f on S_1, \dots, S_k .

We use this characterization to obtain a number of corollaries, namely:

- A graph property \mathcal{P} is testable with a constant number of samples if and only if whether a graph G satisfies \mathcal{P} is (essentially) determined by the edge density of G .
- An affine-invariant property \mathcal{P} of functions $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is testable with a constant number of samples if and only if whether f satisfies \mathcal{P} is (essentially) determined by the density of f .
- For every constant $d \geq 1$, monotonicity of functions $f : [n]^d \rightarrow \{0, 1\}$ on the d -dimensional hypergrid is testable with a constant number of samples.

1 Introduction

Property testing [22, 17] is concerned with the general question: for which properties of combinatorial objects can we efficiently distinguish the objects that have the property from those that are “far” from having the same property? Consideration of this question has led to surprisingly powerful results: many natural graph properties [17, 2], algebraic properties of functions on finite fields [22, 27], and structural properties of Boolean functions [12, 10], for example, can be tested with a *constant* number of queries to the object being tested. Nearly all of these results appear to rely critically on the testing algorithm’s ability to query the unknown object at locations of its choosing. The goal of the present work is to determine to what extent this is actually the case: is it possible that some (or most!) of these properties can still be tested efficiently even if the tester has *no* control over the choice of queries that it makes?

*Supported by an NSERC Discovery Grant

†Supported by JSPS Grant-in-Aid for Young Scientists (B) (No. 26730009), MEXT Grant-in-Aid for Scientific Research on Innovative Areas (24106003), and JST, ERATO, Kawarabayashi Large Graph Project.

Our main question is made precise in the *sample-based* property testing model originally introduced by Goldreich, Goldwasser, and Ron [17]. For a finite set \mathcal{X} , let $\{0, 1\}^{\mathcal{X}}$ denote the set of Boolean-valued functions on \mathcal{X} endowed with the normalized Hamming distance $d(f, g) := |\{x \in \mathcal{X} : f(x) \neq g(x)\}|/|\mathcal{X}|$. A *property* of functions mapping \mathcal{X} to $\{0, 1\}$ is a subset $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$. A function is ϵ -close to \mathcal{P} if it is in the set $\mathcal{P}_\epsilon := \{f \in \{0, 1\}^{\mathcal{X}} : \exists g \in \mathcal{P} \text{ s.t. } d(f, g) \leq \epsilon\}$; otherwise, it is ϵ -far from \mathcal{P} . An *s-sample ϵ -tester* for a property \mathcal{P} is a randomized algorithm with bounded error that observes s pairs $(x_1, f(x_1)), \dots, (x_s, f(x_s)) \in \mathcal{X} \times \{0, 1\}$ with x_1, \dots, x_s drawn independently and uniformly at random from \mathcal{X} and then accepts when $f \in \mathcal{P}$ and rejects when f is ϵ -far from \mathcal{P} . The *sample complexity* of \mathcal{P} for some $\epsilon > 0$ is the minimum value of s such that it has an s -sample ϵ -tester. When the sample complexity of \mathcal{P} is independent of the domain size $|\mathcal{X}|$ of the functions for every $\epsilon > 0$, we say that \mathcal{P} is *constant-sample testable*.

In the 20 years since the original introduction of the sample-based property testing model, a number of different properties have been shown to be constant-sample testable, including all symmetric properties (folklore; see the discussion in [19]), unions of intervals [20], decision trees over low-dimensional domains [20], and convexity of images [21, 5]. Conversely, many other properties, such as monotonicity of Boolean functions [16], linearity [3, 18], linear threshold functions [4], and k -colorability of graphs [18], are known to not be constant-sample testable. Yet, despite a widespread belief that most properties are not testable with a constant number of samples, it remained open until now to determine whether this is actually the case or not. Our main result settles this question, and in the process unifies all of the above results and explains what makes various properties testable with a constant number of samples or not.

1.1 Main result

We show that constant-sample testability is closely tied to a particular notion of symmetry—or invariance—of properties. Let $\mathcal{S}_{\mathcal{X}}$ denote the set of permutations on a finite set \mathcal{X} , and for any subset $S \subseteq \mathcal{X}$, let $\mathcal{S}_{\mathcal{X}}^{(S)}$ denote the set of permutations on \mathcal{X} that preserves the elements in S . A permutation $\pi \in \mathcal{S}_{\mathcal{X}}$ acts on functions $f : \mathcal{X} \rightarrow \{0, 1\}$ in the obvious way: πf is the function that satisfies $(\pi f)(x) = f(\pi x)$ for every $x \in \mathcal{X}$. The property $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$ is *invariant* under a permutation $\pi \in \mathcal{S}_{\mathcal{X}}$ if for every $f \in \mathcal{P}$, we also have $\pi f \in \mathcal{P}$. \mathcal{P} is (*fully*) *symmetric* if it is invariant under all permutations in $\mathcal{S}_{\mathcal{X}}$. The following definition relaxes this condition to obtain a notion of “partial” symmetry.

Definition 1. For any $k \in \mathbb{N}$, the property $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$ is *k-part symmetric* if there is a partition of \mathcal{X} into k parts X_1, \dots, X_k such that \mathcal{P} is invariant under all the permutations in $\mathcal{S}_{X_1} \times \dots \times \mathcal{S}_{X_k}$, where \mathcal{S}_{X_i} is the set of permutations over \mathcal{X} that preserves elements in X_i .

Equivalently, \mathcal{P} is k -part symmetric if there exists a partition X_1, \dots, X_k of \mathcal{X} such that the event $f \in \mathcal{P}$ is completely determined by the density $\frac{|f^{-1}(1) \cap X_i|}{|X_i|}$ of f in each of the sets X_1, \dots, X_k . Our main result shows that $O(1)$ -part symmetry is also essentially equivalent to constant-sample testability.

Theorem 1. *The property \mathcal{P} of a function $f : \mathcal{X} \rightarrow \{0, 1\}$ is constant-sample testable if and only if for any $\epsilon > 0$, there exists a constant $k = k_{\mathcal{P}}(\epsilon)$ that is independent of \mathcal{X} and a k -part symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$.*

In words, Theorem 1 says that constant-sample testable properties are the properties \mathcal{P} that can be covered by some $O(1)$ -part symmetric property \mathcal{P}' that does not include any function that

is ϵ -far from \mathcal{P} . Note that this characterization cannot be replaced with the condition that \mathcal{P} itself is k -part symmetric. To see this, consider the *function non-identity* property $\text{NOTEQ}(g)$ that includes every function except some non-constant function $g : \mathcal{X} \rightarrow \{0, 1\}$. This property is not k -part symmetric for any $k = O(1)$, but the trivial algorithm that accepts every function is a valid ϵ -tester for $\text{NOTEQ}(g)$ for any constant $\epsilon > 0$.

Theorem 1 can easily be generalized to apply to properties of functions mapping \mathcal{X} to any finite set \mathcal{Y} . We restrict our attention to the range $\mathcal{Y} = \{0, 1\}$ for simplicity and clarity of presentation. The sample-based property testing model is naturally extended to non-uniform distributions over the input domain. It appears likely that Theorem 1 can be generalized to this more general setting as well, though we have not attempted to do so.

The proof of Theorem 1 follows the general outline of previous characterizations of the properties testable in the query-based model (e.g., [2, 27, 6]). As with those results, the most interesting part of the proof is the direction showing that constant-sample testability implies coverage by an $O(1)$ -part symmetric property, and this proof is established with a regularity lemma. Our proof departs from previous results in both the type of regularity lemma that we use and in how we use it. For a more detailed discussion of the proof, see Section 2.1.

1.2 Applications

The characterization of constant-sample testability in Theorem 1 can be used to derive a number of different corollaries. We describe a few of these.

When \mathcal{X} is identified with the $\binom{n}{2}$ pairs of vertices in V , the function $f : \mathcal{X} \rightarrow \{0, 1\}$ represents a graph $G = (V, E)$ where $E = f^{-1}(1)$. A *graph property* is a property of these functions that is invariant under relabelling of the vertices. Observant readers will have noted that no non-symmetric graph property was present in the list of properties that have been determined to be constant-sample testable. Using Theorem 1, we can show that this is unavoidable: the only graph properties that are constant-sample testable are those that are (essentially) fully symmetric.

Corollary 1. *For every $\epsilon > 0$, if \mathcal{P} is a graph property that is ϵ -testable with a constant number of samples, then there is a symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$.*

When \mathcal{X} is identified with a finite field \mathbb{F}_p^n , a property $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$ is *affine-invariant* if it is invariant under any affine transformation over \mathbb{F}_p^n . Theorem 1 can be used to show that symmetric properties are essentially the only constant-sample testable affine-invariant properties.

Corollary 2. *For every $\epsilon > 0$, if \mathcal{P} is an affine-invariant property of functions $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ that is ϵ -testable with a constant number of samples, then there is a symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$.*

These two corollaries show that Theorem 1 can be used to show that some properties are not constant-sample testable. In our third application, we show that the other direction of the characterization can also be used to show that some properties are constant-sample testable.

Fix a constant $d \geq 1$. Two points $x, y \in [n]^d := \{1, 2, \dots, n\}^d$ satisfy $x \preceq y$ when $x_1 \leq y_1, \dots$, and $x_d \leq y_d$. The function $f : [n]^d \rightarrow \{0, 1\}$ is *monotone* if for every $x \preceq y \in [n]^d$, we have $f(x) \leq f(y)$. When $d = 1$, it is folklore knowledge that monotonicity of Boolean-valued functions on the line is constant-sample testable. Using Theorem 1, we give an easy proof showing that the same holds for every other constant dimension d .

Corollary 3. *For every constant $d \geq 1$ and constant $\epsilon > 0$, we can ϵ -test monotonicity of functions $f : [n]^d \rightarrow \{0, 1\}$ on the d -dimensional hypergrid with a constant number of samples.*

Chen, Servedio, and Tan [8] showed that the number of queries (and thus also of samples) required to test monotonicity of $f : [n]^d \rightarrow \{0, 1\}$ must depend on d . (The same result for the case where $n = 2$ was first established by Fischer et al. [14].) Combined with the result above, this shows that monotonicity of Boolean-valued functions on the hypergrid is constant-sample testable if and only if the number of dimensions of the hypergrid is constant.

1.3 Related work

Sample-based property testing. The first general result regarding constant-sample testable properties goes back to the original work of Goldreich, Goldwasser, and Ron [17]. They showed that every property with constant VC dimension (and, more generally, every property that corresponds to a class of functions that can be properly learned with a constant number of samples) is constant-sample testable. As they also show, this condition is not necessary for constant-sample testability—in fact, there are even properties that are testable with a constant number of samples whose corresponding class require a *linear* number of samples to learn [17, Prop. 3.3.1].

More general results on sample-based testers were obtained by Balcan *et al.* [4]. In particular, they defined a notion of *testing dimension* of a property \mathcal{P} in terms of the total variation distance between the distributions on the tester’s observations when a function is drawn from distributions π_{yes} and π_{no} essentially supported on \mathcal{P} and $\overline{\mathcal{P}}_\epsilon$, respectively. They show that this testing dimension captures the sample complexity of \mathcal{P} up to constant factors, and observe that it can be interpreted as an “average VC dimension”-type of complexity measure. It would be interesting to see whether the combinatorial characterization in Theorem 1 could be combined with these results to offer new insights into the connections between invariance and VC dimension-like complexity measures.

Finally, Goldreich and Ron [18] and Fischer *et al.* [11, 13] established connections between the query- and sample-based models of property testing giving sufficient conditions for *sublinear*-sample testability of properties. The exact bounds between sample complexity and partial symmetry in the proof of Theorem 1 yield another sufficient condition for sublinear-sample testability: every property \mathcal{P} that can be ϵ -covered by an $o(\log \log |\mathcal{X}|)$ -part symmetric function \mathcal{P}' has sublinear sample complexity $o(|\mathcal{X}|)$. As far as we can tell, these two characterizations are incomparable.

Symmetry and testability. The present work was heavily influenced by the systematic exploration of connections between the invariances of properties and their testability initiated by Kaufman and Sudan [19]. (See also [24].) In that work, the authors showed that such connections yield new insights into the testability of algebraic properties in the query-based property testing model, and advocated for further study of the invariance of properties as a means to better understand their testability. Theorem 1 provides evidence that this approach is a critical tool in the study of sample-based property testing model as well.

The notion of partial symmetry and its connections to computational efficiency has a long history—it goes back at least to the pioneering work of Shannon [23]. Partial symmetry also appeared previously in a property testing context in the authors’ joint work with Amit Weinstein on characterizing the set of functions for which isomorphism testing is constant-query testable [7]. However, it should be noted that the notion of partial symmetry considered in [7] does *not* correspond to the notion of k -part symmetry studied here. In fact, as mentioned in the conference

version of that paper, there are 2-part symmetric functions for which isomorphism testing is not constant-query testable, so the two characterizations inherently require different notions of partial symmetry.

1.4 Organization

The proof of Theorem 1 is presented in Section 2. The proofs of the application results are in Section 3. Finally, since the weak regularity lemma that we use in the proof of Theorem 1 is not completely standard, we include its proof in Section 4 for completeness.

2 Proof of Theorem 1

We prove the two parts (sufficiency and necessity) of Theorem 1 in Sections 2.2 and 2.3, respectively. But first, we provide a high-level overview of the proof and discuss the connections with regularity lemmas in Section 2.1.

2.1 Overview of the proof

Symmetry implies testability. The proof of this direction of the theorem is straightforward and is obtained by generalizing the following folklore proof that symmetric properties can be tested with a constant number of samples. Let \mathcal{P} be any symmetric property. A tester can estimate the density $\mathbb{E}_{x \in \mathcal{X}}[f(x)]$ up to additive accuracy γ for any small $\gamma > 0$ with a constant number of samples. This estimated density can be used to accept or reject the function based on how close it is to the density of the functions in \mathcal{P} . The validity of this tester is established by showing that a function can be ϵ -far from \mathcal{P} only when its density is far from the density of every function in \mathcal{P} .

Consider now a property \mathcal{P} that is k -part symmetric for some constant k . Let X_1, \dots, X_k be a partition of \mathcal{X} associated with \mathcal{P} . We show that a tester which estimates the densities $\mu_i(f) := \mathbb{E}_{x \in X_i}[f(x)]$ for each $i = 1, \dots, k$ and then uses these densities to accept or reject is a valid tester for \mathcal{P} . We do so by showing that any function that is ϵ -far from \mathcal{P} must have a density vector that is far from those of every function in \mathcal{P} .

Testability implies symmetry. To establish the second part of the theorem, we want to show that the existence of a constant-sample tester \mathcal{T} for a property \mathcal{P} implies that there is a partition of \mathcal{X} into a constant number of parts for which \mathcal{P} is nearly determined by the density of functions within those parts. We do so by using a variant of the Frieze–Kannan weak regularity lemma [15] for hypergraphs. An s -uniform weighted hypergraph is a hypergraph $G = (V, E)$ on $|V|$ vertices where $E : V^s \rightarrow [0, 1]$ denotes the weight associated with each hyperedge. Given a partition V_1, \dots, V_k of V and a multi-index $I = (i_1, \dots, i_s)$ with $i_1, \dots, i_s \in [k]$, the *expected weight of hyperedges of G in V_I* is $w_G(V_I) = w_G(V_{i_1}, \dots, V_{i_s}) = |\{E(v_1, \dots, v_s) : v_1 \in V_{i_1}, \dots, v_s \in V_{i_s}\}| / \prod_{i \in I} |V_i|$. Furthermore, for any subset $S \subseteq V$, we define $S \cap V_I = (S \cap V_{i_1}, \dots, S \cap V_{i_s})$.

Lemma 1 (Weak regularity lemma). *For every $\epsilon > 0$ and every s -uniform weighted hypergraph $G = (V, E)$, there is a partition V_1, \dots, V_k of V with $k = 2^{O(\log(\frac{1}{\epsilon})/\epsilon^2)}$ parts such that for every subset $S \subseteq V$,*

$$\sum_{I \in [k]^s} \frac{\prod_{i \in I} |S \cap V_i|}{|V|^s} |w_G(S \cap V_I) - w_G(V_I)| \leq \epsilon. \quad (1)$$

This specific formulation of the weak regularity lemma seems not to have appeared previously in the literature, but its proof is essentially the same as that of usual formulations of the weak regularity lemma. For completeness, we provide a proof of Lemma 1 in Section 4.

Lemma 1 is best described informally when we consider the special case where we consider unweighted graphs. In this setting, the weak regularity lemma says that for every graph G , there is a partition of the vertices of G into $k = O(1)$ parts V_1, \dots, V_k such that for every subset S of vertices, the density of edges between $S \cap V_i$ and $S \cap V_j$ is close to the density between V_i and V_j *on average* over the choice of V_i and V_j . Regularity lemmas where this density-closeness condition is satisfied for *almost all* pairs of parts V_i and V_j are known as “strong” regularity lemmas. To the best of our knowledge, all previous characterization results in property testing that relied on regularity lemmas (e.g., [2, 27, 6]) used strong regularity lemmas. This approach unavoidably introduces tower-type dependencies between the query complexity and the characterization parameters. By using a weak regularity lemma instead, we get a much better (though still doubly-exponential) dependence between the sample complexity and the partial symmetry parameter.

The second point of departure of our proof from previous characterizations is in how we use the regularity lemma. In the prior work, the regularity lemma was applied to the tested object itself (e.g., the dense graphs being tested in [2]) and the testability of the property was used to show that the objects with the given property could be described by some combinatorial characteristics related to the regular partition whose existence is promised by the regularity lemma. Instead, in our proof of Theorem 1, we apply Lemma 1 to a hypergraph associated with the *tester* itself, not with the tested object.

Specifically, let \mathcal{T} be an s -sample ϵ -tester for some property $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$. We associate \mathcal{T} with an s -uniform weighted hypergraph $G_{\mathcal{T}}$ on the set of vertices $\mathcal{X} \times \{0, 1\}$. The weight of each s -hyperedge of $G_{\mathcal{T}}$ is the acceptance probability of \mathcal{T} when its s observations correspond to the s vertices covered by the hyperedge. By associating each function $f : \mathcal{X} \rightarrow \{0, 1\}$ with the subset $S \subseteq \mathcal{X} \times \{0, 1\}$ that includes all $2^{\mathcal{X}}$ vertices of the form $(x, f(x))$, we see that the probability that \mathcal{T} accepts f is the expected value of a hyperedge whose s vertices are drawn uniformly and independently at random from the set S . We can use Lemma 1 to show that there is a partition of V into a constant number of parts such that for each function f with associated set S , this probability is well approximated by some function of the density of S in each of the parts. We then use this promised partition of V to partition the original input domain \mathcal{X} into a constant number of parts where membership in \mathcal{P} is essentially determined by the density of a function in each of these parts, as required.

2.2 Proof that symmetry implies testability

We now begin the proof of Theorem 1 with the easy direction.

Lemma 2. *Let $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$ be a property where for every $\epsilon > 0$, there exists a constant $k = k_{\mathcal{P}}(\epsilon)$ that is independent of \mathcal{X} and a k -part symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_{\epsilon}$. Then \mathcal{P} is constant-sample testable.*

Proof. Fix $\epsilon > 0$. We show that we can distinguish functions in \mathcal{P} from functions that are ϵ -far from \mathcal{P} with a constant number of samples. From the condition in the statement of the lemma, there exists a k -part symmetric property \mathcal{P}' with $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_{\epsilon/2}$ for some $k = k_{\mathcal{P}}(\epsilon/2)$. Let S_1, \dots, S_k be a partition of \mathcal{X} such that whether a function $f : \mathcal{X} \rightarrow \{0, 1\}$ satisfies \mathcal{P}' is determined by $\mu_{S_1}(f), \dots, \mu_{S_k}(f)$. For a set $S \subseteq \mathcal{X}$, let $c_S(f) = \mu_S(f)|S|$ be the number of $x \in S$ with $f(x) = 1$.

Our algorithm for testing \mathcal{P} is as follows. For each $i \in [k]$, we draw $q := O(k^2 \log k / \epsilon^2)$ samples x_1, \dots, x_q and computes the estimates $\tilde{c}_{S_i}(f) := \frac{|\mathcal{X}|}{q} \sum_{j \in [q]: x_j \in S_i} f(x_j)$ for each $i \in [k]$. We accept if there exists $g \in \mathcal{P}'$ such that

$$\sum_{i \in [k]} |\tilde{c}_{S_i}(f) - c_{S_i}(g)| < \frac{\epsilon}{4} |\mathcal{X}|,$$

and reject otherwise.

Let us now establish the correctness of the algorithm. By Hoeffding's bound, for each $i \in [k]$, we have $|c_{S_i}(f) - \tilde{c}_{S_i}(f)| < \frac{\epsilon}{4k} |\mathcal{X}|$ with probability at least $1 - \frac{1}{3k}$ by choosing the hidden constant in the definition of q sufficiently large. By union bound, with probability at least $2/3$, we have $|c_{S_i}(f) - \tilde{c}_{S_i}(f)| < \frac{\epsilon}{4k} |\mathcal{X}|$ for every $i \in [k]$. In what follows, we assume this happens.

If $f \in \mathcal{P}$, then the algorithm accepts f because $\sum_{i \in [k]} |\tilde{c}_{S_i}(f) - c_{S_i}(f)| < \frac{\epsilon}{4} |\mathcal{X}|$ and $f \in \mathcal{P} \subseteq \mathcal{P}'$.

If f is ϵ -far from satisfying \mathcal{P} , then for any $g \in \mathcal{P}'$, the triangle inequality and the fact that $\mathcal{P}' \subseteq \mathcal{P}_{\epsilon/2}$ imply that

$$\sum_{i \in [k]} |\tilde{c}_{S_i}(f) - c_{S_i}(g)| \geq \sum_{i \in [k]} \left(|c_{S_i}(f) - c_{S_i}(g)| - |\tilde{c}_{S_i}(f) - c_{S_i}(f)| \right) > \frac{\epsilon}{2} |\mathcal{X}| - \frac{\epsilon}{4} |\mathcal{X}| = \frac{\epsilon}{4} |\mathcal{X}|$$

and, therefore, the algorithm rejects f . □

2.3 Testability implies symmetry

Suppose that a property \mathcal{P} is testable by a tester \mathcal{T} with sample complexity s . We want to show that for any $\epsilon > 0$, there exists a k -part symmetric property \mathcal{P}' for $k = k(\epsilon)$ such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$.

For any $\mathbf{x} = (x_1, \dots, x_s) \in \mathcal{X}^s$, we define $f(\mathbf{x}) = (f(x_1), \dots, f(x_s))$ and we let $T(\mathbf{x}, f(\mathbf{x})) \in [0, 1]$ denote the acceptance probability of the tester \mathcal{T} of f when the samples drawn are \mathbf{x} . The overall acceptance probability of f by \mathcal{T} is

$$p_{\mathcal{T}}(f) = \mathbb{E}_{\mathbf{x}} [T(\mathbf{x}, f(\mathbf{x}))]. \quad (2)$$

We show that there is a family \mathcal{S} of a constant number of subsets of \mathcal{X} such that, for every function $f : \mathcal{X} \rightarrow \{0, 1\}$, the acceptance probability $p_{\mathcal{T}}(f)$ is almost completely determined by the density of f on the subsets in \mathcal{S} .

Lemma 3. *For any $\epsilon > 0$ and any s -sample tester \mathcal{T} , there is a family $\mathcal{S} = \{S_1, \dots, S_m\}$ of $m \leq 2^{O(2^{2s}/\epsilon^2)}$ subsets of \mathcal{X} such that for every $f : \mathcal{X} \rightarrow \{0, 1\}$,*

$$|p_{\mathcal{T}}(f) - \varphi_{\mathcal{T}}(\mu_{S_1}(f), \dots, \mu_{S_m}(f))| \leq \epsilon$$

where $\mu_S(f) = \mathbb{E}_{x \in S} [f(x)]$ is the density of f on the subset $S \subseteq \mathcal{X}$ and $\varphi_{\mathcal{T}} : [0, 1]^m \rightarrow [0, 1]$ is some fixed function.

Proof. We consider the weighted hypergraph $G = (V, E)$ where $V = \mathcal{X} \times \{0, 1\}$ and E is constructed by adding a hyperedge $((x_1, y_1), \dots, (x_s, y_s))$ of weight $T(\mathbf{x}, \mathbf{y})$ for each $\mathbf{x} = (x_1, \dots, x_s) \in \mathcal{X}^s$ and $\mathbf{y} = (y_1, \dots, y_s) \in \{0, 1\}^s$.

A function $f : \mathcal{X} \rightarrow \{0, 1\}$ corresponds to a subset $S \subseteq V$ of size $|\mathcal{X}| = |V|/2$, that is, $S := \{(x, f(x)) \mid x \in \mathcal{X}\}$. The probability that \mathcal{T} accepts f is

$$\begin{aligned} p_{\mathcal{T}}(f) &= \mathbb{E}_{\mathbf{x} \in \mathcal{X}^s} [T(\mathbf{x}, f(\mathbf{x}))] = \mathbb{E}_{\mathbf{v} \in V^s} [E(\mathbf{v}) \mid v \in S^s] = \sum_{I \in [k]^s} \frac{\prod_{i \in I} |V_i \cap S|}{|S|^s} w_G(S \cap V_I) \\ &= \sum_{I \in [k]^s} 2^s \frac{\prod_{i \in I} |V_i \cap S|}{|V|^s} w_G(S \cap V_I) = \sum_{I \in [k]^s} 2^s \frac{\prod_{i \in I} |V_i \cap S|}{|V|^s} w_G(V_I) \pm \epsilon \end{aligned}$$

where the last step is by Lemma 1 applied with approximation parameter $\epsilon/2^s$. For every part V_i , let $V_i^1 = \{x \in \mathcal{X} : (x, 1) \in V_i\}$ and let $V_i^0 = \{x \in \mathcal{X} : (x, 0) \in V_i\}$. Then the value of $\sum_{I \in [k]^s} 2^s \frac{\prod_{i \in I} |V_i \cap S|}{|V|^s} w(V_I)$ is completely determined by the density of f on $V_1^1, V_1^0, V_2^1, V_2^0, \dots, V_k^1, V_k^0$. \square

We are now ready to complete the second part of the proof of Theorem 1.

Lemma 4. *Let $\mathcal{P} \subseteq \{0, 1\}^{\mathcal{X}}$ be constant-sample testable. Then for every $\epsilon > 0$, there exists a constant $k = k_{\mathcal{P}}(\epsilon)$ that is independent of \mathcal{X} and a k -part symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_{\epsilon}$.*

Proof. Fix any $\epsilon > 0$ and let \mathcal{T} be an s -sample ϵ -tester for \mathcal{P} . Let $\gamma < \frac{1}{3}$ be any constant that is less than $\frac{1}{3}$. By Lemma 3 applied with the parameter γ , there is a family $\mathcal{S} = \{S_1, \dots, S_m\}$ with $m = 2^{O(2^{2s}/\gamma^2)}$ sets such that for every $f : \mathcal{X} \rightarrow \{0, 1\}$,

$$|p_{\mathcal{T}}(f) - \varphi_{\mathcal{T}}(\mu_{S_1}(f), \dots, \mu_{S_k}(f))| \leq \gamma. \quad (3)$$

Define

$$\mathcal{P}' = \{f : \mathcal{X} \rightarrow \{0, 1\} : \exists g \in \mathcal{P} \text{ s.t. } (\mu_{S_1}(f), \dots, \mu_{S_k}(f)) = (\mu_{S_1}(g), \dots, \mu_{S_k}(g))\}.$$

This construction trivially guarantees that $\mathcal{P}' \supseteq \mathcal{P}$. Furthermore, (3) guarantees that for every $f \in \mathcal{P}'$, if we let $g \in \mathcal{P}$ be one of the elements with the same density profile as f ,

$$p_{\mathcal{T}}(f) \geq \varphi_{\mathcal{T}}(\mu_{S_1}(f), \dots, \mu_{S_k}(f)) - \gamma = \varphi_{\mathcal{T}}(\mu_{S_1}(g), \dots, \mu_{S_k}(g)) - \gamma \geq p_{\mathcal{T}}(g) - 2\gamma.$$

Since $\gamma < 1/3$, the fact that \mathcal{T} is an ϵ -tester for \mathcal{P} and that $g \in \mathcal{P}$ imply that $f \in \mathcal{P}_{\epsilon}$.

Let S'_1, \dots, S'_k be the family of sets obtained by taking intersections and complements of S_1, \dots, S_m . Note that S'_1, \dots, S'_k forms a partition of \mathcal{X} and $\mu_{S_1}, \dots, \mu_{S_m}$ is completely determined by $\mu_{S'_1}, \dots, \mu_{S'_k}$. Furthermore, $k = O(2^m)$. Hence, \mathcal{P}' is a k -part symmetric property induced by the partition S'_1, \dots, S'_k with $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_{\epsilon}$. \square

Theorem 1 follows immediately from Lemmas 2 and 4.

3 Applications

3.1 Affine-invariant properties

For an affine transformation $A : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ and a function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$, we define $Af : \mathbb{F}_p^n \rightarrow \{0, 1\}$ to be the function that satisfies $Af(x) = f(Ax)$ for every $x \in \mathbb{F}_p^n$. A property \mathcal{P} of functions

$f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is *affine-invariant* if for any affine transformation $A : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^n$ and $f \in \mathcal{P}$, we have $Af \in \mathcal{P}$. Our characterization shows that the only affine-invariant properties of functions $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ that are testable with a constant number of samples are the (fully symmetric) properties that are determined by the density of f .

Corollary 2 (Restated). *For every $\epsilon > 0$, if \mathcal{P} is an affine-invariant property of functions $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ that is ϵ -testable with a constant number of samples, then there is a symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$.*

Proof. By Theorem 1, if \mathcal{P} is testable with $O(1)$ samples, then there are subsets S_1, \dots, S_k of \mathbb{F}_p^n with $k = O(1)$ and a property \mathcal{P}'' such that $\mathcal{P} \subseteq \mathcal{P}'' \subseteq \mathcal{P}_\epsilon$ and \mathcal{P}'' is invariant under all permutations of S_1, \dots, S_k . Let \mathcal{P}' be the closure of \mathcal{P}'' under all affine transformations over \mathbb{F}_p^n . (I.e., $\mathcal{P}' = \{Af : f \in \mathcal{P}'', A \text{ is affine}\}$.) Since \mathcal{P} itself is invariant under affine transformations, we have that $\mathcal{P}' \subseteq \mathcal{P}_\epsilon$. We want to show that \mathcal{P}' is symmetric. To show this, it suffices to show that \mathcal{P}' is closed under transpositions. I.e., that for every $x, y \in \mathbb{F}_p^n$ and $f \in \mathcal{P}'$, the function g obtained by setting $g(x) = f(y)$, $g(y) = f(x)$, and $g(z) = f(z)$ for every other $z \notin \{x, y\}$ is also in \mathcal{P}' . We write $g = (xy)f$ to denote the action of the transposition (xy) on f .

If $x, y \in S_i$ for some $i \in [k]$, then our conclusion follows immediately from the invariance of \mathcal{P}'' over permutations on S_i . Otherwise, let $w, z \in S_j$ for some $j \in [k]$. Since the set of affine transformations is a doubly-transitive action on \mathbb{F}_p^n , there is a transformation A such that $A(x) = w$ and $A(y) = z$. Then $(xy)f = A^{-1}(wz)Af \in \mathcal{P}'$, as we wanted to show. \square

3.2 Graph properties

Similarly, our characterization shows that the only graph properties that are testable with a constant number of samples are the (fully symmetric) properties that are determined by the edge density of the graph.

Corollary 1 (Restated). *For every $\epsilon > 0$, if \mathcal{P} is a graph property that is ϵ -testable with a constant number of samples, then there is a symmetric property \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$.*

Proof. The proof is nearly identical. By Theorem 1, if \mathcal{P} is testable with $O(1)$ samples, then there are subsets S_1, \dots, S_k of the edge set with $k = O(1)$ and a property \mathcal{P}'' such that $\mathcal{P} \subseteq \mathcal{P}'' \subseteq \mathcal{P}_\epsilon$ and \mathcal{P}'' is invariant under all permutations of S_1, \dots, S_k . Let \mathcal{P}' be the closure of \mathcal{P}'' under all permutations of the vertex set. We again have that $\mathcal{P} \subseteq \mathcal{P}' \subseteq \mathcal{P}_\epsilon$ and we want to show that \mathcal{P}' is invariant under every transposition of the edge set.

The one change with the affine-invariant property proof is that the set of permutations of the vertices of a graph is not a 2-transitive action on the set of edges of this graph. But the same idea still works because we can always find a vertex permutation to send two edges on disjoint vertices to a same part S_i that also contains a pair of edges on disjoint vertices, and we can find a vertex permutation to send two edges that share a common vertex to a part S_j that also contains two edges that share a common vertex. So for every pair of edges e_1, e_2 , we have a vertex permutation π_V and a pair of edges $e_3, e_4 \in S_\ell$ such that if $G \in \mathcal{P}'$, then $(e_1 e_2)G = \pi_V^{-1}(e_3 e_4)\pi_V G \in \mathcal{P}'$. \square

3.3 Testing monotonicity

With Theorem 1, to show that monotonicity of functions $f : [n]^d \rightarrow \{0, 1\}$ is constant-sample testable, it suffices to identify an $O(1)$ -part symmetric function that covers all the monotone functions and does not include any function that is far from monotone. This is what we do below.

Corollary 3 (Restated). *For every constant $d \geq 1$ and constant $\epsilon > 0$, we can ϵ -test monotonicity of functions $f : [n]^d \rightarrow \{0, 1\}$ on the d -dimensional hypergrid with a constant number of samples.*

Proof. For any $\epsilon > 0$, fix $k = \lceil d/\epsilon \rceil$ and let \mathcal{R} be a partition of the space $[n]^d$ into k^d subgrids of side length (at most) $\lfloor \epsilon n \rfloor$ each. We identify the parts in \mathcal{R} with the points in $[k]^d$. For an input $x \in [n]^d$, let $\phi_{\mathcal{R}}(x)$ denote the part of \mathcal{R} that contains x .

Given some function $f : [n]^d \rightarrow \{0, 1\}$, define the \mathcal{R} -granular representation of f to be the function $f_{\mathcal{R}} : [k]^d \rightarrow \{0, 1, *\}$ defined by

$$f_{\mathcal{R}}(x) = \begin{cases} 0 & \text{if } \forall y \in [n]^d \text{ with } \phi_{\mathcal{R}}(y) = x, f(y) = 0 \\ 1 & \text{if } \forall y \in [n]^d \text{ with } \phi_{\mathcal{R}}(y) = x, f(y) = 1 \\ * & \text{otherwise.} \end{cases}$$

Let $\mathcal{P} = \{f : [n]^d \rightarrow \{0, 1\} : \exists \text{ monotone } g \text{ s.t. } f_{\mathcal{R}} = g_{\mathcal{R}}\}$ be the property that includes every function whose \mathcal{R} -granular representation equals that of a monotone function. By construction \mathcal{P} includes all the monotone functions and is invariant under any permutations within the $O(1)$ parts of \mathcal{R} . To complete the proof of the theorem, we want to show that every function in \mathcal{P} is ϵ -close to monotone.

Fix any $f \in \mathcal{P}$ and let $g : [n]^d \rightarrow \{0, 1\}$ be a monotone function for which $f_{\mathcal{R}} = g_{\mathcal{R}}$. The distance between f and g is bounded by

$$d(f, g) \leq \frac{|g_{\mathcal{R}}^{-1}(*)|}{k^d}.$$

Now consider the poset P on $[k]^d$ where $x \prec y$ iff $x_i < y_i$ for every $i \in [d]$. We first observe that the set $g_{\mathcal{R}}^{-1}(*)$ forms an anti-chain on this poset. Indeed, if there exist $x, y \in [k]^d$ with $x \prec y$ and $g(x) = g(y) = *$, then there exist $x', y' \in [n]^d$ such that $\phi_{\mathcal{R}}(x') = x$, $\phi_{\mathcal{R}}(y') = y$, $g(x') = 1$, and $g(y') = 0$. But this contradicts the monotonicity of g because $x' \leq y'$ holds from $x \prec y$.

For $x \in [k]^d$, let $x_{\max} = \max_{i \in [d]} x_i$ and $x_{\min} = \min_{i \in [d]} x_i$. Define $\mathbf{1} = (1, 1, 1, \dots, 1) \in [k]^d$ and $S = \{x \in [k]^d : x_{\min} = 1\}$. We can partition $[k]^d$ into $|S| = k^d - (k-1)^d \leq dk^{d-1}$ chains $(x, x + \mathbf{1}, x + 2 \cdot \mathbf{1}, \dots, x + (x_{\max} - 1) \cdot \mathbf{1})$, one for each $x \in S$. Therefore, by Dilworth's theorem, every anti-chain on P has size at most dk^{d-1} . In particular, this bound holds for the anti-chain $g^{-1}(*)$ so $d(f, g) \leq \frac{dk^{d-1}}{k^d} = \frac{d}{k} \leq \epsilon$. \square

4 Proof of the weak regularity lemma

4.1 Information theory

The proof we provide for Lemma 1 is information-theoretic. In this section, we will use bold fonts to denote random variables. We assume that the reader is familiar with the basic concepts of entropy and mutual information; a good introduction to these definitions is [9]. The only facts we use about these concepts is the chain rule for mutual information and the fact that the entropy of a random variable is at most the logarithm of the number of values it can take.

The one non-basic information-theoretic inequality that we use in the proof is an inequality established by Tao [25] and later refined by Ahlswede [1].

Lemma 5 (Tao [25], Ahlswede [1]). *Let \mathbf{y} , \mathbf{z} , and \mathbf{z}' be discrete random variables where $\mathbf{y} \in [-1, 1]$ and $\mathbf{z}' = \phi(\mathbf{z})$ for some function ϕ . Then*

$$\mathbb{E} \left[\left| \mathbb{E}[\mathbf{y} \mid \mathbf{z}'] - \mathbb{E}[\mathbf{y} \mid \mathbf{z}] \right| \right] \leq \sqrt{2 \ln 2 \cdot I(\mathbf{y}; \mathbf{z} \mid \mathbf{z}')}.$$

Tao originally used his inequality to offer an information-theoretic proof of Szemerédi's (strong) Regularity Lemma. The proof we offer below follows (a simplified version of) the same approach. The fact that Tao's proof of the strong regularity lemma can also be applied (with simplifications) to prove the Frieze–Kannan weak regularity lemma was observed previously by Trevisan [26].

4.2 Proof of Lemma 1

For $\tau > 0$, we say that a hypergraph is τ -granular if the weight of each hyperedge is a multiple of τ . When proving Lemma 1, we can assume that the given graph is τ -granular for $\tau = \Theta(\epsilon)$. To see this, let $G' = (V, E')$ be the hypergraph obtained from $G = (V, E)$ by rounding the weight of each hyperedge to a multiple of τ . Then, for any set $S \subseteq V$ and a partition V_1, \dots, V_k of V , we have

$$\begin{aligned} \sum_{I \in [k]^s} \frac{\prod_{i \in I} |S \cap V_i|}{|V|^s} |w_G(S \cap V_I) - w_G(V_I)| &= \sum_{I \in [k]^s} \frac{\prod_{i \in I} |S \cap V_i|}{|V|^s} |(w_{G'}(S \cap V_I) \pm \tau) - (w_{G'}(V_I) \pm \tau)| \\ &= \sum_{I \in [k]^s} \frac{\prod_{i \in I} |S \cap V_i|}{|V|^s} |w_{G'}(S \cap V_I) - w_{G'}(V_I)| \pm 2\tau. \end{aligned}$$

In order to make the right-hand side less than ϵ , it suffices to show that the claim of Lemma 1 holds for $\epsilon/3$ -granular hypergraphs with an error parameter $\epsilon/3$. In what follows, we assume the input graph G is $\epsilon/3$ -granular.

Let V_1, \dots, V_ℓ be any partition of V . Draw $\mathbf{v} \in V^s$ uniformly at random. Define $\mathbf{y} = E(\mathbf{v})$. Let $\psi : V^s \rightarrow [\ell]^s$ be the function that identifies the parts containing each of s vertices. For any set $S \subseteq V$, let $1_S : V^s \rightarrow \{0, 1\}^s$ be the indicator function of S for s -tuples of vertices. Define $\mathbf{z}_S = (\psi(\mathbf{v}), 1_S(\mathbf{v}))$ and $\mathbf{z}' = \psi(\mathbf{v})$. We consider two cases.

First, consider the situation where for every set $S \subseteq V$,

$$I(\mathbf{y}; \mathbf{z}_S \mid \mathbf{z}') \leq \frac{(\epsilon/3)^2}{2 \ln 2}. \quad (4)$$

Then by Tao's lemma, for every set S we also have

$$\begin{aligned} \sum_{I \in [k]^s, b \in \{0, 1\}^s} \frac{\prod_{j \in [s]} |S_{b_j} \cap V_{i_j}|}{|V|^s} |w_G(S_b \cap V_I) - w_G(V_I)| &= |\mathbb{E}[E(\mathbf{v}) \mid \psi(\mathbf{v})] - \mathbb{E}[E(\mathbf{v}) \mid \psi(\mathbf{v}), 1_S(\mathbf{v})]| \\ &= |\mathbb{E}[\mathbf{y} \mid \mathbf{z}'] - \mathbb{E}[\mathbf{y} \mid \mathbf{z}_S]| \leq \epsilon/3. \end{aligned}$$

And clearly

$$\sum_{I \in [k]^s} \frac{\prod_{j \in [s]} |S \cap V_{i_j}|}{|V|^s} |w_G(S \cap V_I) - w_G(V_I)| \leq \sum_{I \in [k]^s, b \in \{0, 1\}^s} \frac{\prod_{j \in [s]} |S_{b_j} \cap V_{i_j}|}{|V|^s} |w_G(S_b \cap V_I) - w_G(V_I)|$$

since the expression in the left-hand side is one of the terms in the sum (over b) on the right-hand side. So in this case the lemma holds.

Second, we need to consider the case where there is some set $S \subseteq V$ for which

$$I(\mathbf{y}; \mathbf{z}_S \mid \mathbf{z}') > \frac{(\epsilon/3)^2}{2 \ln 2}. \quad (5)$$

Then by the chain rule for mutual information

$$I(\mathbf{y}; \mathbf{z}_S) = I(\mathbf{y}; \mathbf{z}') + I(\mathbf{y}; \mathbf{z}_S \mid \mathbf{z}') \geq I(\mathbf{y}; \mathbf{z}') + \frac{(\epsilon/3)^2}{2 \ln 2}.$$

Define the *information value* of a partition with indicator function ψ as $I(E(\mathbf{v}); \psi(\mathbf{v}))$. Then the above observation shows that when (5) holds, we can obtain a refined partition with 2ℓ parts whose information value increases by at least $\frac{(\epsilon/3)^2}{2 \ln 2}$. The information value of any partition is bounded above by $H(\mathbf{y})$, so after at most $\frac{2 \ln 2 \cdot H(\mathbf{y})}{(\epsilon/3)^2}$ refinements, we must obtain a partition (with at most $2^{\frac{2 \ln 2 \cdot H(\mathbf{y})}{(\epsilon/3)^2}}$ parts) that satisfies (4). Since G is $\epsilon/3$ -granular, we have $H(\mathbf{y}) = O(\log(1/\epsilon))$, and the lemma follows. \square

References

- [1] Rudolf Ahlswede. The final form of Tao’s inequality relating conditional expectation and conditional mutual information. *Advances in Mathematics of Communications*, 1(2):239–242, 2007.
- [2] 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.
- [3] Noga Alon, Rani Hod, and Amit Weinstein. On active and passive testing. *Combinatorics, Probability and Computing*, 25(1):1–20, 2016.
- [4] Maria-Florina Balcan, Eric Blais, Avrim Blum, and Liu Yang. Active property testing. *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 21–30, 2012.
- [5] Piotr Berman, Meiram Murzabulatov, and Sofya Raskhodnikova. Tolerant testers of image properties. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 90:1–90:14, 2016.
- [6] Arnab Bhattacharyya and Yuichi Yoshida. An algebraic characterization of testable Boolean CSPs. In *Proceedings of the 40th International Colloquium Conference on Automata, Languages, and Programming (ICALP)*, pages 123–134, 2013.
- [7] Eric Blais, Amit Weinstein, and Yuichi Yoshida. Partially symmetric functions are efficiently isomorphism testable. *SIAM Journal on Computing*, 44(2):411–432, 2015.
- [8] Xi Chen, Rocco A Servedio, and Li-Yang Tan. New algorithms and lower bounds for monotonicity testing. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 286–295, 2014.

- [9] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [10] Ilias Diakonikolas, Homin K. Lee, Kevin Matulef, Krzysztof Onak, Ronitt Rubinfeld, Rocco A. Servedio, and Andrew Wan. Testing for concise representations. In *Proc. 48th Symposium on Foundations of Computer Science*, pages 549–558, 2007.
- [11] Eldar Fischer, Yonatan Goldhirsh, and Oded Lachish. Partial tests, universal tests and decomposability. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 483–500, 2014.
- [12] Eldar Fischer, Guy Kindler, Dana Ron, Shmuel Safra, and Alex Samorodnitsky. Testing juntas. *J. Comput. Syst. Sci.*, 68(4):753–787, 2004.
- [13] Eldar Fischer, Oded Lachish, and Yadu Vasudev. Trading query complexity for sample-based testing and multi-testing scalability. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1163–1182, 2015.
- [14] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 474–483, 2002.
- [15] Alan M. Frieze and Ravi Kannan. The regularity lemma and approximation schemes for dense problems. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 12–20, 1996.
- [16] Oded Goldreich, Shafi Goldwasser, Eric Lehman, Dana Ron, and Alex Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.
- [17] 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.
- [18] Oded Goldreich and Dana Ron. On sample-based testers. In *Proceedings of the 6th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 337–345, 2015.
- [19] Tali Kaufman and Madhu Sudan. Algebraic property testing: the role of invariance. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 403–412, 2008.
- [20] Michael Kearns and Dana Ron. Testing problems with sublearning sample complexity. *Journal of Computer and System Sciences*, 61(3):428–456, 2000.
- [21] Sofya Raskhodnikova. Approximate testing of visual properties. In *Proceedings of the 7th International Workshop on Randomization and Computation (RANDOM)*, pages 370–381, 2003.
- [22] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.
- [23] Claude Shannon. The synthesis of two-terminal switching circuits. *Bell System Technical Journal*, 28(1):59–98, 1949.

- [24] Madhu Sudan. Invariance in property testing. In Oded Goldreich, editor, *Property Testing: Current Research and Surveys*, pages 211–227. Springer, 2010.
- [25] Terence Tao. Szemerédi’s regularity lemma revisited. *Contributions to Discrete Mathematics*, 1(1), 2006.
- [26] Luca Trevisan. Entropy and the weak regularity lemma. <https://lucatrevisan.wordpress.com/2006/09/10/entropy-and-the-weak-regularity-lemma/>, 2006.
- [27] Yuichi Yoshida. A characterization of locally testable affine-invariant properties via decomposition theorems. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 154–163, 2014.