

Testing k -Monotonicity

(The Rise and Fall of Boolean Functions)*

Clément L. Canonne[†] Elena Grigorescu[‡] Siyao Guo[§] Akash Kumar[¶]

Karl Wimmer^{||}

March 10, 2017

Abstract

A Boolean k -monotone function defined over a finite poset domain \mathcal{D} alternates between the values 0 and 1 at most k times on any ascending chain in \mathcal{D} . Therefore, k -monotone functions are natural generalizations of the classical *monotone* functions, which are the 1-monotone functions.

Motivated by the recent interest in k -monotone functions in the context of circuit complexity and learning theory, and by the central role that monotonicity testing plays in the context of property testing, we initiate a systematic study of k -monotone functions, in the property testing model. In this model, the goal is to distinguish functions that are k -monotone (or are close to being k -monotone) from functions that are far from being k -monotone.

Our results include the following:

1. We demonstrate a separation between testing k -monotonicity and testing monotonicity, on the hypercube domain $\{0, 1\}^d$, for $k \geq 3$;
2. We demonstrate a separation between testing and learning on $\{0, 1\}^d$, for $k = \omega(\log d)$: testing k -monotonicity can be performed with $2^{O(\sqrt{d} \cdot \log d \cdot \log 1/\varepsilon)}$ queries, while learning k -monotone functions requires $2^{\Omega(k \cdot \sqrt{d} \cdot 1/\varepsilon)}$ queries (Blais et al. (RANDOM 2015)).
3. We present a tolerant test for functions $f: [n]^d \rightarrow \{0, 1\}$ with complexity independent of n , which makes progress on a problem left open by Berman et al. (STOC 2014).

Our techniques exploit the testing-by-learning paradigm, use novel applications of Fourier analysis on the grid $[n]^d$, and draw connections to distribution testing techniques.

*An extended abstract of this work is to appear in the Proceedings of the 8th Innovations in Theoretical Computer Science (ITCS '17).

[†]Columbia University. Email: cannonne@cs.columbia.edu. Research supported by NSF CCF-1115703 and NSF CCF-1319788.

[‡]Purdue University. Email: elena-g@purdue.edu. Research supported in part by NSF CCF-1649515.

[§]Simons Institute for the Theory of Computing, UC Berkeley. Email: siyao.guo@berkeley.edu.

[¶]Purdue University. Email: akumar@purdue.edu. Research supported in part by NSF CCF-1649515.

^{||}Duquesne University. Email: wimmerk@duq.edu.

Contents

1	Introduction	3
1.1	Our results	3
1.1.1	Testing k -monotonicity on the hypercube $\{0, 1\}^d$	3
1.1.2	Testing k -monotonicity on the hypergrid $[n]^d$	4
1.2	Proof overview and technical contribution	7
1.3	Discussion and open problems	8
1.4	Related work	9
1.5	Organization of the paper	9
2	Preliminaries	10
3	Structural results	11
4	On the Boolean hypercube	14
4.1	Upper bound: beating the learning approach	14
4.2	Lower bounds	15
4.2.1	One-sided lower bounds	16
4.2.2	Two-sided lower bounds	18
5	On the line	20
5.1	Upper and lower bounds for one-sided testers	20
5.2	Upper bound for two-sided testers: proof of Theorem 1.5	23
5.2.1	Testing k -monotonicity over $[Ck]$	23
5.2.2	Reducing $[n] \rightarrow \{0, 1\}$ to $[Ck] \rightarrow \{0, 1\}$.	26
6	On the grid	27
6.1	The case $k = 2$	27
6.2	Possible extensions	33
7	On the high-dimensional grid	33
7.1	Fully tolerant testing with $O(kd/(\varepsilon_2 - \varepsilon_1))^d$ queries	34
7.2	Tolerant testing <i>via</i> agnostic learning	36
7.3	Preliminary Fourier analysis tools	37
7.4	Finishing the proof of Proposition 7.3	38
8	Tolerant testing and applications to L_1-testing	39
A	Previous work on monotonicity testing	47

1 Introduction

A function $f: \mathcal{D} \rightarrow \{0, 1\}$, defined over a finite partially ordered domain (\mathcal{D}, \preceq) is said to be k -monotone, for some integer $k \geq 0$, if there does not exist $x_1 \preceq x_2 \preceq \dots \preceq x_{k+1}$ in \mathcal{D} such that $f(x_1) = 1$ and $f(x_i) \neq f(x_{i+1})$ for all $i \in [k]$. Note that 1-monotone functions are the classical *monotone* functions, satisfying $f(x_1) \leq f(x_2)$, whenever $x_1 \preceq x_2$.

Monotone functions have been well-studied on multiple fronts in computational complexity due to their natural structure. They have been celebrated for decades in the property testing literature [GGL⁺00, DGL⁺99, FLN⁺02, BCGM12, CS13b, CS13a, CS14], where we have recently witnessed ultimate results [KMS15, CDST15, BB16], in the circuit complexity literature, where we now have strong lower bounds [RW92, Raz85], and in computational learning, where we now have learning algorithms in numerous learning models [BT96, Ang87, KV94, Ser04, OS07, OW09].

The generalized notion of k -monotonicity has also been studied in the context of circuit lower bounds for more than 50 years. In particular, Markov [Mar57] showed that any k -monotone function (even with multiple outputs) can be computed using circuits containing only $\log k$ negation gates. The presence of negation gates appears to be a challenge in proving circuit lower bounds: “the effect of such gates on circuit size remains to a large extent a mystery” [Juk12]. The recent results of Blais *et al.* [BCO⁺15] on learning lower bounds have prompted renewed interest in understanding k -monotone functions from multiple angles, including cryptography, circuit complexity, learning theory, and Fourier analysis ([Ros15, GMOR15, GK15, LZ16]).

Motivated by the exponential lower bounds on PAC learning k -monotone functions due to [BCO⁺15], we initiate the study of k -monotonicity in the closely related *Property Testing* model. In this model, given query access to a function, one must decide if the function is k -monotone, or is far from being k -monotone, by querying the input only in a small number of places.

1.1 Our results

We focus on testing k -monotonicity of Boolean functions defined over the d -dimensional hypergrid $[n]^d$, and the hypercube $\{0, 1\}^d$. We begin our presentation with the results for the hypercube, in order to build intuition into the difficulty of the problem, while comparing our results with the current literature on testing monotonicity. Our stronger results concern the hypergrid $[n]^d$.

1.1.1 Testing k -monotonicity on the hypercube $\{0, 1\}^d$

In light of the recent results of [KMS15] that provide a $O(\sqrt{d})$ -query tester for monotonicity, we first show that testing k -monotonicity is strictly harder than testing monotonicity on $\{0, 1\}^d$, for $k \geq 3$.

Theorem 1.1. *For $1 \leq k \leq d^{1/4}/2$, any one-sided non-adaptive tester for k -monotonicity of functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$ must make $\Omega(d/k^2)^{k/4}$ queries.*

Both **Theorem 1.1** and its proof generalize the $\Omega(d^{1/2})$ lower bound for testing monotonicity, due to Fischer *et al.* [FLN⁺02].

On the upper bounds side, while the monotonicity testing problem is providing numerous potential techniques for approaching this new problem [GGL⁺00, DGL⁺99, CS13a, BCGM12, CST14, KMS15], most common techniques appear to resist generalizations to k -monotonicity. However, our upper bounds demonstrate a separation between testing and PAC learning k -monotonicity, for large enough values of $k = \omega(\log d)$.

Theorem 1.2. *There exists a one-sided non-adaptive tester for k -monotonicity of functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$ with query complexity $q(d, \varepsilon, k) = 2^{O(\sqrt{d} \cdot \log d \cdot \log \frac{1}{\varepsilon})}$.*

Indeed, in the related PAC learning model, [BCO⁺15] shows that learning k -monotone functions on the hypercube requires $2^{\Omega(k \cdot \sqrt{d} \cdot 1/\varepsilon)}$ many queries.

We further observe that the recent non-adaptive and adaptive 2-sided lower bounds of [CDST15, BB16], imply the same bounds for k -monotonicity, using black box reductions. We summarize the state of the art for testing k -monotonicity on the hypercube in Table 1.

	upper bound	1.s.-n.a. lower bound	2.s.-n.a. lower bound	2.s.-a. lower bound
$k = 1$	$O(\sqrt{d})$ [KMS15]	$\Omega(d^{1/2})$ [FLN ⁺ 02]	$\Omega(d^{1/2-o(1)})$ [CDST15]	$\tilde{\Omega}(d^{1/4})$ [BB16]
$k \geq 2$	$d^{O(k\sqrt{d})}$ [BCO ⁺ 15], $d^{O(\sqrt{d})}$ Thm 1.2	$\Omega(d/k^2)^{k/4}$ Thm 1.1 ($k = O(d^{1/4})$)	$\Omega(d^{1/2-o(1)})$ Cor 4.6	$\tilde{\Omega}(d^{1/4})$ Cor 4.6

Table 1: Testing k -monotonicity of a function $f: \{0, 1\}^d \rightarrow \{0, 1\}$

1.1.2 Testing k -monotonicity on the hypergrid $[n]^d$

The remainder of the paper focuses on functions defined over the d -dimensional hypergrid domain $[n]^d$, where we denote by $(i_1, i_2, \dots, i_d) \preceq (j_1, j_2, \dots, j_d)$ the partial order in which $i_1 \leq j_1, i_2 \leq j_2, \dots, i_d \leq j_d$. Testing monotonicity has received a lot of attention over the d -dimensional hypergrids [GGL⁺00, EKK⁺00, Fis04, BRW05, AC06, HK08, BBM12, CS13b, CS13a, CS14, BRY14], where the problem is well-understood, and we refer the reader to Table 4 in the appendix for a detailed review on the state of the art in the area. We summarize our results on testing k -monotonicity over $[n]^d$ in Table 2.

	General k	$k = 2$	$k = 1$ (monotonicity)
$d = 1$	$\Theta(\frac{k}{\varepsilon})$ 1.s.-n.a., $\tilde{O}(\frac{1}{\varepsilon^7})$ 2.s.-n.a.	$O(\frac{1}{\varepsilon})$ 1.s.-n.a.	$\Theta(\frac{1}{\varepsilon})$ 1.s.-n.a.
$d = 2$	$\tilde{O}(\frac{k^2}{\varepsilon^3})$ 2.s.-n.a. (from below)	$\Theta(\frac{1}{\varepsilon})$ 2.s.-a.	$\Theta(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ 1.s.-n.a., $\Theta(\frac{1}{\varepsilon})$ 1.s.-a.
$d \geq 3$	$\tilde{O}(\frac{1}{\varepsilon^2} (\frac{5kd}{\varepsilon})^d)$ 2.s.-n.a., $2^{\tilde{O}(k\sqrt{d}/\varepsilon^2)}$ 2.s.-n.a.	$\tilde{O}(\frac{1}{\varepsilon^2} (\frac{10d}{\varepsilon})^d)$ 2.s.-n.a., $2^{\tilde{O}(\sqrt{d}/\varepsilon^2)}$ 2.s.-n.a.	$O(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon})$ 1.s.-n.a.

Table 2: Summary of our results: testing k -monotonicity of a function $f: [n]^d \rightarrow \{0, 1\}$ (first two columns). The last column contains known bounds on monotonicity testing (see Appendix A) and is provided for comparison.

Testing k -monotonicity on the line and the 2-dimensional grid We begin with a study of functions $f: [n] \rightarrow \{0, 1\}$. As before, note that 1-sided tests should always accept k -monotone functions, and so, they must accept unless they discover a violation to k -monotonicity in the form of a sequence $x_1 \preceq x_2 \preceq \dots \preceq x_{k+1}$ in $[n]^d$, such that $f(x_1) = 1$ and $f(x_i) \neq f(x_{i+1})$. Therefore, lower bounds for 1-sided k -monotonicity testing must grow at least linearly with k . We show that this is indeed the case for both adaptive and non-adaptive tests, and moreover, we give a tight non-adaptive algorithm. Consequently, our results demonstrate that adaptivity does not help in testing k -monotonicity with one-sided error on the line domain.

Theorem 1.3. *Any one-sided (possibly adaptive) tester for k -monotonicity of functions $f: [n] \rightarrow \{0, 1\}$ must have query complexity $\Omega\left(\frac{k}{\varepsilon}\right)$.*

The upper bound generalizes the $O(1/\varepsilon)$ tester for monotonicity on the line.

Theorem 1.4. *There exists a one-sided non-adaptive tester for k -monotonicity of functions $f: [n] \rightarrow \{0, 1\}$ with query complexity $q(n, \varepsilon, k) = O\left(\frac{k}{\varepsilon}\right)$.*

Testing with 2-sided error, however, does not require a dependence on k . In fact the problem has been well-studied in the machine learning literature in the context of testing/learning “union of intervals” [KR00, BBY12], and in testing geometric properties, in the context of testing surface area [KNOW14, Nee14],¹ resulting in an $O(1/\varepsilon^{7/2})$ -query algorithm. Namely, the starting point of [BBY12] (later improved by [KNOW14]) is a “Buffon Needle’s”-type argument, where the crucial quantity to analyze is the noise sensitivity of the function, that is the probability that a randomly chosen pair of nearby points cross a “boundary” – i.e., have different values. (Moreover, the algorithm of [BBY12] works in the *active testing* setting: it only requires a weaker access model than the standard query model).

We provide an alternate proof of a $\text{poly}(1/\varepsilon)$ bound (albeit with a worse exponent) that reveals a surprising connection with *distribution testing*, namely with the problem of estimating the support size of a distribution.

Theorem 1.5. *There exists a two-sided non-adaptive tester for k -monotonicity of functions $f: [n] \rightarrow \{0, 1\}$ with query complexity $q(n, \varepsilon, k) = \tilde{O}(1/\varepsilon^7)$, independent of k .*

An immediate implication of [Theorem 1.5](#) is that one can test even $n^{1-\alpha}$ -monotonicity of $f: [n] \rightarrow \{0, 1\}$, for every $\alpha > 0$, with a constant number of queries. Hence, there is a separation between 1-sided and 2-sided testing, for $k = \omega(1)$.

Turning to the 2-dimensional grid, we show that 2-monotone functions can be tested with the minimum number of queries one could hope for:

Theorem 1.6. *There exists a two-sided adaptive tester for 2-monotonicity of functions $f: [n]^2 \rightarrow \{0, 1\}$ with query complexity $q(n, \varepsilon) = O\left(\frac{1}{\varepsilon}\right)$.*

We also discuss possible generalizations of [Theorem 1.6](#) to general k or d in [Section 6.2](#).

Testing k -monotonicity on $[n]^d$, tolerant testing, and distance approximation Moving to the general grid domain $[n]^d$, we show that k -monotonicity is testable with $\text{poly}(1/\varepsilon, k)$ queries in constant-dimension grids.

Theorem 1.7. *There exists a non-adaptive tester for k -monotonicity of functions $f: [n]^d \rightarrow \{0, 1\}$ with query complexity $q(n, d, \varepsilon, k) = \min\left(\tilde{O}\left(\frac{1}{\varepsilon^2} \left(\frac{5kd}{\varepsilon}\right)^d\right), 2^{\tilde{O}(k\sqrt{d}/\varepsilon^2)}\right)$.*

In fact, we obtain more general testing algorithms than in [Theorem 1.7](#), namely our results hold for *tolerant* testers (as we define next).

The notion of tolerant testing was first introduced in [PRR06] to account for the possibility of noisy data. In this notion, a test should accept inputs that are ε_1 -close to the property, and reject

¹We thank Eric Blais for mentioning the connection, and pointing us to these works.

inputs that are ε_2 -far from the property, where ε_1 and ε_2 are given parameters. Tolerant testing is intimately connected to the notion of distance approximation: given tolerant testers for every $(\varepsilon_1, \varepsilon_2)$, there exists an algorithm that estimates the distance to the property within any (additive) ε , while incurring only a $\tilde{O}\left(\log \frac{1}{\varepsilon}\right)$ factor blow up in the number of queries. Furthermore, [PRR06] shows that both tolerant testing and distance approximation are no harder than agnostic learning. We prove the following general result.

Theorem 1.8. *There exist*

- a non-adaptive (fully) tolerant tester for k -monotonicity of functions $f: [n]^d \rightarrow \{0, 1\}$ with query complexity $q(n, d, \varepsilon_1, \varepsilon_2, k) = \tilde{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \left(\frac{5kd}{\varepsilon_2 - \varepsilon_1}\right)^d\right)$;
- a non-adaptive tolerant tester for k -monotonicity of functions $f: [n]^d \rightarrow \{0, 1\}$ with query complexity $q(n, d, \varepsilon_1, \varepsilon_2, k) = 2^{\tilde{O}(k\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2)}$, under the restriction that $\varepsilon_2 > 3\varepsilon_1$.

To the best of our knowledge, the only previous results for tolerant testing for monotonicity on $[n]^d$ are due to Fattal and Ron [FR10]. They give both additive and multiplicative distance approximations algorithms, and obtain $O(d)$ -multiplicative and ε -additive approximations with query complexity $\text{poly}\left(\frac{1}{\varepsilon}\right)$. While very efficient, these results only give fully tolerant testers for dimensions $d = 1$ and $d = 2$. Our results generalize the work of [FR10] showing existence of tolerant testers for k -monotonicity (and hence for monotonicity) for any dimension $d \geq 1$, and any $k \geq 1$, but paying the price in the query complexity.

As a consequence to [Theorem 1.8](#), we make progress on an open problem of Berman *et al.* [BRY14], as explained next.

Testing k -monotonicity under L_p distance The property of being a monotone Boolean function has a natural extension to real-valued functions. Indeed, a real-valued function defined over a finite domain D is monotone if $f(x) \leq f(y)$ whenever $x \preceq y$. For real-valued functions the more natural notion of distance is L_p distance, rather than Hamming distance. The study of monotonicity has been extended to real-valued functions in a recent work by Berman *et al.* [BRY14]. They give tolerant testers for grids of dimension $d = 1$ and $d = 2$, and leave open the problem of extending the results to general d , as asked explicitly at the recent Sublinear Algorithms Workshop 2016 [Sub16].

We make progress towards solving this open problem, by combining our [Theorem 1.8](#) with a reduction from L_p testing to Hamming testing inspired by [BRY14]. This reduction relates L_1 -distance of a function $f: [n]^d \rightarrow [0, 1]$ to monotonicity to Hamming distance to monotonicity of a “rounded” function $\tilde{f}: [n]^d \times [m] \rightarrow \{0, 1\}$, essentially trading the range for an extra dimension (where m is a rounding parameter to be suitably chosen). Moreover, simulating query access to \tilde{f} can be performed efficiently given query access to f .

Theorem 1.9. *There exists a non-adaptive tolerant L_1 -tester for monotonicity of functions $f: [n]^d \rightarrow [0, 1]$ with query complexity*

- $\tilde{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \left(\frac{5d}{\varepsilon_2 - \varepsilon_1}\right)^d\right)$, for any $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$;
- $2^{\tilde{O}(\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2)}$, for any $0 \leq 3\varepsilon_1 < \varepsilon_2 \leq 1$.

1.2 Proof overview and technical contribution

Structural properties and the separation between testing and learning on $\{0, 1\}^d$. We first observe that basic structural properties, such as *extendability* (i.e. the feature that a function that is monotone on a sub-poset of $[n]^d$ can be extended into a monotone function on the entire poset domain), and properties of the *violation graph* (i.e., the graph whose edges encode the violations to monotonicity), extend easily to k -monotonicity (see [Section 3](#)). These properties help us to argue the separation between testing and learning ([Theorem 1.2](#)). However, unlike the case of monotonicity testing, these properties do not seem to be enough for showing upper bounds that grow polynomially in d .

Grid coarsening and testing by implicit/explicit learning. One pervading technique, which underlies all the hypergrid upper bounds in this work, is that of *gridding*: i.e., partitioning the domain into “blocks” whose size no longer depends on the main parameter of the problem, n . This technique generalizes the approach of [\[FR10\]](#) who performed a similar gridding for dimension $d = 2$. By simulating query access to the “coarsened” version of the unknown function (with regard to these blocks), we are able to leverage methods such as testing-by-learning (either fully or partially learning the function), or reduce our testing problem to a (related) question on these nicer “coarsenings.” (The main challenge here lies in providing efficient and consistent oracle access to the said coarsenings.)

At a high-level, the key aspect of k -monotonicity which makes this general approach possible is reminiscent of the concept of *heredity* in property testing. Specifically, we rely upon the fact that “gridding preserves k -monotonicity:” if f is k -monotone, then so will be its coarsening g – but now g is much simpler to handle. This allows us to trade the domain $[n]^d$ for what is effectively $[m]^d$, with $m \ll n$. We point out that this differs from the usual paradigm of *dimension reduction*: indeed, the latter would reduce the study of a property of functions on $[n]^d$ to that of functions on $[n]^{d'}$ for $d' \ll d$ (usually even $d' = 1$) by projecting f on a lower-dimensional domain. In contrast, we do not take the dimension down, but instead reduce the size of the *alphabet*. Moreover, it is worth noting that this gridding technique is also orthogonal to that of *range reduction*, as used e.g. in [\[DGL⁺99\]](#). Indeed, the latter is a reduction of the range of the function from $[R]$ to $\{0, 1\}$, while gridding is only concerned about the domain size.

Estimating the support of distributions. Our proof of the $\text{poly}(1/\varepsilon)$ upper bound for testing k -monotonicity on the line ([Theorem 1.5](#)) rests upon an unexpected connection to *distribution testing*, namely to the question of support size estimation of a probability distribution. In more detail, we describe how to reduce k -monotonicity testing to the support size estimation problem in (a slight modification of) the *Dual access model* introduced by Canonne and Rubinfeld [\[CR14\]](#), where the tester is granted samples from an unknown distribution as well as query access to its probability mass function.

For our reduction to go through, we first describe how any function $f: [n] \rightarrow \{0, 1\}$ determines a probability distribution D_f (on $[n]$), whose effective support size is directly related to the k -monotonicity of f . We then show how to implement dual access to this D_f from queries to f : in order to avoid any dependence on k and n in this step, we resort both to the gridding approach outlined above (allowing us to remove n from the picture) and to a careful argument to “cap” the values of D_f returned by our simulated oracle. Indeed, obtaining the exact value of $D_f(x)$ for arbitrary x may require $\Omega(k)$ queries to f , which we cannot afford; instead, we argue that

only returning $D_f(x)$ whenever this value is “small enough” is sufficient. Finally, we show that implementing this “capped” dual access oracle is possible with no dependence on k whatsoever, and we can now invoke the support size estimation algorithm of [CR14] to conclude.

Fourier analysis on the hypergrid. We give an algorithm for fully tolerantly testing k -monotonicity whose query complexity is exponential in d . We also describe an alternate tester (with a slightly worse tolerance guarantee) whose query complexity is instead exponential in $\tilde{O}(k\sqrt{d})$ for constant distance parameters. As mentioned above, we use our gridding approach combined with tools from learning theory. Specifically, we employ an agnostic learning algorithm of [KKMS08] using polynomial regression. Our coarsening methods allow us to treat the domain as if it were $[m]^d$ for some m that is independent of n . To prove that this agnostic learning algorithm will succeed, we turn to Fourier analysis over $[m]^d$. We extend the bound on average sensitivity of k -monotone functions over the Boolean hypercube from [BCO⁺15] to the hypergrid, and we show that this result implies that the Fourier coefficients are concentrated on “simple” functions.

1.3 Discussion and open problems

This is the first work to study k -monotonicity, a natural and well-motivated generalization of monotonicity. Hence this work opens up many intriguing questions in the area of property testing, with potential applications to learning theory, circuit complexity and cryptography.

As previously mentioned, the main open problem prompted by our work is the following:

Can k -monotonicity on the hypercube $\{0, 1\}^d$ be tested with $\text{poly}(d^k)$ queries?

A natural 1-sided tester for k -monotonicity is a *chain tester*: it queries points along a random chain, and rejects only if it finds a violation to k -monotonicity, in the form of a sequence $x_1 \preceq x_2 \preceq \dots \preceq x_{k+1}$ in $\{0, 1\}^d$, such that $f(x_1) = 1$ and $f(x_i) \neq f(x_{i+1})$. In particular, the testers in [GGL⁺00, CS13a, CST14, KMS15] all directly imply a chain tester. We conjecture that there exists a chain tester for k -monotonicity that succeeds with probability $d^{-O(k)}$.

Another important open question concerns the hypergrid domain, and in particular it pushes for a significant strengthening of [Theorem 1.7](#) and [Theorem 1.9](#):

Can k -monotonicity on the hypergrid $[n]^d$ be (tolerantly) tested with $2^{o_k(\sqrt{d})}$ queries?

Answering this question would imply further progress on the L_1 -testing question for monotonicity, left open in [BRY14, Sub16].

There also remains the question of establishing two-sided lower bounds that would go beyond those of monotonicity. Specifically:

Is there an $d^{\Omega(k)}$ -query two-sided lower bound for k -monotonicity on the hypercube $\{0, 1\}^d$?

In this work we also show surprising connections to distribution testing (e.g. in the proof of [Theorem 1.5](#)), and to testing union of intervals and testing surface area (as discussed in [Section 5.2](#)). An intriguing direction is to generalize this connection to union of intervals and surface area in higher dimensions, to leverage or gain insight on k -monotonicity on the d -dimensional hypergrid.

Finally, while we only stated here a few directions, we emphasize that every question that is relevant to monotonicity is also relevant and interesting in the case of k -monotonicity.

1.4 Related work

As mentioned, k -monotonicity has deep connections with the notion of *negation complexity* of functions, which is the minimum number of negation gates needed in a circuit to compute a given function. The power of negation gates is intriguing and far from being understood in the context of circuit lower bounds. Quoting from Jukna’s book [Juk12], *the main difficulty in proving nontrivial lower bounds on the size of circuits using AND, OR, and NOT is the presence of NOT gates: we already know how to prove even exponential lower bounds for monotone functions if no NOT gates are allowed. The effect of such gates on circuit size remains to a large extent a mystery.*

This gap has motivated the study of circuits with *few* negations. Two notable works successfully extend lower bounds in the monotone setting to negation-limited setting: in [AM05], Amano and Maruoka show superpolynomial circuit lower bounds for $(1/6) \log \log n$ negations using the CLIQUE function; and recently the breakthrough work of Rossman [Ros15] establishes circuit lower bounds for NC^1 with roughly $\frac{1}{2} \log n$ negations by drawing upon his lower bound for monotone NC^1 .

The divide between the understanding of monotone and non-monotone computation exists in general: while we usually have a fairly good understanding of the monotone case, many things get murky or fail to hold even when a single negation gate is allowed. In order to get a better grasp on negation-limited circuits, a body of recent work has been considering this model in various contexts: Blais *et al.* [BCO⁺15] study negation-limited circuits from a computational learning viewpoint, Guo *et al.* [GMOR15] study the possibility of implementing cryptographic primitives using few negations, and Lin and Zhang [LZ16] are interested in verifying whether some classic Boolean function conjectures hold for the subset of functions computed by negation-limited circuits.

Many of these results implicitly or explicitly rely on a simple but powerful tool: the decomposition of negation-limited circuits into a composition of some “nice” function with monotone components. Doing so enables one to apply results on separate monotone components, and finally to carefully combine the outcomes (e.g., [GK15]). Though these techniques can yield results for as many as $O(\log n)$ negations, they also leave open surprisingly basic questions:

- [BCO⁺15] Can we have an efficient weak learning algorithm for functions computed by circuits with a *single* negation?
- [GMOR15] Can we obtain pseudorandom generators when allowing only a *single* negation?

In contexts where the circuit size is not the quantity of interest, the equivalent notion of 2-monotone functions is more natural than that of circuits allowing only one negation. Albeit seemingly simple, even the class of 2-monotone functions remains largely a mystery: as exemplified above, many basic yet non-trivial questions, ranging from the structure of their Fourier spectrum to their expressive power of k -monotone functions, remain open.

1.5 Organization of the paper

After recalling some notations and definitions in Section 2 and providing some insight into the structural properties of functions far from k -monotonicity in Section 3, we consider the case of the Boolean hypercube in Section 4, where we establish lower bounds on testing k -monotonicity of functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$ for both one- and two-sided algorithms, and provide an algorithm which “beats” the testing-by-learning approach, showing that testing is provably easier than learning.

Next, we establish our results for functions on the line in Section 5, starting with the lower and upper bounds for one-sided testers before turning in Section 5.2 to the two-sided upper bound

of [Theorem 1.3](#). We then describe in [Section 6](#) our results for functions on the grid $[n]^2$, focusing on the case $k = 2$; and discussing possible extensions in [Section 6.2](#).

[Section 7](#) contains our general algorithms for k -monotonicity on the hypergrid $[n]^d$, for arbitrary k and d . We prove [Theorem 1.8](#) in two parts. We establish its first item (general tolerant testing algorithm with exponential dependence in d) in [Section 7.1](#) ([Proposition 7.2](#)). The second item (with query complexity exponential in $k\sqrt{d}$) is proven in [Section 7.2](#), where we analyze the Fourier-based tolerant tester of [Proposition 7.11](#). We then apply these results to the question of tolerant L_1 -testing of monotonicity in [Section 8](#), after describing a reduction between monotonicity of functions $[n]^d \rightarrow [0, 1]$ and of $[n]^{d+1} \rightarrow \{0, 1\}$.

Except maybe [Section 8](#) which depends on [Section 7](#), all sections are independent and self-contained, and the reader may choose to read them in any order.

2 Preliminaries

We denote by \log the binary logarithm, and use $\tilde{O}(\cdot)$ to hide polylogarithmic factors in the argument (so that $\tilde{O}(f) = O(f \log^c f)$ for some $c \geq 0$).

Given two functions $f, g: \mathcal{X} \rightarrow \mathcal{Y}$ on a finite domain \mathcal{X} , we write $\text{dist}(f, g)$ for the (normalized) Hamming distance between them, i.e.

$$\text{dist}(f, g) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \mathbb{1}_{\{f(x) \neq g(x)\}} = \Pr_{x \sim \mathcal{X}} [f(x) \neq g(x)]$$

where $x \sim \mathcal{X}$ refers to x being drawn from the uniform distribution on \mathcal{X} . A *property* of functions from \mathcal{X} to \mathcal{Y} is a subset $\mathcal{P} \subseteq \mathcal{X}^{\mathcal{Y}}$ of these functions; we define the distance of a function f to \mathcal{P} as the minimum distance of f to any $g \in \mathcal{P}$:

$$\text{dist}(f, \mathcal{P}) = \inf_{g \in \mathcal{P}} \text{dist}(f, g).$$

For some of our applications, we will also use another notion of distance specific to real-valued functions, the L_1 distance (as introduced in the context of property testing in [\[BRY14\]](#)). For $f, g: \mathcal{X} \rightarrow [0, 1]$, we write

$$L_1(f, g) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} |f(x) - g(x)| = \mathbb{E}_{x \sim \mathcal{X}} [|f(x) - g(x)|] \in [0, 1]$$

and extend the definition to $L_1(f, \mathcal{P})$, for $\mathcal{P} \subseteq \mathcal{X}^{[0,1]}$, as before.

Property testing. We recall the standard definition of testing algorithms, as well as some terminology:

Definition 2.1. Let \mathcal{P} be a property of functions from \mathcal{X} to \mathcal{Y} . A *q-query testing algorithm for \mathcal{P}* is a randomized algorithm \mathcal{T} which takes as input $\varepsilon \in (0, 1]$ as well as query access to a function $f: \mathcal{X} \rightarrow \mathcal{Y}$. After making at most $q(\varepsilon)$ queries to the function, \mathcal{T} either outputs ACCEPT or REJECT, such that the following holds:

- if $f \in \mathcal{P}$, then \mathcal{T} outputs ACCEPT with probability at least $2/3$; (Completeness)
- if $\text{dist}(f, \mathcal{P}) \geq \varepsilon$, then \mathcal{T} outputs REJECT with probability at least $2/3$; (Soundness)

where the probability is taken over the algorithm’s randomness. If the algorithm only errs in the second case but accepts any function $f \in \mathcal{P}$ with probability 1, it is said to be a *one-sided* tester; otherwise, it is said to be *two-sided*. Moreover, if the queries made to the function can only depend on the internal randomness of the algorithm, but not on the values obtained during previous queries, it is said to be *non-adaptive*; otherwise, it is *adaptive*.

Additionally, we will also be interested in *tolerant* testers – roughly, algorithms robust to a relaxation of the first item above:

Definition 2.2. Let \mathcal{P} , \mathcal{X} , and \mathcal{Y} be as above. A q -query tolerant testing algorithm for \mathcal{P} is a randomized algorithm \mathcal{T} which takes as input $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$, as well as query access to a function $f: \mathcal{X} \rightarrow \mathcal{Y}$. After making at most $q(\varepsilon_1, \varepsilon_2)$ calls to the oracle, \mathcal{T} outputs either ACCEPT or REJECT, such that the following holds:

- if $\text{dist}(f, \mathcal{P}) \leq \varepsilon_1$, then \mathcal{T} outputs ACCEPT with probability at least $2/3$; (Completeness)
- if $\text{dist}(f, \mathcal{P}) \geq \varepsilon_2$, then \mathcal{T} outputs REJECT with probability at least $2/3$; (Soundness)

where the probability is taken over the algorithm’s randomness. The notions of one-sidedness and adaptivity of [Definition 2.1](#) extend to tolerant testing algorithms as well.

Note that as stated, in both cases the algorithm “knows” \mathcal{X}, \mathcal{Y} , and \mathcal{P} ; so that the query complexity q can be parameterized by these quantities. More specifically, when considering $\mathcal{X} = [n]^d$ and the property \mathcal{P} of k -monotonicity, we will allow q to depend on n, d , and k . Finally, we shall sometimes require a probability of success $1 - \delta$ instead of the (arbitrary) constant $2/3$; by standard techniques, this can be obtained at the cost of a multiplicative $O(\log(1/\delta))$ in the query complexity.

PAC and agnostic learning [[Val84](#)] A learning algorithm \mathcal{A} for a *concept class* \mathcal{C} of functions $f: \mathcal{X} \rightarrow \mathcal{Y}$ (under the uniform distribution) is given parameters $\varepsilon, \delta > 0$ and sample access to some target function $f \in \mathcal{C}$ via labeled samples $\langle x, f(x) \rangle$, where x is drawn uniformly at random from \mathcal{X} . The algorithm should output a hypothesis $h: \mathcal{X} \rightarrow \mathcal{Y}$ such that $\text{dist}(h, f) \leq \varepsilon$ with probability at least $1 - \delta$. The algorithm is *efficient* if it runs in time $\text{poly}(n, 1/\varepsilon, 1/\delta)$. If \mathcal{A} must output $h \in \mathcal{C}$ we say it is a *proper learning algorithm*, otherwise, we say it is an *improper learning* one.

Moreover, if \mathcal{A} still succeeds when f does not actually belong to \mathcal{C} , we say it is an *agnostic learning algorithm*. Specifically, the hypothesis function h that it outputs must satisfy $\text{dist}(f, h) \leq \text{OPT}_f + \varepsilon$ with probability at least $1 - \delta$, where $\text{OPT}_f = \min_{g \in \mathcal{C}} \text{dist}(f, g)$.

3 Structural results

In this section, we will prove that the distance to k -monotonicity of a Boolean function f can be expressed in a combinatorial way – which does not require measuring the distance between f and the closest k -monotone function to f . We will prove this for a general finite poset domain building up on the ideas of [[FLN⁺02](#)]. In the rest of this section, we denote by $\mathcal{P} = (V, \preceq)$ an arbitrary poset, the underlying domain of the function.

Definition 3.1. We define the *violated pattern* K as the sequence of alternating bits $K = (b_1, b_2, \dots, b_k, b_{k+1})$ of length $(k + 1)$, where $b_1 = 1$ and all the bits in the sequence alternate, i.e., $b_i \neq b_{i+1} \forall i \in [k]$.

A function $f: \mathcal{P} \rightarrow \{0, 1\}$ is k -monotone only if it avoids K . That is, for any $x_1 \prec x_2 \prec \dots \prec x_{k+1} \in \mathcal{P}$ we have $f(x_i) \neq K(i)$ for some $i \in [k+1]$.

Using insights from the literature on monotonicity testing, we show that functions far from k -monotonicity have a large matching of “violated hyperedges” in the “violation hypergraph” which we define shortly. Let us recall the definition of “violation graph” which has been extremely useful with monotonicity testing as seen in [EKK⁺00, PRR06, HK08, ACCL07, FLN⁺02].

Definition 3.2 (Violation graph). Given a function $f: \mathcal{P} \rightarrow \{0, 1\}$, the *violation graph* of f is defined as $G_{\text{viol}}(f) = (\mathcal{P}, E(G_{\text{viol}}))$ where $(x, y) \in E(G_{\text{viol}})$ if $x, y \in \mathcal{P}$ form a monotonicity violating pair in f – that is $x \preceq y$ but $f(x) > f(y)$.

The following theorem, proved in [FLN⁺02], about violation graphs has been extremely useful in monotonicity testing literature.

Theorem 3.3. *Let $f: \mathcal{P} \rightarrow \{0, 1\}$ be a function that is ε -far from monotone. Then, there exists a matching of edges in the violation graph for f of size at least $\varepsilon|\mathcal{P}|/2$.*

Now let us define a generalization of this concept, the violation hypergraph.

Definition 3.4 (Violation hypergraph). Given a function $f: \mathcal{P} \rightarrow \{0, 1\}$, the *violation hypergraph* of f is $H_{\text{viol}}(f) = (\mathcal{P}, E(H_{\text{viol}}))$ where $(x_1, x_2, \dots, x_k) \in E(H_{\text{viol}})$ if the ordered $(k+1)$ -tuple $x_1 < x_2 < \dots < x_{k+1}$ (which is a $(k+1)$ -uniform hyperedge) forms a violation to k -monotonicity in f .

Now, we state the main theorem we prove in this section. This theorem offers an alternate characterization of distance to k -monotonicity that we seek. We recall that a set of edges forms a matching in a hypergraph if any pair of hyperedges is pairwise disjoint.

Theorem 3.5. *Let $f: \mathcal{P} \rightarrow \{0, 1\}$ be function that is ε -far from k -monotone. Then, there exists a matching of $(k+1)$ -uniform hyperedges of size at least $\frac{\varepsilon|\mathcal{P}|}{k+1}$ in the violation hypergraph.*

To prove this theorem we first exploit the key notion of *extendability* (also explored in the context of testing monotonicity in [FLN⁺02]) which we define below. Later we will show that k -monotone functions are extendable.

Definition 3.6 (Extendability). A property of Boolean functions is said to be *extendable* over a poset domain if the following holds for any set $X \subseteq \mathcal{P}$: given a function $f: X \rightarrow \{0, 1\}$ which has the property (on X), it is possible to define a function $g: \mathcal{P} \rightarrow \{0, 1\}$ such that $g(x) = f(x), \forall x \in X$ and g has the property.

In other words, a property is extendable if for any subset $X \subseteq \mathcal{P}$, given a function defined over the set X which respects the property, it is possible to fill in values outside X such that the new function obtained continues to respect the property. Next, we show that k -monotonicity is an extendable property:

Lemma 3.7. *k -monotonicity is an extendable property.*

Proof of Lemma 3.7. Consider $X \subseteq \mathcal{P}$ and some function $f: X \rightarrow \{0, 1\}$ which is k -monotone over X . That is, for any $x_1 < x_2 < \dots < x_{k+1} \in X$ there exists $i \in [k+1]$ such that $f(x_i) \neq K[i]$. Take a minimal point $v \in \mathcal{P} \setminus X$. That is, for any other point $v' \in \mathcal{P} \setminus X$ either $v \leq v'$ or v and v' are not comparable. We will use the following result:

Claim 3.8. *There exists a function $g: X \cup \{v\} \rightarrow \{0, 1\}$ such that $g(x) = f(x)$ for all $x \in X$, and g respects k -monotonicity over its domain.*

Before proving this claim, we show how it implies [Lemma 3.7](#). Namely, starting with any function $f: X \rightarrow \{0, 1\}$ which is k -monotone on its domain X , we just keep applying the [Claim 3.8](#) inductively until we get a function defined over the entire poset which respects k -monotonicity.

Proof of Claim 3.8. We will show this by contradiction. Suppose there is no good assignment available for $g(v)$, that is that both the choices $g(v) = 0$ and $g(v) = 1$ lead to a violation to k -monotonicity in g . Consider the choice $g(v) = 0$. Since this results in a violation to k -monotonicity, we know that there is a path $P_0 = (x_1 \prec x_2 \prec \dots \prec x_{k+1})$ which is a violation to k -monotonicity. It is clear that $v \in P_0$; let i be such that $x_i = v$. Similarly, there is path $P_1 = (y_1 \prec y_2 \prec \dots \prec y_{k+1})$ corresponding to $g(v) = 1$ which also contains the violated pattern, and some j such that $y_j = v$. And thus, $g(x_t) = g(y_t)$, for all $t \in [k+1]$ (as both of the paths indexed by x and y form a violation to k -monotonicity).

By the above discussion 2 paths, P_0 and P_1 , meet at v . We will see that one of the two paths

$$P'_0 = (x_1 \prec x_2 \prec \dots \prec x_{i-1} \prec y_i \prec y_{i+1} \prec \dots \prec y_{k+1})$$

or

$$P'_1 = (y_1 \prec y_2 \prec \dots \prec y_{j-1} \prec x_j \prec x_{j+1} \prec \dots \prec x_{k+1})$$

is already a violation to k -monotonicity in f . To see this, let us begin by recalling that we let v be the i^{th} vertex on P_0 and the j^{th} vertex on P_1 . Now it is clear that $i \neq j$. Without loss of generality, suppose $i < j$. In this case, the evaluations of f along path P'_1 form the violated pattern. This is because the function values alternate along the segment $(y_1 \prec y_2 \prec \dots \prec y_{j-1})$. Also, the function values alternate along the segment $(x_i \prec x_{i+1} \prec x_{i+2} \prec \dots \prec x_{k+1})$. And finally note that since $f(y_{j-1}) \neq f(y_j)$ and $f(y_j) = f(x_j)$ we get that $f(y_{j-1}) \neq f(x_j)$ as well. So, the path P'_1 indeed contains a violation to k -monotonicity as claimed. The other case, $i > j$, is analogous. Hence the claim follows. \square

\square

In the next lemma, we show that there is a nice characterization of distance to k -monotonicity in terms of the size of the *minimum vertex cover* of the violation hypergraph.²

Lemma 3.9. *Let \mathcal{M}_k denote the set of k -monotone functions over the poset \mathcal{P} , and $f: \mathcal{P} \rightarrow \{0, 1\}$. Then $\text{dist}(f, \mathcal{M}_k) = \varepsilon_f$ if, and only if, the size of the minimum vertex cover in $H_{\text{viol}}(f)$ is $\varepsilon_f |\mathcal{P}|$.*

Proof of Lemma 3.9. We establish separately the two inequalities.

Claim 3.10. $\text{dist}(f, \mathcal{M}_k) \geq |VC_{\min}(H_{\text{viol}})|$

Proof. Suppose the distance to k -monotonicity is ε_f , and let g be a k -monotone function achieving it, so that $\text{dist}(f, g) = \varepsilon_f$. Define $X = \{x \in \mathcal{P} : f(x) \neq g(x)\}$ (thus, $|X| = \varepsilon_f |\mathcal{P}|$). Let us consider the violation hypergraph for f given as $H_{\text{viol}}(f) = (\mathcal{P}, E(H_{\text{viol}}))$. Now, delete vertices in X and the hyperedges containing any vertex $v \in X$ from this hypergraph. Because X is the smallest set of vertices changing values at which gives a k -monotone function, it follows that every hyperedge in $E(H_{\text{viol}})$ must contain a vertex in X . Thus, X indeed forms a vertex cover in $H_{\text{viol}}(f)$. \square

²Recall that a vertex cover in a hypergraph is just a set of vertices such that every hyperedge contains at least one of the vertices from this set.

Claim 3.11. $\text{dist}(f, \mathcal{M}_k) \leq |VC_{\min}(H_{\text{viol}})|$

Proof. Suppose the minimum vertex cover in the violation hypergraph has size $\varepsilon_f |\mathcal{P}|$. We will show that the distance of the function to k -monotonicity is ε_f . To see this, let $C \subseteq \mathcal{P}$ be the smallest vertex cover of the violation hypergraph of the said size. Observe that deleting C from $H_{\text{viol}}(f)$ removes all the hyperedges, and therefore that the function f restricted to $X = \mathcal{P} \setminus C$ is k -monotone. And by the extendability of k -monotone functions established in [Lemma 3.7](#), it follows that the function can be extended to the rest of the domain (by providing values in the cover C) such that it keeps respecting k -monotonicity. Thus, the distance to k -monotonicity is at most $|C|/|\mathcal{P}|$. \square

\square

Having characterized distance to k -monotonicity as the size of the smallest vertex cover in the violation graph, we are ready to establish [Theorem 3.5](#). To do so, we will require the following standard fact:

Fact 3.12. *Let $G = (V, E)$ be a t -uniform hypergraph. Let M be the maximum matching in G . Then, $|M| \leq VC_{\min}(G) \leq t|M|$*

Proof of Theorem 3.5. By [Lemma 3.9](#), we know that the violation hypergraph, $H_{\text{viol}}(f)$ has a minimum vertex cover of size at least $\varepsilon_f |\mathcal{P}|$. And by [Fact 3.12](#), it is seen that it contains a matching of t -uniform hyperedges of size at least $\frac{\varepsilon_f}{t} |\mathcal{P}|$. \square

4 On the Boolean hypercube

In this section, we focus on k -monotonicity of Boolean functions over the hypercube $\{0, 1\}^d$. We begin in [Section 4.1](#) with a tester with query complexity $2^{\tilde{O}(\sqrt{d})}$, establishing a strict separation between learning and testing. [Section 4.2](#) is then dedicated to our lower bounds on k -monotonicity testing.

4.1 Upper bound: beating the learning approach

In this section, we prove the following theorem:³

Theorem 1.2. *There exists a one-sided non-adaptive tester for k -monotonicity of functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$ with query complexity $q(d, \varepsilon, k) = 2^{O(\sqrt{d} \cdot \log d \cdot \log \frac{1}{\varepsilon})}$.*

Let us recall the following standard fact which we use in the proof.

Fact 4.1. *There exists an absolute constant $C > 0$ such that the number of points of $\{0, 1\}^d$ that do not have integer weights in the middle levels $[\frac{d}{2} - \sqrt{d} \log \frac{C}{\varepsilon}, \frac{d}{2} + \sqrt{d} \log \frac{C}{\varepsilon}]$ is at most $\varepsilon 2^{d-1}$.*

Proof of Theorem 1.2. Let us begin by describing the following tester. We will later see that this tester has the claimed query complexity.

- (1) Sample $O(1/\varepsilon)$ random points from the middle levels

³We note that this result is only interesting in the regime $k \leq \sqrt{d}$: indeed, for $k = \Omega(\sqrt{d} \log \frac{1}{\varepsilon})$ every function is ε -close to k -monotone.

- (2) For each of the queries in the first step, query all points with Hamming weight in the middle levels which fall in the subcube below and in the subcube above each such random point. We call each of these $O(1/\varepsilon)$ collections of queries a *superquery*.

The key idea behind the tester is that there is a big “matching of violated hyperedges” in a function far from k -monotonicity, as seen in [Theorem 3.5](#). The tester tries to find one of the “violated hyperedges”. We recall that the violation hypergraph for a function f over the hypercube has $\{0, 1\}^d$ as its vertex set. And a tuple $(x_1 < x_2 < \dots < x_{k+1})$ of $(k + 1)$ vertices forms an edge iff it is a violation to k -monotonicity.

In the rest of the analysis, we assume that k is odd.⁴ In this case, given a function f ε -far from k -monotonicity we can assume without loss of generality that f equals 0 on points with Hamming weight $< \frac{d}{2} - \sqrt{d} \log \frac{C}{\varepsilon}$, and equals 1 on points with Hamming weight $> \frac{d}{2} + \sqrt{d} \log \frac{C}{\varepsilon}$. From [Fact 4.1](#) it follows that the resulting function is still $\frac{\varepsilon}{2}$ -far from being k -monotone.

Now, by [Theorem 3.5](#), we know that there exists a matching M_f of violations to k -monotonicity in f of size $\frac{\varepsilon/2 \cdot 2^d}{k+1} = \frac{\varepsilon 2^{d-1}}{k+1}$. The set of vertices that participate in these violations has cardinality

$$(k + 1) \cdot \frac{\varepsilon 2^{d-1}}{k + 1} = \varepsilon 2^{d-1}.$$

With constant probability, in $O(1/\varepsilon)$ queries in the first step, the tester queries a vertex that belongs to some violation from M_f . Then in the next step, the tester finds this violation. Now we bound the number of queries made in a single superquery. Since it involves only $2\sqrt{d} \log \frac{C}{\varepsilon}$ levels of the cube, the number of points queried in a single superquery is no more than $b = d^{O(\sqrt{d} \log \frac{1}{\varepsilon})}$. The total query complexity of the tester can therefore be upperbounded by $b/\varepsilon = d^{O(\sqrt{d} \log \frac{1}{\varepsilon})}$. We emphasize that the number of queries made by this tester has no dependence on k . □

Remark 4.2. We can design and analyze another standalone tester with an almost identical asymptotic upper bound on query complexity. Our analysis of the performance of this tester *does not use* the extendability property. The tester makes $O(1/\varepsilon)$ iterations. In every iteration, it picks a *random seed* $x \in \{0, 1\}^n$ and queries all points on chains going through x in the truncated cube. The idea behind our analysis is a notion of “rank” for points in the cube. We say $\text{rank}(x) \geq \ell$ if there is some ascending chain ending at x which has $\geq \ell$ flips on it. Observe that if the function is indeed far, then there is at least an ε fraction of points with rank at least $k + 1$. Thus, with constant probability, some iteration hits a seed with large rank. And in the next step, the tester finds the violation. To upperbound the query complexity, note that the number of chains in the truncated cube going through any initial seed x is at most $b = d^{O(\sqrt{d} \log \frac{1}{\varepsilon})}$. This gives an upperbound of $O(b/\varepsilon)$.

4.2 Lower bounds

We now turn to lower bounds for testing k -monotonicity of Boolean functions over the hypercube $\{0, 1\}^d$. In [Section 4.2.1](#), we show that for $1 \leq k \leq d^{1/4}/2$, any one-sided non-adaptive tester for k -monotonicity requires $\Omega(d/k^2)^{k/4}$ queries, generalizing the $\Omega(\sqrt{d})$ lower bound for monotonicity

⁴The case k even is similar. In this case, one may assume that f evaluates to 0 on all points outside the middle levels.

due to Fischer *et al.* [FLN⁺02]. This bound suggests the problem become strictly harder when k increases: specifically, for $2 < k \leq d^{1/4}/2$ testing k -monotonicity requires $\omega(\sqrt{d})$ queries, while an $O(\sqrt{d})$ one-sided non-adaptive upper bound holds for monotonicity testing [KMS15].

We then describe in Section 4.2.2 a general reduction from monotonicity testing to k -monotonicity testing, for arbitrary constant k . This blackbox reduction allows us to carry any lower bound (possibly 2-sided or adaptive) for monotonicity testing to k -monotonicity. In particular, combining it with the recent lower bounds [CDST15, BB16] for 2-sided monotonicity testing, we obtain an $\Omega(d^{1/2-o(1)})$ lower bound for non-adaptive k -monotonicity testers, and an $\Omega(d^{1/4})$ lower bound for adaptive ones.

4.2.1 One-sided lower bounds

Theorem 1.1. *For $1 \leq k \leq d^{1/4}/2$, any one-sided non-adaptive tester for k -monotonicity of functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$ must make $\Omega(d/k^2)^{k/4}$ queries.*

One natural way to prove Theorem 1.1 is to extend the monotonicity lower bound in [FLN⁺02]. To this end, we consider the $\binom{d}{k}$ -sized family of functions $\{g_S: S \subseteq [d] \text{ of size } k\}$ where every single function in the family is a truncated anti-parity over some subset S of size k . By an argument analogous to the one used in [FLN⁺02], one can show $\Omega(d^{k/4})$ lower bound on query complexity. However, functions in this family are only $\Omega(2^{-k})$ -far from being k -monotone. We refine this argument to obtain functions that are $\Omega(1)$ -far from being k -monotone, for all k .

Proof of Theorem 1.1. Consider the family of functions $\{f_S: S \subseteq [d] \text{ of size } t\}$ where $t \geq k$ is a parameter to be determined later and f_S is a truncated anti-parity over t input coordinates indexed by S , namely,

$$f_S = \begin{cases} \bigoplus_{i \in S} x_i & \text{if } ||x| - d/2| \leq \sqrt{d}, \\ 0 & \text{otherwise.} \end{cases}$$

The theorem immediately follows from the two claims below.

Claim 4.3. *For any non-adaptive q -query algorithm \mathcal{A} , there exists f_S where $|S| = t$ such that \mathcal{A} reveals a violation on f_S with probability at most $q^2 \binom{2\sqrt{d}}{k} \binom{t}{k} / \binom{d}{k}$.*

Claim 4.4. *For $k \leq t \leq \sqrt{d}$, f_S is $\Omega(\sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{i} / 2^t)$ -far from any k -monotone functions.*

In particular, let $t = 4k^2$. By Claim 4.4, for $t \leq \sqrt{d}$, f_S is $\Omega(1)$ -far from any k -monotone function, and by Claim 4.3, to reject every f_S with $\Omega(1)$ probability, q needs to be at least $d^{k/4} \cdot \left(\frac{k}{2e^{2t}}\right)^{k/2} = \Omega(d/k^2)^{k/4}$.

Proof of Claim 4.3. Let Q be an arbitrary set of q queries. Without loss of generality, we assume every query z in Q has Hamming weight $|z| \in [\frac{d}{2} - \sqrt{d}, \frac{d}{2} + \sqrt{d}]$. We define $Q_{x,z} = \{y \in Q : x \preceq y \preceq z\}$ and $\text{Rej}(Q) = \{f_S : Q \text{ contains a violation for } f_S\}$. Note that $\text{Rej}(Q) = \bigcup_{(x,z): x \preceq z \in Q} \text{Rej}(Q_{x,z})$. Hence we can bound the size of $\text{Rej}(Q)$ by

$$|\text{Rej}(Q)| \leq \sum_{(x,z) \in Q^2: x \preceq z} |\text{Rej}(Q_{x,z})|. \quad (1)$$

Fix any $x \preceq z \in Q$. Because any violation for f_S in $Q_{x,z}$ contains at least two points x' and z' such that $x \preceq x' \preceq z' \preceq z$, and $x'_{S'} = 0^k$ and $z'_{S'} = 1^k$, we have $x_{S'} = 0^k$ and $z_{S'} = 1^k$ for some $S' \subseteq S$ of size k . Note that x and z differ in at most $2\sqrt{d}$ coordinates so that there are at most $\binom{2\sqrt{d}}{k}$ distinct S' on which $x_{S'}$ is 0^k and $z_{S'}$ is 1^k . Moreover, for each S' there are at most $\binom{d-k}{t-k}$ distinct S such that $S' \subseteq S$. Therefore we can bound the size of $\text{Rej}(Q_{x,z})$ by

$$|\text{Rej}(Q_{x,z})| \leq \binom{2\sqrt{d}}{k} \binom{d-k}{t-k}. \quad (2)$$

Combining (1) and (2), $|\text{Rej}(Q)| \leq q^2 \binom{2\sqrt{d}}{k} \binom{d-k}{t-k}$. It follows that for any non-adaptive algorithm making at most q queries,

$$\sum_{S \subseteq [d]: |S|=t} \Pr[\mathcal{A} \text{ reveals a violation for } f_S] \leq \mathbb{E}[|\text{Rej}(Q)|] \leq q^2 \binom{2\sqrt{d}}{k} \binom{d-k}{t-k}.$$

Hence there exists f_S such that $\Pr[\mathcal{A} \text{ reveals a violation for } f_S] \leq q^2 \frac{\binom{2\sqrt{d}}{k} \binom{d-k}{t-k}}{\binom{d}{t}} = q^2 \frac{\binom{2\sqrt{d}}{k} \binom{t}{k}}{\binom{d}{t}}$. \square

Proof of Claim 4.4. Let f'_S be the closest k -monotone function to f_S . Let Z denote the set $\{z \in \{0,1\}^{d-t} : d/2 - \sqrt{d} \leq |z| \leq d/2 + \sqrt{d} - t\}$. For any $t \leq \sqrt{d}$, $[\frac{d-t}{2} - \frac{\sqrt{d-t}}{2}, \frac{d-t}{2} + \frac{\sqrt{d-t}}{2}]$ is contained in $[d/2 - \sqrt{d}, d/2 + \sqrt{d} - t]$ so that $|Z| = \Omega(2^{d-t})$. For any assignment $z \in Z$ on coordinates indexed by $[d] \setminus S$, $f_S(\cdot, z)$ agrees with $\bigoplus_{i \in S} x_i$ and $f'_S(\cdot, z)$ is k -monotone. To finish the proof, it suffices to show that an anti-parity over t inputs is $\sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{k} / 2^t$ -far from any k -monotone function over t inputs. Indeed, this will imply that, for every $z \in Z$, $f_S(\cdot, z)$ differs from $f'_S(\cdot, z)$ on $\sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{k}$ points, and thus finally that f_S differs from $f_{S'}$ on $|Z| \cdot \sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{k} = \Omega(2^{d-t} \cdot \sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{k})$ points.

Now we show it is the case that an anti-parity function over t inputs $h(x_1, \dots, x_t) = \bigoplus_{i=1}^t x_i$ is far from any k -monotone function g (over x_1, \dots, x_t). We begin by noting that we can sample a random point from $\{0,1\}^t$ by first sampling a random chain $\mathcal{C} = (x^0 = 0^t, x^1, \dots, x^t = 1^t)$ (from all possible chains from 0^t to 1^t) then outputting x^i with probability $\binom{t}{i} / 2^t$. Thus the distance between h and g , namely $\Pr_x[h(x) \neq g(x)]$, is the same as $\Pr_{\mathcal{C}, i}[h(x^i) \neq g(x^i)]$ which can be further expanded as follows

$$\mathbb{E}_{\mathcal{C}} \left[\sum_{i=0}^t \frac{\binom{t}{i}}{2^t} \cdot \mathbb{1}_{\{h(x^i) \neq g(x^i)\}} \right] = \mathbb{E}_{\mathcal{C}} \left[\frac{1}{2^t} \sum_{i=1}^t \binom{t-1}{i-1} \cdot (\mathbb{1}_{\{h(x^i) \neq g(x^i)\}} + \mathbb{1}_{\{h(x^{i-1}) \neq g(x^{i-1})\}}) \right] \quad (3)$$

where the last inequality relies on identity $\binom{t}{i} = \binom{t-1}{i-1} + \binom{t-1}{i}$. For any fixed chain \mathcal{C} , because g alternates at most k times, there are at least $t-k$ choices of $i \in [t]$ such that $g(x^{i-1}) = g(x^i)$, which implies $\mathbb{1}_{\{h(x^i) \neq g(x^i)\}} + \mathbb{1}_{\{h(x^{i-1}) \neq g(x^{i-1})\}} \geq 1$ (due to $h(x^{i-1}) \neq h(x^i)$). Thus $\sum_{i=1}^t \binom{t-1}{i-1} \cdot (\mathbb{1}_{\{h(x^i) \neq g(x^i)\}} + \mathbb{1}_{\{h(x^{i-1}) \neq g(x^{i-1})\}})$ is at least the sum of smallest $t-k$ binomials among $\binom{t-1}{0}, \dots, \binom{t-1}{t-1}$ which is

$$\sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t-1}{i} + \sum_{i=0}^{\lfloor \frac{t-k-2}{2} \rfloor} \binom{t-1}{t-1-i} = \sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t-1}{i} + \sum_{i=0}^{\lfloor \frac{t-k-2}{2} \rfloor} \binom{t-1}{i} \geq \sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{i},$$

which implies $\Pr_x[h(x) \neq g(x)] \geq \frac{1}{2^t} \sum_{i=0}^{\lfloor \frac{t-k-1}{2} \rfloor} \binom{t}{i}$ by combining the above with (3). \square

□

4.2.2 Two-sided lower bounds

The following theorem gives a construction that enables us to convert monotone functions into k -monotone functions, and functions that are far from monotone into functions that are far from k -monotone.

Theorem 4.5. *There exists an efficiently computable function $h: \{0, 1\}^{d/2} \rightarrow \{0, 1\}$ such that for any $g: \{0, 1\}^{d/2} \rightarrow \{0, 1\}$, then $g||h: \{0, 1\}^d \rightarrow \{0, 1\}$ is a Boolean function (defined as $g||h(x, y) \stackrel{\text{def}}{=} g(x) \oplus h(y)$ for any $x, y \in \{0, 1\}^{d/2}$) satisfying the following.*

- if g is monotone, then $g||h$ is a k -monotone function;
- if g is ε -far from monotone, then $g||h$ is $\Omega(\varepsilon/k)$ -far from being a k -monotone function;

The above theorem reduces testing monotonicity to testing k -monotonicity (for arbitrary constant k) with the same number of queries to the input function. As we will see, it carries any lower bound for testing monotonicity over to testing k -monotonicity, while preserving the characteristics (two-sidedness, adaptivity) of the original lower bound. In particular, combining it with the recent of [CDST15, BB16], we obtain the following corollary.

Corollary 4.6. *For any $c > 0$ and $k \geq 1$, there exists $\varepsilon = \varepsilon(k, c) > 0$ such that any 2-sided non-adaptive algorithm for testing whether f is k -monotone or ε -far from it requires $\Omega(d^{1/2-c})$ queries. Any 2-sided adaptive algorithm requires $\tilde{\Omega}(d^{1/4})$ queries.*

To prove **Theorem 4.5**, we prove the following three claims. **Claim 4.7** and **Claim 4.8** assume the existence of a $(k - 1)$ -monotone function h which satisfies an interesting structural property. This property is then exploited to establish **Theorem 4.5**. Finally in **Claim 4.9**, we establish the existence of such h , and **Theorem 4.5** follows.

Claim 4.7. *Let $h: \{0, 1\}^d \rightarrow \{0, 1\}$ be a $(k - 1)$ -monotone function. Then for any monotone $g: \{0, 1\}^d \rightarrow \{0, 1\}$, $f = g||h$ is a k -monotone function.*

Claim 4.8. *Suppose there exists a $h: \{0, 1\}^d \rightarrow \{0, 1\}$ such that the following holds. There exists at least M paths of length $k - 1$ such that (i) all paths are vertex disjoint and (ii) for every path $y_1 \preceq \dots \preceq y_k$, $h(y_1) = 0$ and $h(y_i) \neq h(y_{i+1})$ for $1 \leq i \leq k - 1$. Then, for any $g: \{0, 1\}^d \rightarrow \{0, 1\}$ which is ε -far from being a monotone function, the function $f = g||h$ is $\Omega(\frac{M}{2^d} \cdot \varepsilon)$ -far from being a k -monotone function.*

Claim 4.9. *For any constant k , there exists an efficient computable $h: \{0, 1\}^d \rightarrow \{0, 1\}$ such that h is a $(k - 1)$ -monotone function and h contains at least $\frac{(1-o_d(1))2^d}{k}$ paths of length $k - 1$ such that all paths are vertex disjoint and for every path $y_1 \preceq \dots \preceq y_k$, $h(y_1) = 0$ and $h(y_i) \neq h(y_{i+1})$ for $1 \leq i \leq k - 1$.*

Proof of Claim 4.7. Suppose $f = g||h$ is not k -monotone, then there exist $(x_1, y_1), \dots, (x_{k+1}, y_{k+1})$ such that $(x_1, y_1) \preceq \dots \preceq (x_{k+1}, y_{k+1})$, $f(x_1, y_1) = 1$ and $f(x_i, y_i) \neq f(x_{i+1}, y_{i+1})$ for any $1 \leq i \leq k$. Because g is monotone, either g is constant on x_1, \dots, x_{k+1} , or there exists an index $1 < j \leq k + 1$ such that $g(x_i) = 0$ for $i < j$ and $g(x_i) = 1$ for $i \geq j$. In the first case, if g is the constant 0 on x_1, \dots, x_{k+1} (resp. constant 1) either $h(y_i) = f(x_i, y_i)$ (resp. or $h(y_i) = 1 - f(x_i, y_i)$) for any i and h

alters exactly k times on points $y_1 \preceq \dots \preceq y_{k+1}$. For the second case, $h(y_1) = f(x_1, y_1) = 1$, and h alters exactly $k - 1$ times on $y_1 \preceq \dots \preceq y_{j-1} \preceq y_{j+1} \preceq \dots \preceq y_{k+1}$. Both cases contradict h being $(k - 1)$ -monotone. \square

Proof of Claim 4.8. Let M_h be the maximal set of paths of length k such that all paths are vertex disjoint and for every path $y_1 \preceq \dots \preceq y_k$, $h(y_1) = 0$ and $h(y_i) \neq h(y_{i+1})$ for $1 \leq i \leq k - 1$. Let M_g be the maximal set of pairs such that all pairs are vertex disjoint and every pair $x_1 \preceq x_2$ is a violation for g , i.e., $g(x_1) = 1$ and $g(x_2) = 0$. For each path $(y_1 \preceq \dots \preceq y_k) \in M_h$ and any pair $(x_1 \preceq x_2) \in M_g$, it is easy to see the following path is a violation for $f||g$ being k -monotone :

$$(x_1, y_1), \dots, (x_1, y_k), (x_2, y_k).$$

Let f' be the closest k -monotone function to f . For each violating path, f and f' differ on at least 1 point. Because both M_h and M_g are vertex disjoint, violating paths constructed by taking every path in M_h and every pair in M_g are vertex disjoint. Thus f and f' differ on at least $|M_h| \times |M_g|$ points and f is $(|M_h| \cdot |M_g| / 2^{2d})$ -far from f' . It is known ([GGL⁺00]) that for any $g: \{0, 1\}^d \rightarrow \{0, 1\}$ which is ε -far from monotone, $|M_g| \geq 2^{d-1}\varepsilon$. The desired conclusion follows. \square

Proof of Claim 4.9. Let B_1, B_2, \dots, B_k be consecutive blocks each consisting of consecutive hamming levels of the hypercube. Letting, $|B_i|$ denote the number of points in the i^{th} block, note that it is possible to have all blocks satisfy,

$$\forall i \in [k] \left(1 - \frac{k}{\sqrt{d}}\right) \frac{2^d}{k} \leq |B_i| \leq \left(1 + \frac{k}{\sqrt{d}}\right) \frac{2^d}{k}.$$

Because k is a constant and every layer contains at most $2^d/\sqrt{d}$ points, we can always greedily find B_1, \dots, B_k one by one. Let h be a function such that h is constant over every block and alternates on successive blocks with evaluates to 0 over all of B_1 . Note that h is $(k - 1)$ -monotone. Next we inductively argue that for any $1 \leq j \leq k$, B_1, \dots, B_j contain at least $\frac{(1-o_d(1))2^d}{k}$ vertex disjoint paths of length $j - 1$ such the i th point on every path is in B_i . Claim 4.9 follows from the case $j = k$.

For $j = 1$, the statement holds by taking all points in B_1 . For $j > 1$, assume that B_1, \dots, B_{j-1} contain a set P_{j-1} of such $\frac{(1-o_d(1))2^d}{k}$ vertex disjoint paths of length $j - 2$. Let M be the maximal matching between B_{j-1} and B_j and let P_j be the set of paths of length $j - 1$ constructed in following way: for each path in P_{j-1} with an endpoint y_{j-1} , if there exists y_j such that $(y_{j-1}, y_j) \in M$, we add the path appended with y_j into P_j . Because no points in B_j will be added into two different paths in P_j and P_{j-1} are vertex disjoint, paths in P_j are vertex disjoint.

Now we show $|M| = \min(|B_{j-1}|, |B_j|)$ which implies $|P_j| \geq \frac{(1-o_d(1))2^d}{k}$. Suppose $|B_{j-1}| \leq |B_j|$ (the argument is analogous in the other case). For any subset S of B_{j-1} , let f_S be the indicator function of the upper closure of S denoted as $N(S)$. It is not hard to check that f_S is monotone and thus $\Pr[f_S(x) = 1 | x \in B_j] \geq \Pr[f_S(x) = 1 | x \in B_{j-1}]$. It follows

$$|N(S) \cap B_j| = |B_j| \Pr[f_S(x) = 1 | x \in B_j] \geq |B_{j-1}| \Pr[f_S(x) = 1 | x \in B_{j-1}] \geq |S|.$$

By Hall's theorem, $|M| = |B_{j-1}|$. By similar argument, we can show $|M| = |B_j|$ when $|B_{j-1}| > |B_j|$. Thus $|M| = \min(|B_{j-1}|, |B_j|)$. \square

5 On the line

In this section we prove our results on testing k -monotonicity on the line, that is of functions $f: [n] \rightarrow \{0, 1\}$. We start with [Theorem 1.4](#), which establishes that this can be done non-adaptively with one-sided error, with only $O(k/\varepsilon)$ queries; we then turn to [Theorem 1.3](#), which shows that this is the best one can hope for if we insist on one-sidedness. The last result of this section is [Theorem 1.5](#), where we show that – perhaps unexpectedly – *two-sided* algorithms, even non-adaptive, can break this barrier and test k -monotonicity with *no* dependence on k .

5.1 Upper and lower bounds for one-sided testers

We first prove the upper bound, restated below:

Theorem 1.4. *There exists a one-sided non-adaptive tester for k -monotonicity of functions $f: [n] \rightarrow \{0, 1\}$ with query complexity $q(n, \varepsilon, k) = O\left(\frac{k}{\varepsilon}\right)$.*

Proof. We assume that $\frac{\varepsilon n}{50k}$ is an integer,⁵ and partition the domain into $K \stackrel{\text{def}}{=} \frac{50k}{\varepsilon}$ intervals of size $\frac{\varepsilon n}{50k}$, the consecutive “blocks” B_1, \dots, B_K . We then define $g: [n] \rightarrow \{0, 1, *\}$ as the function constant on each block $B_i = \{b_i, \dots, b_{i+1} - 1\}$, such that

- If $f(b_i) = f(b_{i+1} - 1)$, then $g(j) = f(b_i)$ for all $j \in B_i$;
- otherwise, $g(j) = *$ for all $j \in B_i$.

We say that a block B_i such that $g|_{B_i} = *$ is a *changepoint block* for g . Clearly, given query access to f one can obtain the value of g on any point $j \in [n]$ with only two queries to f . Moreover, defining $\tilde{g}: [n] \rightarrow \{0, 1\}$ to be the function obtained from g by replacing $*$ by 0, we observe the following:

- If f is k -monotone, then (i) so is \tilde{g} , and (ii) f and \tilde{g} differ in at most k blocks (namely the changepoint blocks of g), so that $\text{dist}(f, \tilde{g}) \leq k \cdot \frac{1}{K} = \frac{\varepsilon}{50}$;
- If f is ε -far from k -monotone, then either (i) \tilde{g} is not k -monotone, or (ii) $\text{dist}(f, \tilde{g}) > \varepsilon$.

We first discuss a natural idea, why it fails, and how to fix it. The above suggests a very simple algorithm: first, learn the function \tilde{g} (exactly) using $O(k/\varepsilon)$ queries, and then distinguish between $\text{dist}(f, \tilde{g}) \leq \frac{\varepsilon}{50}$ and $\text{dist}(f, \tilde{g}) > \varepsilon$ by random sampling, using $O(1/\varepsilon)$ queries. This immediately results in a non-adaptive testing algorithm with query complexity $O(k/\varepsilon)$; however, this tester is inherently 2-sided, due to the second (distance estimation) step. To obtain one-sidedness, we thus take a slightly different route, and consider the number of blocks on which f and \tilde{g} differ instead of estimating the distance. As finding at least $k + 1$ such blocks can only happen when f is not k -monotone, this will never result in rejecting a k -monotone function; but this one-sidedness comes at the price of a slightly less straightforward analysis.

As mentioned above, we start by learning g (and thus \tilde{g}) exactly, using $2K = O(k/\varepsilon)$ non-adaptive queries. Setting $m \stackrel{\text{def}}{=} C \frac{k}{\varepsilon}$, we also sample $m' \sim \text{Poisson}(m)$ points⁶ $j_1, \dots, j_{m'}$ independently and uniformly from $[n]$, where $C > 0$ is a constant to be determined in the course of the analysis, and

⁵If not, we consider instead $\varepsilon' \stackrel{\text{def}}{=} \frac{50k}{n} \lfloor \frac{\varepsilon n}{50k} \rfloor > \varepsilon - \frac{50k}{n} > \frac{\varepsilon}{2}$ if $\varepsilon > \frac{100k}{n}$; while if $\varepsilon \leq \frac{100k}{n}$ we query the entire function, for a total of $n = O\left(\frac{k}{\varepsilon}\right)$ queries.

⁶The fact that we sample $\text{Poisson}(m)$ instead of m is for ease of the analysis; note that due to the tight concentration of Poisson random variables, with probability $1 - o(1)$ we will have $m' \leq 2m$. If this does not happen, the tester can output ACCEPT, incurring only a small additional error probability (and not affecting the one-sidedness).

query the value of f (and g) on all of them. Then, we reject if either (i) \tilde{g} is not k -monotone; or (ii) there exist at least $k + 1$ distinct blocks which contain a sample s_j such that $\tilde{g}(s_j) \neq f(s_j)$.

By definition, this tester is non-adaptive; and it is not difficult to see it accepts any k -monotone function with probability 1, since in that case f and g (and a fortiori \tilde{g}) differ in at most k blocks: indeed, these blocks can only be changepoint blocks for g , i.e. blocks where f changes value.

It remains to argue soundness: we will show that if f is ε -far from k -monotone, the tester will reject with probability at least $2/3$. By the first check made, (i), we can assume in the following that \tilde{g} is k -monotone – as otherwise f is rejected with probability 1 – and we need to show that (ii) will reject with probability at least $2/3$. For each block B_i (where $i \in [K]$), let $p_i \in [0, \frac{1}{K}]$ be defined as the (normalized) number of points in B_i on which f and \tilde{g} differ (we henceforth refer to such a point as a *giveaway point*):

$$p_i \stackrel{\text{def}}{=} \frac{1}{n} \sum_{j \in B_i} \mathbb{1}_{\{f(j) \neq \tilde{g}(j)\}}.$$

Since f is ε -far from the k -monotone function \tilde{g} , we have $\sum_{i=1}^K p_i \geq \varepsilon$. Now, letting Z_i be the indicator of the event that among the m' samples, at least one is giveaway point from B_i , and $Z = \sum_{i=1}^K Z_i$, we can write $Z_i = \mathbb{1}_{\{Y_i \neq 0\}}$, where the $(Y_i)_{i \in [K]}$ are *independent* Poisson random variables with $Y_i \sim \text{Poisson}(mp_i)$. The expected number of blocks in which a giveaway point is sampled is then

$$\mathbb{E}Z = \sum_{i=1}^K \mathbb{E}Z_i = \sum_{i=1}^K \Pr[Y_i \neq 0] = \sum_{i=1}^K (1 - e^{-mp_i})$$

Since for every $i \in [K]$ it holds that $e^{-mp_i} \leq 1 - \frac{m}{2}p_i$ (the inequality holding since $0 \leq mp_i \leq 1$, which is verified for $m \leq K$), we get

$$\mathbb{E}Z = \sum_{i=1}^K \mathbb{E}Z_i \geq \sum_{i=1}^K \frac{m}{2}p_i \geq \frac{m\varepsilon}{2} \geq \frac{C}{2}k.$$

Moreover, by a Chernoff bound, we get that

$$\Pr\left[Z < \frac{C}{4}k\right] \leq e^{-\frac{Ck}{16}} \leq e^{-\frac{C}{16}}$$

which is less than $1/4$ for $C \geq 23$. Setting $C \stackrel{\text{def}}{=} 30$ satisfies both conditions that $m \leq K$ and $C \geq 23$, and results in a one-sided non-adaptive tester which rejects functions far from k -monotone with probability at least $1 - \frac{1}{4} + o(1) \geq \frac{2}{3}$. □

Turning to the lower bounds against one-sided testers, we show the following:

Theorem 1.3. *Any one-sided (possibly adaptive) tester for k -monotonicity of functions $f: [n] \rightarrow \{0, 1\}$ must have query complexity $\Omega\left(\frac{k}{\varepsilon}\right)$.*

Proof. Since a lower bound of $\Omega\left(\frac{1}{\varepsilon}\right)$ straightforwardly holds, we can restrict ourselves to $k \geq 8$, and $\varepsilon < \frac{1}{12}$; moreover, we assume without loss of generality that $\frac{\varepsilon n}{k}$ is an integer, and partition the domain into $K \stackrel{\text{def}}{=} \frac{k}{\varepsilon}$ intervals of size $\frac{\varepsilon n}{k}$, the consecutive “blocks” B_1, \dots, B_K . For $v \in \{0, 1\}^{K/2}$, we

define $g_v: [n] \rightarrow \{0, 1\}$ as the function which has constant value v_i on block B_{2i-1} and has constant value 1 on the remaining blocks.

Consider the distribution over $\{g_v\}_v$ where each coordinate of v is independently set to 0 with probability $p \stackrel{\text{def}}{=} 6\varepsilon$, and 1 otherwise. We next show that g_v is at least ε -far from any k -monotone function with very high probability over the choice of v . By a Chernoff bound, with probability at least $1 - e^{-pK/16} = 1 - e^{-3k/8}$, g_v has at least $pK/4 = 3k/2$ blocks that are 0 blocks. Conditioned on this, it is easy to see that g_v is ε -far from k -monotone: indeed, to make it k -monotone one has to flip its value on at least k blocks, and each block contains an ε/k fraction of the domain.

Fix any deterministic adaptive algorithm with query complexity $q \leq k/(24\varepsilon)$ queries, and denote by x_1, \dots, x_q the sequence of queries made (when given query access to some function g_v). Without loss of generality, we assume that x_1, \dots, x_q are made into distinct odd blocks. x_1, \dots, x_q reveals a violation if and only if number of 0's in corresponding answers is at least $k/2 + 1$.

If x_1, \dots, x_q are non-adaptive queries, the probability that the number of 0's in $f(x_1), \dots, f(x_q)$ is at least $k/2 + 1$ can be upper bounded by Chernoff Bound in a straightforward way. Because each odd block is sampled independently, the same bound can be obtained in the adaptive setting. Note, for arbitrary $a \in \{0, 1\}^q$,

$$\Pr[f(x_1) = a_1, \dots, f(x_q) = a_q] = \Pr[f(x_1) = a_1] \cdot \prod_{i=2}^q \Pr[f(x_i) = a_i \mid f(x_1) = a_1, \dots, f(x_{i-1}) = a_{i-1}].$$

Because x_i is determined by $f(x_1) = a_1, \dots, f(x_{i-1}) = a_{i-1}$ and each odd block is sampled independently, for every i it holds that $\Pr[f(x_i) = a_i \mid f(x_1) = a_1, \dots, f(x_{i-1}) = a_{i-1}] = 6\varepsilon$ if $a_i = 0$ and $\Pr[f(x_i) = a_i \mid f(x_1) = a_1, \dots, f(x_{i-1}) = a_{i-1}] = 1 - 6\varepsilon$ if $a_i = 1$. Thus,

$$\Pr[f(x_1) = a_1, \dots, f(x_q) = a_q] = (1 - 6\varepsilon)^{|a|} (6\varepsilon)^{q-|a|}. \quad (4)$$

Let X_i be the indicator that $f(x_i) = 0$. We get that, writing $F(i, N, p)$ for the cumulative distribution function of a Binomial with parameters N and p ,

$$\begin{aligned} \Pr\left[\sum_{i=1}^q X_i \geq \frac{k}{2} + 1\right] &\leq \sum_{a \in \{0,1\}^q: |\bar{a}| \geq k/2} (1 - 6\varepsilon)^{|a|} (6\varepsilon)^{q-|a|} \\ &= \sum_{\ell=0}^{q-k/2} \binom{q}{\ell} (1 - 6\varepsilon)^\ell (6\varepsilon)^{q-\ell} = F\left(q - \frac{k}{2}, q, 1 - 6\varepsilon\right) \\ &= F(q(1-x), q, 1 - 6\varepsilon) \quad \left(x \stackrel{\text{def}}{=} \frac{k}{2q} \in (12\varepsilon, 1)\right) \\ &\leq e^{-qD(1-x \parallel 1-6\varepsilon)} \quad (\text{Relative entropy Chernoff bound}^7) \end{aligned}$$

where $D(p \parallel p') \stackrel{\text{def}}{=} p \ln \frac{p}{p'} + (1-p) \ln \frac{1-p}{1-p'}$, and we used the fact that $\frac{k}{2} + 1 \leq q \leq \frac{k}{24\varepsilon}$. Rewriting slightly the right-hand-side, we obtain

$$\Pr\left[\sum_{i=1}^q X_i \geq \frac{k}{2} + 1\right] \leq e^{-\frac{k}{2} \Phi(x)}$$

⁷Recall that the relative entropy version of the Chernoff bound states that $F(m, N, p) \leq e^{-mD(\frac{m}{N} \parallel p)}$ as long as $0 \leq \frac{m}{N} \leq p$.

for $\Phi(x) \stackrel{\text{def}}{=} \frac{1}{x} \left((1-x) \ln \frac{1-x}{1-6\varepsilon} + x \ln \frac{x}{6\varepsilon} \right)$. It is not hard to see that Φ is increasing on $[6\varepsilon, 1)$, and since $x \geq 12\varepsilon$ the right-hand-side is at most $e^{-\frac{k}{2}\Phi(12\varepsilon)}$. It then suffices to observe that, for $\varepsilon \leq \frac{1}{12}$, it holds that $\Phi(12\varepsilon) \geq \Phi(0) = \ln 2 - \frac{1}{2} > \frac{1}{8}$ to conclude that

$$\Pr \left[\sum_{i=1}^q X_i \geq \frac{k}{2} + 1 \right] \leq e^{-\frac{k}{16}}$$

and therefore obtain

$$\Pr[f(x_1), \dots, f(x_q) \text{ contains at least } (k/2 + 1) \text{ zeros}] \leq e^{-\frac{k}{16}}.$$

Combining the two, this shows that the probability that x_1, \dots, x_q does not reveal a violation for g_v while g_v is ε -far from k -monotone is at least $1 - e^{-k/16} - e^{-3k/8} > 1/3$ (since $k \geq 8$). By Yao's principle, for any (possibly randomized) non-adaptive algorithm \mathcal{A} making at most $k/(24\varepsilon)$ there exists a fixed v such that g_v is ε -far from k -monotone yet \mathcal{A} rejects g_v with probability less than $2/3$. The desired conclusion follows. \square

5.2 Upper bound for two-sided testers: proof of **Theorem 1.5**

In this section, we prove the two-sided non-adaptive upper bound of **Theorem 1.5**, restated below:

Theorem 1.5. *There exists a two-sided non-adaptive tester for k -monotonicity of functions $f: [n] \rightarrow \{0, 1\}$ with query complexity $q(n, \varepsilon, k) = \tilde{O}(1/\varepsilon^7)$, independent of k .*

In what follows, we assume that $k > 20/\varepsilon$, as otherwise we can use for instance the $O(k/\varepsilon)$ -query (non-adaptive, one-sided) tester of **Theorem 1.4** to obtain an $O(1/\varepsilon^2)$ query complexity.

5.2.1 Testing k -monotonicity over $[Ck]$

In this section, we give a $\text{poly}(C/\varepsilon)$ -query tester for k -monotonicity over the domain $[Ck]$, where C is a parameter to be chosen (for our applications, we will eventually set $C = \text{poly}(1/\varepsilon)$).

Lemma 5.1. *There exists a two-sided non-adaptive tester for k -monotonicity of functions $f: [Ck] \rightarrow \{0, 1\}$ with query complexity $O\left(\frac{C^3}{\varepsilon^3}\right)$.*

The tester proceeds by reducing to support size estimation and using (a slight variant of) an algorithm of Canonne and Rubinfeld [CR14]. Let $f: [Ck] \rightarrow \{0, 1\}$, and suppose f is s -monotone but not $(s-1)$ -monotone. Then there is a unique partition of $[Ck]$ into $s+1$ disjoint intervals I_1, I_2, \dots, I_{s+1} such that f is constant on each interval; note that this constant value alternates in consecutive intervals. We can then define a distribution D_f over $[s+1]$ such that $D_f(i) = |I_i|/(Ck)$.

Our next claims, **Claim 5.2** and **Claim 5.3**, provide the basis for the reduction (from testing k -monotonicity of f to support size estimation of D_f).

Claim 5.2. *If f is ε -far from k -monotone, then it is not $(1 + \frac{\varepsilon}{4})k$ -monotone, and in particular $|\text{supp}(D_f)| > (1 + \frac{\varepsilon}{4})k + 1$.*

Claim 5.3. *To ε -test k -monotonicity of f , it suffices to estimate $|\text{supp}(D_f)|$ to within $\frac{\varepsilon k}{10}$.*

The proofs of above claims are relatively straightforward. So we defer these proofs to the end of this section and now proceed to use algorithm of [CR14] to do support size estimation. The algorithm of [CR14] uses “dual access” to D ; an oracle that provides a random sample from D , and an oracle that given an element of D , returns the probability mass assigned to this element by D .

Theorem 5.4 ([CR14, Theorem 14 (rephrased)]). *In the access model described above, there exists an algorithm that, on input a threshold $n \in \mathbb{N}^*$ and a parameter $\varepsilon > 0$, and given access to a distribution D (over an arbitrary set) satisfying*

$$\min_{x \in \text{supp}(D)} D(x) \geq \frac{1}{n}$$

estimates the support size $|\text{supp}(D)|$ up to an additive εn , with query complexity $O\left(\frac{1}{\varepsilon^2}\right)$.

Note however that we only have access to D_f through query access to f , and thus have to manage to simulate (efficiently) access to the former. One difficulty is that, to access $D_f(i)$, we need to determine where I_i lies in f . For example, finding $D_f(k/2)$ requires finding $I_{k/2}$, which might require a large number of queries to f . We circumvent this by weakening the “dual access” model in two ways, arguing for each of these two relaxations that the algorithm of [CR14] can still be applied:

- we rewrite the support size as in [CR14], as $|\text{supp}(D_f)| = \mathbb{E}_{x \sim D_f}[1/D_f(x)]$. We want to estimate it to within $\pm O(\varepsilon k)$ which we can do by random sampling;
- the quantity inside the expectation depends on $D_f(x)$ but not x itself, so “labels” are unnecessary for our random sampling. Thus, it will be sufficient to be able to compute $D_f(x)$ (and thus $1/D_f(x)$) for a random $x \sim D_f$, even if not actually knowing x itself;
- actually, even calculating $D_f(x)$ may possibly too expensive, so instead we will estimate $\mathbb{E}_{x \sim D_f}[1/\tilde{D}_f(x)]$ where $\tilde{D}_f(x) = \min(\frac{20}{\varepsilon k}, D_f(x))$. Note that \tilde{D}_f might no longer define a probability distribution; but this expectation is only off by at most $\frac{\varepsilon k}{20}$, since $1/D'_f(x) = \max(\varepsilon k/20, 1/D_f(x))$ and $1/D_f(x)$ is positive.

More details follow.

First, we note as discussed above that the algorithm does not require knowing the “label” of any element in the support of the distribution: the only access required is being able to randomly sample elements according to D_f , and evaluate the probability mass on the sampled points. This, in turn, can be done, as the following two lemmas explain:

Lemma 5.5 (Sampling from D_f). *Let $i \in [n]$ be chosen uniformly at random, and let j be such that $i \in I_j$. Then, the distribution of j is exactly D_f .*

Lemma 5.6 (Evaluating $D_f(j)$). *Suppose $I_j = \{a, a + 1, \dots, b\}$. Given i such that $i \in I_j$, we can find I_j by querying $f(i + 1) = f(i + 2) = \dots = f(b)$ and $f(b + 1) \neq f(b)$, as well as $f(i - 1) = f(i - 2) = \dots = f(a)$ and $f(a - 1) \neq f(a)$. The number of queries to f is $b - a + 3 = |I_j| + 3$.*

Here comes the second difficulty: if we straightforwardly use these approaches to emulate the required oracles to estimate the support size of D_f , the number of queries is potentially very large. For instance, if we attempt to query $D_f(j)$ where $|I_j| = \Omega(k)$, we will need $\Omega(k)$ queries to f . This is where comes the second relaxation: specifically, we shall argue that it will be enough for us to “cap” the size of the interval (as per our next lemma).

Lemma 5.7 (Evaluating $D_f(j)$ with a cap). *Given i such that $i \in I_j$, we will query f on every point in $[i - 20C/\varepsilon, i + 20C/\varepsilon]$. If $|I_j| \leq 20C/\varepsilon$, then I_j will be determined by these queries. If these queries do not determine I_j , we know $|I_j| > 20C/\varepsilon$. Beyond querying i , this requires $40C/\varepsilon$ (nonadaptive) queries.*

We now can put all the above pieces together and give the proof for Lemma 5.1:

Proof of Lemma 5.1: As previously discussed, we use the algorithm of [CR14] for estimating support size. Inspecting their algorithm, we see that our cap of $20C/\varepsilon$ for interval length (and therefore $20/(\varepsilon k)$ for maximum probability reported) might result in further error of the estimate. The algorithm interacts with the unknown function by estimating the expected value of $1/D_f(j)$ over random choices of j with respect to D_f . Our cap can only decrease this expectation by at most $(\varepsilon k)/20$. Indeed, the algorithm works by estimating the quantity $\mathbb{E}_{x \sim D_f}[\frac{1}{D_f(x)} \mathbb{1}_{\{D_f(x) > \tau\}}]$, for some suitable parameter $\tau > 0$. By capping the value of $1/D_f(x)$ to $20/(\varepsilon k)$, we can therefore only decrease the estimate, and by at most $20/(\varepsilon k) \cdot D_f(\{x : D_f(x) > (\varepsilon k)/20\}) \leq 20/(\varepsilon k)$.

The condition for their algorithm to estimate support size to within $\pm \varepsilon m$ is that all elements in the support have a probability mass of at least $1/m$. Since each nonempty interval has length at least 1, we have $\min_j D_f(j) \geq (1/Ck)$. In order for their algorithm to report an estimate within $\pm \varepsilon k/20$ of support size, we set $\varepsilon' = (\varepsilon/20C)$ in their algorithm.

The total error in support size is at most $\varepsilon k/20 + \varepsilon k/20 = \varepsilon k/10$. By Claim 5.3, this suffices to test ε -test k -monotonicity of f .

Using the algorithm of [CR14], we need $O(1/\varepsilon'^2) = O((C/\varepsilon)^2)$ queries to D_f . For every query to D_f , we need to make $O(C/\varepsilon)$ queries to f , so the overall query complexity is $O(C^3/\varepsilon^3)$. \square

Proof of Claim 5.2. The last part of the statement is immediate from the first, so it suffices to prove the first implication. We show the contrapositive: assuming f is $(1 + \frac{\varepsilon}{4})k$ -monotone, we will “fix” it into a k -monotone function by changing at most εn points. In what follows, we assume $\frac{\varepsilon k}{4} \geq 1$, as otherwise the statement is trivial (any function that is ε -far from k -monotone is *a fortiori* not k -monotone).

Let as before ℓ^* be the minimum integer ℓ for which f is ℓ -monotone: we can assume $k < \ell^* \leq (1 + \frac{\varepsilon}{4})k$ (as if $\ell^* \leq k$ we are done.) We will show that taking the set of $(\ell^* - k)$ and altering them all at once produces a k -monotone function *eps*-close to f . Consider as above the maximal consecutive monochromatic intervals I_1, \dots, I_{ℓ^*} , and let i be the index of the shortest one. In particular, it must be the case that $|I_i| \leq \frac{n}{\ell^* + 1}$. Flipping the value of f on I_i therefore has “cost” at most $\frac{n}{\ell^* + 1}$, and the resulting function f' is now exactly $(\ell^* - 2)$ -monotone if $1 < i < \ell^*$, and $(\ell^* - 1)$ -monotone if $i \in \{1, \ell^*\}$. This means in particular that repeating the above $\frac{\varepsilon}{4}k$ times is enough to obtain a k -monotone function, and the total cost is upperbounded by

$$\sum_{j=0}^{\varepsilon k/4} \frac{n}{\ell^* + 1 - 2j} \leq \sum_{j=0}^{\varepsilon k/4} \frac{n}{k + 1 - 2j} = \sum_{j=k(1-\frac{\varepsilon}{2})+1}^{k+1} \frac{n}{j} \leq n \frac{\frac{\varepsilon}{2}k + 1}{(1 - \frac{\varepsilon}{4})k + 1} \leq n \frac{\frac{3\varepsilon}{4}k}{(1 - \frac{\varepsilon}{4})k}$$

where for the last inequality (for the numerator) we used that $1 \leq \frac{\varepsilon k}{4}$. But this last RHS is upperbounded by εn (as $\frac{3}{4}x \leq x(1 - \frac{1}{4}x)$ for $x \in [0, 1]$), showing that therefore, f was ε -close to k -monotone to begin with, which is a contradiction. \square

Proof of Claim 5.3. If f is ε -far from k -monotone, then $|\text{supp}(D_f)| > (1 + \frac{\varepsilon}{4})k = k + \frac{\varepsilon}{4}k$, and if f is k -monotone, then $|\text{supp}(D_f)| \leq k + 1$. The fact that $k > 20/\varepsilon$ then allows us to conclude. \square

5.2.2 Reducing $[n] \rightarrow \{0, 1\}$ to $[Ck] \rightarrow \{0, 1\}$.

Now we show how to reduce ε -testing k -monotonicity of $f: [n] \rightarrow \{0, 1\}$ to ε' -testing k -monotonicity of a function $g: [Ck] \rightarrow \{0, 1\}$ for $C = \text{poly}(1/\varepsilon)$ and $\varepsilon' = \text{poly}(\varepsilon)$, resulting in a $\text{poly}(1/\varepsilon)$ -query algorithm for ε -testing k -monotonicity.

The first step is (as before) to divide $[n]$ into blocks (disjoint intervals) of size $\frac{\varepsilon n}{4k}$ if $\varepsilon > \frac{8k}{n}$ (again assuming without loss of generality that $\frac{\varepsilon n}{4k}$ is an integer), and blocks of size 1 otherwise (in which case $n \leq \frac{8k}{\varepsilon}$ and we can directly apply the result of [Lemma 5.1](#), with $C = n/k \leq 8/\varepsilon$). Let $m = 4k/\varepsilon$ be the number of resulting blocks, and define $f_m: [n] \rightarrow \{0, 1\}$ as the m -block-coarsening of f : namely, for any $j \in B_i$, we set

$$f_m(j) = \operatorname{argmax}_{b \in \{0,1\}} \Pr_{k \in B_i} [f_m(k) = b] \quad (\text{majority vote})$$

Ordering the blocks B_1, B_2, \dots, B_m , we also define $g: [m] \rightarrow \{0, 1\}$ such that $g(i) = \min_{a \in B_i} f_m(a)$.

It is easy to see that if f is k -monotone, then f has at most k non-constant blocks, and f_m is k -monotone. Because the function f only changes values k times; for a block to be non-constant, the block must contain a pair of points with a value change. We call a block *variable* if the minority points comprise at least an $\varepsilon/100$ -fraction of the block; formally, B is variable if $\min_{b \in \{0,1\}} \Pr_{j \in B} [f(j) = b] \geq \varepsilon/100$.

We need the following claims (their proofs are at the end of the section) to prove [Theorem 1.5](#).

Claim 5.8. *Suppose f has s variable blocks. Then $\operatorname{dist}(f, f_m) \leq s/m + \varepsilon/100$.*

Claim 5.9. *Suppose f is promised to be either (i) k -monotone or (ii) such that f_m has more than $\frac{5}{4}k$ variable blocks. Then we can determine which with $O\left(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon}\right)$ queries, and probability $9/10$.*

Proof of Theorem 1.5. We use the estimation/tester from the previous claim as the first part of our tester. Assuming f passes, we can assume that f_m has less than $\frac{5}{4}k$ variable blocks. By [Claim 5.8](#), $\operatorname{dist}(f, f_m) \leq \frac{5k}{4}/m + \frac{\varepsilon}{100} = \frac{5\varepsilon}{16} + \frac{\varepsilon}{100} \leq \frac{\varepsilon}{3}$. This part takes $O\left(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon}\right)$ queries.

Now, we apply the tester of [Lemma 5.1](#) (with probability of success amplified to $9/10$ by standard arguments) to $(\varepsilon/6)$ -test k -monotonicity of $g: [m] \rightarrow \{0, 1\}$, where $g(i)$ is the constant value of f_m on B_i , and $m = (4k)/\varepsilon$. Let q be the query complexity of the tester, and set $\delta = 1/(10q)$; to query $g(i)$, we randomly query f on $O\left(\frac{1}{\varepsilon} \log \frac{1}{\delta}\right)$ points in B_i and take the majority vote. With probability at least $1 - \delta$, we get the correct value of $g(i)$, and by a union bound all q simulated queries have the correct value with probability at least $9/10$.

Therefore, to get a single query to g , we use $O((\log q)/\varepsilon)$ queries. In the context of our previous section, we have $C = 4/\varepsilon$, so $q = O(C^3/\varepsilon^3) = O(1/\varepsilon^6)$ and the overall query complexity of this part is $O((q \log q)/\varepsilon) = O\left(\frac{1}{\varepsilon^7} \log \frac{1}{\varepsilon}\right)$. This dominates the query complexity of the other part of the tester, from [Claim 5.9](#), which is $O\left(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon}\right)$. By a union bound over the part from [Claim 5.9](#), the simulation of g , and the call to the tester of [Lemma 5.1](#), the algorithm is correct with probability at least $1 - 3/10 > 2/3$. \square

Proof of Claim 5.8. We will estimate the error of f_m in computing f on variable blocks and non-variable blocks separately. Each non-variable block B can contribute error on at most $\varepsilon|B|/100$ points. Each variable block B can contribute error on at most $|B| = n/m$ points. The total number of errors is at most $\varepsilon n/100 + s(n/m) = n(\varepsilon/100 + s/m)$, yielding the upper bound on $\operatorname{dist}(f, f_m)$. \square

Proof of Claim 5.9. We first note that given any fixed block B , it is easy to detect whether it is variable (with probability of failure at most δ) by making $O\left(\frac{1}{\varepsilon} \log \frac{1}{\delta}\right)$ uniformly distributed queries in B . Doing so, a variable block will be labelled as such with probability at least $1 - \delta$, while a constant block will never be marked as variable. (If a block is neither constant nor variable, then any answer will do.)

Letting s denote the number of variable blocks, we then want to non-adaptively distinguish between $s \geq \frac{5}{4}k = \frac{5\varepsilon}{16}m$ and $s \leq k = \frac{\varepsilon}{4}m$ (since if f were k -monotone, then f_m had at most k variable blocks). Doing so with probability at least $19/20$ can be done by checking only $q = O\left(\frac{1}{\varepsilon}\right)$ blocks chosen uniformly at random: by the above, setting $\delta = \frac{1}{20q}$ all of the q checks will also yield the correct answer with probability no less than $9/10$, so by a union bound we will distinguish (i) and (ii) with probability at least $9/10$. We conclude by observing that all $O\left(q \cdot \frac{1}{\varepsilon} \log \frac{1}{q}\right) = O\left(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon}\right)$ queries are indeed non-adaptive. \square

6 On the grid

We now turn to the grid, and consider k -monotonicity of functions defined on $[n]^2$. More specifically, in this section we prove [Theorem 1.6](#), giving an adaptive tester for 2-monotonicity with optimal query complexity, before discussing in [Section 6.2](#) possible extensions of these ideas.

6.1 The case $k = 2$

Theorem 1.6. *There exists a two-sided adaptive tester for 2-monotonicity of functions $f: [n]^2 \rightarrow \{0, 1\}$ with query complexity $q(n, \varepsilon) = O\left(\frac{1}{\varepsilon}\right)$.*

Proof. At a high-level, the algorithm relies on two key components: the first is the observation that testing 2-monotonicity of $f: [n]^2 \rightarrow \{0, 1\}$ under some suitable additional assumption on f reduces to (tolerant) testing monotonicity of two *one-dimensional* functions (but with larger range), under the L_1 norm. The second is that, given access to an arbitrary f , one can efficiently provide query access to some function g which satisfies this additional assumption, and such that g will also be close to f whenever f is truly 2-monotone.

Combining the two then enables one to test this function g for 2-monotonicity, and then check whether it is also the case that f and g are sufficiently close. The first step, by the above, can be done efficiently by simulating query access to g , which in turn allows to (with some additional tricks) simulate access to the corresponding one-dimensional functions: and invoke on these two functions the L_1 -tester of [\[BRY14\]](#). (The main challenges there lies in performing this two-level simulation while keeping the number of queries to f low enough; which we achieve by carefully amortizing the queries made overall.)

Details. We hereafter assume without loss of generality that f is identically 0 on the bottom and top rows, that is $f(1, j) = f(n, j) = 0$ for all $j \in [n]$. (Indeed, we can ensure this is the case by adding two extra columns, extending the domain of f to $[n + 2] \times [n]$: note that f remains 2-monotone if it was already, and can only decrease its distance to 2-monotonicity by $O(1/n)$).⁸ For the sake of the proof, we will require the notion of *2-column-wise-monotonicity*, defined below:

⁸This will be used in the proof of [Lemma 6.2](#) and [Lemma 6.4](#).

Definition 6.1. A function $f: [n]^2 \rightarrow \{0, 1\}$ is said to be *2-column-wise-monotone* if, for every $j \in [n]$, its restriction $f_j: [n] \times \{j\} \rightarrow \{0, 1\}$ is 2-monotone. Given such a function f , we define the two sequences $(\bar{\partial}f_j)_{j \in [n]}$ and $(\partial f_j)_{j \in [n]}$ as the sequence of “changepoints” in the columns. More formally, we define

$$\partial f_j = \min \{ i \in [n] : f(i, j) \neq f(1, j) \} - 1, \quad \bar{\partial}f_j = \max \{ i \in [n] : f(i, j) \neq f(n, j) \} + 1$$

for every j such that $f|_{[n] \times \{j\}}$ is not constant. If the column $f|_{[n] \times \{1\}}$ is constant, let $\partial f_1 = 1$ and $\bar{\partial}f_1 = n$; further, if for some $j > 1$ the column $f|_{[n] \times \{j\}}$ is constant, let $\partial f_j = \partial f_{j-1}$ and $\bar{\partial}f_j = \bar{\partial}f_{j-1}$. Note that $\partial f_j \leq \bar{\partial}f_j$ for every column $j \in [n]$.

As it turns out, testing 2-monotonicity of functions *guaranteed to be 2-column-wise-monotone* reduces to testing monotonicity of these two specific subsequences:

Lemma 6.2. *Let $f: [n]^2 \rightarrow \{0, 1\}$ be 2-column-wise-monotone. If f is 2-monotone, then both sequences $(\partial f_j)_{j \in [n]}$ and $(\bar{\partial}f_j)_{j \in [n]}$ are non-increasing. Moreover, if f is ε -far from 2-monotone, then at least one of the two sequences is $\frac{\varepsilon}{2}$ -far from non-increasing (in Hamming distance).*

Proof. The first part of the theorem is straightforward (by contrapositive, if at least one of the two sequences is not non-increasing then we can find a violation of 2-monotonicity). We thus turn to the second part, and show the contrapositive; for this purpose, we require the following result:

Claim 6.3. *If $f: [n]^2 \rightarrow \{0, 1\}$ is a 2-column-wise monotone function such that (i) $f(1, j) = f(n, j) = 0$ for all $j \in [n]$ and (ii) both $(\partial f_j)_{j \in [n]}, (\bar{\partial}f_j)_{j \in [n]} \subseteq [n]$ are non-increasing, then f is 2-monotone.*

Proof. By contradiction, suppose there exists a 2-column-wise monotone function f satisfying (i) and (ii), which is not 2-monotone. This last point implies there exists a triple of comparable elements $x = (i_x, j_x) \prec y = (i_y, j_y) \prec z = (i_z, j_z)$ constituting a violation, i.e. such that $(f(x), f(y), f(z)) = (1, 0, 1)$. Moreover, since (i) holds we must have $1 < i_x \leq i_y \leq i_z < n$; more precisely, $1 \leq \partial f_{j_x} < i_x \leq i_y \leq i_z < \bar{\partial}f_{j_z} \leq n$. As $x \prec y \prec z$, we have $j_x \leq j_y \leq j_z$, which by the non-increasing assumption (ii) implies that $\partial f_{j_x} \geq \partial f_{j_y}$ and $\bar{\partial}f_{j_y} \geq \bar{\partial}f_{j_z}$. But this is not possible, as altogether this leads to $\partial f_{j_y} < i_y < \bar{\partial}f_{j_y}$, i.e. $f(y) = 1$. \square

Assume both sequences $(\partial f_j)_{j \in [n]}, (\bar{\partial}f_j)_{j \in [n]} \subseteq [n]$ are $\frac{\varepsilon}{2}$ -close to non-increasing, and let $L, H \subset [n]$ (respectively) be the set of indices where the two sequences need to be changed in order to become non-increasing. By assumption, $|L|, |H| \leq \frac{\varepsilon n}{2}$, so $|L \cup H| \leq \varepsilon n$. But to “fix” a value of $(\partial f_j)_{j \in [n]}$ or $(\bar{\partial}f_j)_{j \in [n]}$ requires to change the values of the function f inside a single column – and this can be done preserving its 2-column-wise-monotonicity, so that changing the value of f on at most n points is enough. It follows that making both $(\partial f_j)_{j \in [n]}$ and $(\bar{\partial}f_j)_{j \in [n]}$ non-increasing requires to change f on at most εn^2 points, and with **Claim 6.3** this results in a function which is 2-monotone. Thus, f is ε -close to 2-monotone. \square

It is possible to refine the above statement to obtain, in the second case, a more precise characterization in terms of the L_1 distance of these two sequences to monotonicity. For conciseness, in the rest of this section we denote by $\mathcal{M}_k^{(d)}$ the class of k -monotone functions on $[n]^d$, and will omit the subscript when $k = 1$ (i.e., for monotone functions: $\mathcal{M}^{(d)} = \mathcal{M}_1^{(d)}$).

Lemma 6.4. *Let $f: [n]^2 \rightarrow \{0, 1\}$ be 2-column-wise-monotone. If f is ε -far from 2-monotone then at least one of the two sequences is $\frac{\varepsilon}{2}$ -far from non-increasing (in L_1 distance). More precisely:*

$$\text{dist}\left(f, \mathcal{M}_2^{(2)}\right) \leq L_1(\bar{\partial}f, \mathcal{M}^{(1)}) + L_1(\underline{\partial}f, \mathcal{M}^{(1)}) \quad (5)$$

Proof. Recall that we aim at establishing the following:

$$\text{dist}\left(f, \mathcal{M}_2^{(2)}\right) \leq L_1(\bar{\partial}f, \mathcal{M}^{(1)}) + L_1(\underline{\partial}f, \mathcal{M}^{(1)}) \quad (6)$$

For notational convenience, we will view in this proof the sequences $(\underline{\partial}f)_j, (\bar{\partial}f)_j$ as functions $\underline{\partial}f, \bar{\partial}f: [n] \rightarrow [n]$. Let $\ell, h: [n] \rightarrow [n]$ (for “low” and “high,” respectively) be monotone functions achieving $L_1(\underline{\partial}f, \mathcal{M}^{(1)})$ and $L_1(\bar{\partial}f, \mathcal{M}^{(1)})$, respectively.

- As $\underline{\partial}f(j) \leq \bar{\partial}f(j)$ for all $j \in [n]$, we will assume $\ell(j) \leq h(j)$ for all j . Otherwise, one can consider instead the functions $\ell' = \min(\ell, h)$ and $h' = \max(\ell, h)$: both will still be monotone (non-increasing), and by construction

$$|\ell'(j) - \underline{\partial}f(j)| + |h'(j) - \bar{\partial}f(j)| \leq |\ell(j) - \underline{\partial}f(j)| + |h(j) - \bar{\partial}f(j)|$$

for all $j \in [n]$, so that $L_1(\bar{\partial}f, \ell') + L_1(\underline{\partial}f, h') \leq L_1(\bar{\partial}f, \ell) + L_1(\underline{\partial}f, h)$.

- From ℓ and h , we can define a 2-column-wise monotone function $g: [n]^2 \rightarrow [n]$ such that $\underline{\partial}g = \ell$ and $\bar{\partial}g = h$: that is,

$$g(i, j) = \begin{cases} 0 & \text{if } i \geq h(j) \\ 1 & \text{if } \ell(j) < i < h(j) \\ 0 & \text{if } i \leq \ell(j) \end{cases}$$

for $(i, j) \in [n]^2$.

It is clear that g is 2-column-wise monotone with $g(1, j) = g(n, j) = 0$ for all $j \in [n]$; since by construction $\underline{\partial}g, \bar{\partial}g$ are non-increasing, we can invoke [Claim 6.3](#) to conclude g is 2-monotone. It remains to bound the distance between f and g : writing $\Delta_j \in \{0, \dots, n\}$ for the number of points on which f and g differ in the j -th column, we have

$$\begin{aligned} \text{dist}\left(f, \mathcal{M}_2^{(2)}\right) &\leq \text{dist}(f, g) = \frac{1}{n^2} \sum_{j=1}^n \Delta_j \leq \frac{1}{n^2} \sum_{j=1}^n \left(|\ell(j) - \underline{\partial}f(j)| + |h(j) - \bar{\partial}f(j)| \right) \\ &= \frac{1}{n^2} \sum_{j=1}^n |\ell(j) - \underline{\partial}f(j)| + \frac{1}{n^2} \sum_{j=1}^n |h(j) - \bar{\partial}f(j)| = L_1(\underline{\partial}f, \ell) + L_1(\bar{\partial}f, h) \\ &\leq L_1(\underline{\partial}f, \mathcal{M}^{(1)}) + L_1(\bar{\partial}f, \mathcal{M}^{(1)}) \end{aligned}$$

which concludes the proof. \square

Having established these two lemmata, we turn to the proof of [Theorem 1.6](#). We first describe a non-optimal tester making $O(1/\varepsilon^2)$ queries; before explaining how to modify it in order to amortize the number of queries, to yield the desired $O(1/\varepsilon)$ query complexity.

Suppose we are given query access to an arbitrary function $f: [n]^2 \rightarrow \{0, 1\}$. For simplicity, as before we assume without loss of generality that $\frac{\varepsilon n}{16}$ is an integer, and partition each column into

$K \stackrel{\text{def}}{=} \frac{16}{\varepsilon}$ intervals of size $\frac{\varepsilon n}{16}$, that is partition the domain $[n] \times [n]$ into a grid $\lceil \frac{16}{\varepsilon} \rceil \times [n]$ (each column being divided in K blocks B_1, \dots, B_K of size $\frac{\varepsilon n}{16}$). This uniquely defines a function $g: [n]^2 \rightarrow \{0, 1, *\}$: for any point $x = (i, j) \in [n]^2$, we let $\ell \in [K]$ be the index such that $i \in B_\ell = \{b_\ell, \dots, b_{\ell+1} - 1\}$, and set:

- $g(x) = f(i, b_\ell)$, if $f(i, b_\ell) = f(i, b_{\ell+1} - 1)$;
- $g(x) = *$, if $f(i, b_\ell) \neq f(i, b_{\ell+1} - 1)$;

so that g is constant on any “block” $B_\ell \times \{i\}$. Note that we can provide query access to g , at the price of an overhead of 2 queries (to f) per query (to g).

However, this g may not itself be 2-column-wise-monotone; for this reason, we will instead work with a “fixed” version of g which will by construction be 2-column-wise-monotone. In more detail, we define \tilde{g} to be the 2-column-wise-monotone function obtained by the following process.

- First, we (arbitrarily) set the values $*$ to 0, so that g becomes a function $g: [n]^2 \rightarrow \{0, 1\}$.
- Define \tilde{g} by its restriction on each column: letting $(\bar{\partial}g_j)_{j \in [n]}$ and $(\partial g_j)_{j \in [n]}$ be as defined in [Definition 6.1](#) (observing that the quantities are well and uniquely defined even if g is not 2-column-wise-monotone), set

$$\tilde{g}(i, j) = \begin{cases} g(1, j) & \text{if } i \leq \partial g_j \\ 1 - g(1, j) & \text{if } \partial g_j < i < \bar{\partial}g_j \\ g(n, j) & \text{if } i \geq \bar{\partial}g_j \end{cases}$$

From this construction, it is clear that \tilde{g} is 2-column-wise-monotone, and entirely and deterministically determined by g (and therefore by f); moreover we have $\tilde{g} = g$ whenever g is itself 2-column-wise-monotone. Furthermore, any query to \tilde{g} can straightforwardly be answered by making at most queries $O(1/\varepsilon)$ to g , and hence to f .

- If f is 2-monotone, then g is $\frac{\varepsilon}{8}$ -close to f and so is \tilde{g} ; moreover, \tilde{g} is 2-monotone as well. Therefore $\text{dist}(f, \tilde{g}) \leq \frac{\varepsilon}{8}$ and $\text{dist}(\tilde{g}, \mathcal{M}_2^{(2)}) = 0$.
- If f is ε -far from 2-monotone, then \tilde{g} is either (i) $\frac{\varepsilon}{4}$ -far from f or (ii) $\frac{3\varepsilon}{4}$ -far from 2-monotone, since if neither hold then by the triangle inequality f is ε -close to 2-monotone.

The tester (first take). The algorithm now proceeds as follows:

1. simulate query access to \tilde{g} (defined as above) to detect if $\text{dist}(\tilde{g}, \mathcal{M}_2^{(2)}) \geq \varepsilon$ using [Eq.\(5\)](#), with probability of failure $\frac{1}{6}$. More precisely, test monotonicity in L_1 distance of both $\partial \tilde{g}$ and $\bar{\partial} \tilde{g}$ with parameter $\frac{\varepsilon}{64}$, using the (non-adaptive) algorithm of [\[BRY14\]](#); and reject if any of these two tests rejects.
 - If $\text{dist}(f, \mathcal{M}_2^{(2)}) = 0$ (and so in particular \tilde{g} is 2-monotone as well by the first bullet above), then
 - (a) $L_1(\bar{\partial} \tilde{g}, \mathcal{M}^{(1)}) = L_1(\partial \tilde{g}, \mathcal{M}^{(1)}) = 0$, And
 - (b) Estimating that $\text{dist}(f, \tilde{g}) \leq \frac{\varepsilon}{8}$ vs $\text{dist}(f, \tilde{g}) > \frac{\varepsilon}{4}$ reveals the correct answer (the former) with probability at least $5/6$
by [Lemma 6.2](#), and both tests will accept.
 - If $\text{dist}(\tilde{f}, \mathcal{M}_2^{(2)}) > \varepsilon$, (and so in particular \tilde{g} is possibly $\frac{3\varepsilon}{4}$ far from 2-monotone) then

- (a) $\max(L_1(\bar{\partial}\tilde{g}, \mathcal{M}^{(1)}), L_1(\underline{\partial}\tilde{g}, \mathcal{M}^{(1)})) > \frac{3\varepsilon}{8}$, Or
- (b) Estimating that $\text{dist}(f, \tilde{g}) \leq \frac{\varepsilon}{8}$ vs $\text{dist}(f, \tilde{g}) > \frac{\varepsilon}{4}$ reveals the correct answer (the latter) with probability at least $5/6$

So, in this case, at least one of the two tests rejects.

2. return ACCEPT if both of the two tests above passed, REJECT otherwise.

By a union bound, the tester is then correct with probability at least $\frac{2}{3}$; its query complexity is

$$2t \cdot O\left(\frac{1}{\varepsilon}\right) + \left(t' \cdot O\left(\frac{1}{\varepsilon}\right) + O\left(\frac{1}{\varepsilon}\right)\right)$$

where t and t' are respectively the cost of simulating query access to $\bar{\partial}g, \underline{\partial}g$, and to \tilde{g} ; as the first step, testing in L_1 for functions defined on the line $[n]$, has query complexity $O(1/\varepsilon)$ from [BRY14]. Taking $t = t' = O(1/\varepsilon)$ as discussed above then results in a query complexity of $O(1/\varepsilon^2)$.

However, as mentioned previously this is not optimal: as we shall see, we can modify this to obtain instead an $O(1/\varepsilon)$ query complexity. In order to *amortize* the overall query complexity, we define the following process that specifies a 2-column-wise-monotone function \hat{g} :

The Amortized Tester The idea here is to have our previous tester simulate query access to a new function \hat{g} instead \tilde{g} . Simulating a single query \tilde{g} took $O(1/\varepsilon)$ queries to f and this is where amortizing queries using \hat{g} reduces the number of queries. Details follow.

- (*Initialization*) Let $j_1, \dots, j_{1/\varepsilon+1} \in [n]$ be the indices defined by $j_\ell = (\ell - 1) \cdot \varepsilon n + 1$ for $\ell \in [1/\varepsilon]$, and $j_{1/\varepsilon+1} = n$.
- (*\hat{g} on j_1 -st column*) Obtain all the values of \tilde{g} on the j_1 -st column $[n] \times \{j_1\}$, at a cost of $O(1/\varepsilon)$ queries to f , to find $\bar{\partial}\hat{g}_{j_1}, \underline{\partial}\hat{g}_{j_1}$. Define \hat{g} on this column accordingly.
- (*\hat{g} on remaining columns*) Assuming \hat{g} has been defined on the j_ℓ -th column, define it on the $j_{\ell+1}$ -th column: starting at the “vertical” positions of the two changepoints $\bar{\partial}\hat{g}_{j_\ell}, \underline{\partial}\hat{g}_{j_\ell}$ of the previous column, start querying “downwards” the values of g on the $j_{\ell+1}$ -th column until candidate values for $\bar{\partial}\hat{g}_{j_{\ell+1}}, \underline{\partial}\hat{g}_{j_{\ell+1}}$ consistent with g are found or the bottom of the column is reached (in which case the corresponding changepoint $\bar{\partial}\hat{g}_{j_{\ell+1}}$ or $\underline{\partial}\hat{g}_{j_{\ell+1}}$ is set to 1). After this, define \hat{g} on the $j_{\ell+1}$ -th column to be consistent with these (at most) two changepoints.

Note that the queries made in this step are adaptive, and the (partial) function \hat{g} obtained at the end coincides with \tilde{g} if f (and therefore \tilde{g}) is indeed 2-monotone. This is because in this case, by Lemma 6.2 the sequences $(\bar{\partial}\tilde{g}_j)_j, (\underline{\partial}\tilde{g}_j)_j$ will be non-decreasing, and therefore the process outlined above will result in $\bar{\partial}\hat{g}_{j_\ell} = \bar{\partial}\tilde{g}_{j_\ell}$ and $\underline{\partial}\hat{g}_{j_\ell} = \underline{\partial}\tilde{g}_{j_\ell}$ for all $1 \leq \ell \leq 1/\varepsilon + 1$. Moreover, it is not difficult to see that the total number of queries made to f in this initialization step will be $O(1/\varepsilon)$: this is because of the fact that we only search for the current changepoint $\bar{\partial}\hat{g}_{j_\ell}$ (resp. $\underline{\partial}\hat{g}_{j_\ell}$) starting at the position of the previous one $\bar{\partial}\hat{g}_{j_{\ell-1}}$ (resp. $\underline{\partial}\hat{g}_{j_{\ell-1}}$), going downwards only. Since to obtain a changepoint we only query the positions of g (and therefore f) at block endpoints, there are in total at most $1/\varepsilon$ positions to query where starting at one and then only going “down.” Thus, the number of queries made for each column j_ℓ can be written as $O(1) + m_\ell$, where $\sum_{\ell=1}^{1/\varepsilon+1} m_\ell \leq K = O(1/\varepsilon)$.

Query time. When querying the value of \hat{g} on a point (i, j) , first let ℓ be the index such that $j_\ell \leq j < j_{\ell+1}$. Then define $\bar{\partial}\hat{g}_j$ (resp. $\underline{\partial}\hat{g}_j$) by querying the value of g on the j -th column for all at most $1/\varepsilon$ (block) indices between $\bar{\partial}\hat{g}_{j_\ell}$ and $\bar{\partial}\hat{g}_{j_{\ell+1}}$ (resp. $\underline{\partial}\hat{g}_{j_\ell}$ and $\underline{\partial}\hat{g}_{j_{\ell+1}}$) to find the corresponding candidate changepoints:

$$\begin{aligned}\underline{\partial}\hat{g}_j &= \min \left\{ \underline{\partial}\hat{g}_{j_{\ell+1}} \leq i \leq \underline{\partial}\hat{g}_{j_\ell} : g(i, j) \neq g(1, j) \right\} - 1, \\ \bar{\partial}\hat{g}_j &= \max \left\{ \bar{\partial}\hat{g}_{j_{\ell+1}} \leq i \leq \bar{\partial}\hat{g}_{j_\ell} : g(i, j) \neq g(n, j) \right\} + 1\end{aligned}$$

Again, note that after each query the (partial) function \hat{g} obtained so far will coincide with \tilde{g} if f (and therefore \tilde{g}) is indeed 2-monotone; and \hat{g} is uniquely determined by f (and in particular does not depend on the actual queries made nor on their order). Finally, the function \hat{g} thus defined will always by construction be 2-column-wise-monotone.

The tester previously described can then be slightly modified to simulate access to \hat{g} instead of \tilde{g} : by the above discussion, for the completeness case we will have the same guarantees as then $\hat{g} = \tilde{g}$, while the soundness case stays unchanged:

- If f is 2-monotone, then $\hat{g} = \tilde{g}$ is $\frac{\varepsilon}{8}$ -close to f ; so $\text{dist}(f, \hat{g}) \leq \frac{\varepsilon}{8}$ and $\text{dist}(\hat{g}, \mathcal{M}_2^{(2)}) \leq \frac{\varepsilon}{8}$.
- If f is ε -far from 2-monotone, then \hat{g} is either (i) $\frac{\varepsilon}{4}$ -far from f or (ii) $\frac{3\varepsilon}{4}$ -far from 2-monotone.

Thus, the analysis of correctness of the tester carries through with this modification; it only remains to bound the query complexity. We will show that the *expected* number of queries made is $O(1/\varepsilon)$; a bound on the worst-case query complexity will then follow from standard arguments,⁹ at the price of a constant factor in the $O(\cdot)$.

To give this bound on the expected number of queries, we first observe that the algorithm from [BRY14] we rely on in the first stage of the tester is non-adaptive, and moreover all the queries it makes are uniformly distributed (as it works by a reduction, invoking the non-adaptive, one-sided, sample-based monotonicity tester for functions $[n] \rightarrow \{0, 1\}$). Similarly, all the queries made in the second stage are uniformly distributed as well.

Therefore, the expected number of queries a_ℓ made to columns with indices $j_\ell \leq j < j_{\ell+1}$ is the same for each $\ell \in [1/\varepsilon + 1]$, namely

$$a_\ell = \frac{q}{1/\varepsilon} = O(1).$$

where $q = O\left(\frac{1}{\varepsilon}\right)$ is the total number of queries made to \hat{g} (and/or to $\bar{\partial}\hat{g}, \underline{\partial}\hat{g}$) during the second phase of ‘‘Query time.’’ Now, letting $m_\ell = \underline{\partial}\hat{g}_{j_\ell} - \underline{\partial}\hat{g}_{j_{\ell+1}}$ and $m'_\ell = \bar{\partial}\hat{g}_{j_\ell} - \bar{\partial}\hat{g}_{j_{\ell+1}}$, we have that the total expected cost of simulating these queries (in terms of queries to f) is upperbounded by

$$\sum_{\ell=1}^{\frac{1}{\varepsilon}} a_\ell \cdot (O(1) + m_\ell + m'_\ell) = O\left(\frac{1}{\varepsilon}\right) + O(1) \cdot \sum_{\ell=1}^{\frac{1}{\varepsilon}} m_\ell + O(1) \cdot \sum_{\ell=1}^{\frac{1}{\varepsilon}} m'_\ell = O\left(\frac{1}{\varepsilon}\right).$$

Since the total number of queries made to f is the sum of the number of queries made to partly build \hat{g} during the ‘‘Initialization phase’’ (which is $O(1/\varepsilon)$ by the foregoing discussion), the number of queries made to simulate access to \hat{g} or $\bar{\partial}\hat{g}, \underline{\partial}\hat{g}$ during the ‘‘Query time’’ (which was just shown to be $O(1/\varepsilon)$ in expectation), and the number of queries directly made to f when testing the distance of f to \hat{g} (which is also $O(1/\varepsilon)$), the expected total number of queries is indeed $O(1/\varepsilon)$, as claimed. \square

⁹Namely, stopping the algorithm and outputting REJECT if the number of queries made exceeds $\frac{C}{\varepsilon}$ for some absolute constant $C > 0$, found by applying Markov’s inequality.

6.2 Possible extensions

We now discuss two possible extensions of the techniques underlying [Theorem 1.6](#), namely to (i) k -monotonicity testing of functions over $[n]^2$, for general k ; and (ii) 2-monotonicity testing of functions over $[n]^d$, for general d . (Note that we do provide a different tester for general k and d in the next section, [Section 7](#)).

Extending to general k (for $d = 2$). A natural direction would be to first to generalize [Definition 6.1](#) to k -column-wise monotone functions f , defining the k sequences $(\partial f_j^{(1)})_{j \in [n]}, \dots, (\partial f_j^{(k)})_{j \in [n]}$ of column changepoints with $\partial f_j^{(1)} \leq \dots \leq \partial f_j^{(k)}$ for all $j \in [n]$. The next step would then be to obtain an analogue of the key lemma of the previous section, [Lemma 6.4](#) to this setting. An issue is that it appears necessary to consider now the L_1 distance to $(k-1)$ -monotonicity of these k sequences, instead of monotonicity as before. Thus, taking this route requires to generalize the definition of k -monotonicity to real-valued functions, but also to develop L_1 -testers for k -monotonicity over the line.

The testing algorithm now follows the same outline as in the previous section, with the same “amortizing” idea when invoking this newly obtained L_1 -tester for k -monotonicity in parallel on the k subsequences, each with probability of failure $\delta = 1/(10k)$ (for a union bound) and approximation parameter $\varepsilon' = \varepsilon/(10k)$. (Note that some more optimizations may then help further reduce the query complexity, by “sharing” the same set of queries between the k instances of the L_1 -testing algorithm.)

Extending to general d (for $k = 2$). At a very high-level, the tester of [Section 6.1](#) works by reducing 2-monotonicity testing of $f: [n]^2 \rightarrow \{0, 1\}$ to monotonicity L_1 -testing of $\partial f, \bar{\partial} f: [n] \rightarrow [0, 1]$. More generally, one can hope to extend this approach to higher dimensions, reducing 2-monotonicity testing of $f: [n]^d \rightarrow \{0, 1\}$ to monotonicity L_1 -testing of $\partial f, \bar{\partial} f: [n]^{d-1} \rightarrow [0, 1]$: that is, testing monotonicity (in L_1) of the two $(d-1)$ -dimensional “surfaces” of changepoints. This in turn could be done invoking the L_1 non-adaptive tester of [\[BRY14\]](#) for monotonicity over $[n]^d$, which has query complexity $\tilde{O}(d/\varepsilon)$: which may lead to a total query complexity of $\text{poly}(d, 1/\varepsilon)$, that is *polynomial in the dimension*. We leave this possible extension as an interesting direction for future work.

7 On the high-dimensional grid

In this section, we give two algorithms for tolerant testing, that is testing whether a function $f: [n]^d \rightarrow \{0, 1\}$ is ε_1 -close to k -monotone vs. ε_2 -far from k -monotone, establishing [Theorem 1.8](#). The first has query complexity exponential in the dimension d and is *fully tolerant*, that is works for any setting of $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$. The second applies whenever $\varepsilon_2 > 3\varepsilon_1$, and has (incomparable) query complexity exponential in $\tilde{O}(k\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2)$. Both of these algorithms can be used for non-tolerant (“regular”) testing by setting $\varepsilon_1 = 0$ and $\varepsilon_2 = \varepsilon$, which implies [Theorem 1.7](#).

Theorem 1.8. *There exist*

- a non-adaptive (fully) tolerant tester for k -monotonicity of functions $f: [n]^d \rightarrow \{0, 1\}$ with query complexity $q(n, d, \varepsilon_1, \varepsilon_2, k) = \tilde{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \left(\frac{5kd}{\varepsilon_2 - \varepsilon_1}\right)^d\right)$;

- a non-adaptive tolerant tester for k -monotonicity of functions $f: [n]^d \rightarrow \{0, 1\}$ with query complexity $q(n, d, \varepsilon_1, \varepsilon_2, k) = 2^{\tilde{O}(k\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2)}$, under the restriction that $\varepsilon_2 > 3\varepsilon_1$.

As a corollary, this implies [Theorem 1.7](#), restated below:

Theorem 1.7. *There exists a non-adaptive tester for k -monotonicity of functions $f: [n]^d \rightarrow \{0, 1\}$ with query complexity $q(n, d, \varepsilon, k) = \min(\tilde{O}\left(\frac{1}{\varepsilon^2} \left(\frac{5kd}{\varepsilon}\right)^d\right), 2^{\tilde{O}(k\sqrt{d}/\varepsilon^2)})$.*

For convenience, we will view in this part of the paper the set $[n]$ as $[n] = \{0, 1, \dots, n-1\}$. Assuming that m divides n , we let $\mathcal{B}_{m,n}: [n]^d \rightarrow [m]^d$ be the mapping such that $\mathcal{B}_{m,n}(y)_i = \lfloor my_i/n \rfloor$ for $1 \leq i \leq m$. For $x \in [m]^d$, we define the set $\mathcal{B}_{m,n}^{-1}(x)$ to be the inverse image of x . Specifically, $\mathcal{B}_{m,n}^{-1}(x)$ is the set of points of the form $m \cdot x + [n/m]^d$, with standard definitions for scalar multiplication and coordinate-wise addition. That is, $\mathcal{B}_{m,n}^{-1}(x)$ is a ‘‘coset’’ of $[n/m]^d$ points in $[n]^d$. To keep on with the notations of the other sections, we will call these cosets *blocks*, and will say a function $h: [n]^d \rightarrow \{0, 1\}$ is an m -block function if it is constant on each block. Moreover, for clarity of presentation, we will omit the subscripts on \mathcal{B} and \mathcal{B}^{-1} whenever they are not necessary.

We first establish a lemma that will be useful for the proofs of correctness of both algorithms.

Lemma 7.1. *Suppose $f: [n]^d \rightarrow \{0, 1\}$ is k -monotone. Then there is an m -block function $h: [n]^d \rightarrow \{0, 1\}$ such that $\text{dist}(f, h) < kd/m$.*

Proof. Fix any k -monotone function $f: [n]^d \rightarrow \{0, 1\}$. We partition $[m]^d$ into chains of the form

$$C_x = \left\{ x + \ell \cdot \mathbf{1}^d : \ell \in \mathbb{N}, x \in [m]^d \text{ and } x_i = 0 \text{ for some } i \right\}.$$

There are $m^d - (m-1)^d \leq dm^{d-1}$ of these chains: we will show that f can only be nonconstant on at most k blocks of each chain.

By contradiction, suppose there exists $x \in [m]^d$ such that f is nonconstant on $k+1$ different blocks $\mathcal{B}^{-1}(z^{(i)})$, where $z^{(1)} \prec z^{(2)} \prec \dots \prec z^{(k)} \prec z^{(k+1)}$, and each $z^{(i)} \in C_x$. By construction, we have $\mathcal{B}^{-1}(z^{(i)}) \prec \mathcal{B}^{-1}(z^{(j)})$ for $i < j$. For each $1 \leq i \leq k+1$, there are two points $v_*^{(i)}, v_{**}^{(i)} \in \mathcal{B}^{-1}(z_i)$ such that $v_*^{(i)} \prec v_{**}^{(i)}$ and $f(v_*^{(i)}) \neq f(v_{**}^{(i)})$. By construction $v_*^{(1)} \prec v_{**}^{(1)} \prec v_*^{(2)} \prec v_{**}^{(2)} \prec v_*^{(3)} \prec v_{**}^{(3)} \prec \dots \prec v_*^{(k+1)} \prec v_{**}^{(k+1)}$, and there must be at least $k+1$ pairs of consecutive points with differing function values. Out of these $2k+2$ many points, there is a chain of points $\bar{v}^{(1)} \prec \bar{v}^{(2)} \prec \dots \prec \bar{v}^{(k+1)}$ where $f(\bar{v}^{(i)}) \neq f(\bar{v}^{(i+1)})$ for $1 \leq i \leq k$, which is a violation of the k -monotonicity of f .

Thus, in each of the dm^{d-1} many chains of blocks, there can only be k nonconstant blocks. It follows that there are at most $kd m^{d-1}$ nonconstant blocks in total. We now define $h(y)$ to be equal to $f(y)$ if f is constant on $\mathcal{B}(y)$, and arbitrarily set $h(y) = 0$ otherwise. Each set $\mathcal{B}^{-1}(y)$ contains $(n/m)^d = n^d \cdot m^{-d}$ many points, and f is not constant on at most $kd m^{d-1}$ of these. It follows that $\text{dist}(f, h) \leq kd m^{d-1} \cdot m^{-d} = kd/m$. \square

7.1 Fully tolerant testing with $O(kd/(\varepsilon_2 - \varepsilon_1))^d$ queries

Our first algorithm ([Algorithm 1](#)) then proceeds by essentially brute-force learning an m -block function close to the unknown function, and establishes the first item of [Theorem 1.8](#).

Proposition 7.2. *[Algorithm 1](#) accepts all functions ε_1 -close to k -monotone functions, and rejects all functions ε_2 -far from k -monotone (with probability at least $2/3$). Its query complexity is $O\left(\frac{d}{(\varepsilon_2 - \varepsilon_1)^2} \left(\frac{5kd}{\varepsilon_2 - \varepsilon_1} + 1\right)^d \log \frac{kd}{\varepsilon_2 - \varepsilon_1}\right)$.*

Algorithm 1 Fully tolerant testing with $O(kd/(\varepsilon_2 - \varepsilon_1))^d$ queries.

Require: Query access to $f: [n]^d \rightarrow \{0, 1\}$, $\varepsilon_2 > \varepsilon_1 \geq 0$, a positive integer k

- 1: $\alpha \leftarrow (\varepsilon_2 - \varepsilon_1)$, $m \leftarrow \lceil 5kd/\alpha \rceil$, $t \leftarrow \lceil 25 \ln(6m^d)/(2\alpha^2) \rceil$
 - 2: \triangleright Define a distribution D over $[m]^d \times \{0, 1\}$.
 - 3: **for** $x \in [m]^d$ **do**
 - 4: Query f on t random points $T_x \subseteq \mathcal{B}^{-1}(x)$.
 - 5: $D(x, 0) \leftarrow \Pr_{y \in T_x} [f(y) = 0] / m^d$
 - 6: $D(x, 1) \leftarrow \Pr_{y \in T_x} [f(y) = 1] / m^d$
 - 7: **end for**
 - 8: \triangleright Define a distribution D' over $[n]^d \times \{0, 1\}$ such that $D'(y, b) = D(\mathcal{B}(y), b) \cdot m^d/n^d$.
 - 9: **if** there exists a k -monotone m -block function h such that $\Pr_{(y,b) \sim D'} [h(y) \neq b] \leq \varepsilon_1 + \frac{\alpha}{2}$ **then**
 return ACCEPT
 - 10: **end if**
 - 11: **return** REJECT
-

Proof. The algorithm first estimates $\Pr_{y \in \mathcal{B}^{-1}(x)} [f(y) = b]$ for every $x \in [m]^d$ and $b \in \{0, 1\}$ to within $\pm \frac{\alpha}{5}$. We use $t = 25 \ln(6m^d)/2\alpha^2$ points in each block to ensure (by an additive Chernoff bound) that each estimate is correct except with probability at most $m^{-d}/3$. By a union bound, the probability that all estimates are correct is at least $2/3$, and we hereafter condition on this. By construction, $\mathbb{E}_{(x,b) \sim D} [\Pr_{y \in \mathcal{B}^{-1}(x)} [f(y) \neq b]] = \Pr_{(y,b) \sim D'} [f(y) \neq b] \leq \frac{\alpha}{5}$. In this probability experiment, the marginal distribution of D' on y is uniform over $[n]^d$.

Let $f^*: [n]^d \rightarrow \{0, 1\}$ be a k -monotone function minimizing $\Pr[f(y) \neq f^*(y)]$. [Lemma 7.1](#) ensures that there is a k -monotone m -block function $h: [n]^d \rightarrow \{0, 1\}$ such that $\text{dist}(f^*, h) < kd/m \leq \alpha/5$. Let $h^*: [n]^d \rightarrow \{0, 1\}$ be a k -monotone m -block function minimizing $\text{dist}(f^*, h^*)$.

Completeness. Suppose $\text{dist}(f, f^*) \leq \varepsilon_1$. Then by the triangle inequality,

$$\Pr_{(y,b) \sim D'} [h^*(y) \neq b] \leq \Pr_{(y,b) \sim D'} [h^*(y) \neq f^*(y)] + \Pr_{(y,b) \sim D'} [f^*(y) \neq f(y)] + \Pr_{(y,b) \sim D'} [f(y) \neq b] \leq \varepsilon_1 + \frac{2\alpha}{5}.$$

where to bound the first term $\Pr_{(y,b) \sim D'} [h^*(y) \neq f^*(y)]$ by $\text{dist}(f^*, h^*) \leq \alpha/5$ we used the fact that the marginal distribution of y is uniform when $(y, b) \sim D'$. Thus, the algorithm will find a k -monotone m -block function close to D (without using any queries to f) and accept.

Soundness. Suppose $\text{dist}(f, f^*) \geq \varepsilon_2$. Then by the triangle inequality

$$\begin{aligned} \Pr_{(y,b) \sim D'} [h(y) \neq b] &\geq \Pr_{(y,b) \sim D'} [h(y) \neq f(y)] - \Pr_{(y,b) \sim D'} [f(y) \neq b] \\ &\geq \Pr_{(y,b) \sim D'} [f^*(y) \neq f(y)] - \Pr_{(y,b) \sim D'} [f(y) \neq b] \\ &\geq \varepsilon_2 - \frac{\alpha}{5} \end{aligned}$$

for every k -monotone m -block function h . Since $\varepsilon_2 - 2\alpha/5 \geq \varepsilon_1 + 3\alpha/5$, the algorithm never find a k -monotone m -block function h with low error with respect to D , and the algorithm will reject.

Query complexity. The algorithm only makes queries in constructing D ; the number of queries required is $m^d \cdot t = O\left(\frac{d}{\alpha^2} \left(\frac{5kd}{\alpha} + 1\right)^d \log \frac{kd}{\alpha}\right)$. \square

7.2 Tolerant testing *via* agnostic learning

We now present our second algorithm, [Algorithm 2](#), proving the second item of [Theorem 1.8](#). At its core is the use of an *agnostic learning algorithm* for k -monotone functions, which we first describe.¹⁰

Proposition 7.3. *There exists an agnostic learning algorithm for k -monotone functions over $[r]^d \rightarrow \{0, 1\}$ with excess error τ with sample complexity $\exp(\tilde{O}(k\sqrt{d}/\tau^2))$.*

We start the proof of [Proposition 7.3](#) by formally presenting the learning algorithm.

Algorithm 2 Multiplicative approximation with $\exp(\tilde{O}(k\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2))$ queries

Require: Query access to $f: [n]^d \rightarrow \{0, 1\}$, $\varepsilon_2 > 3\varepsilon_1 \geq 0$, a positive integer k

- 1: $\alpha \leftarrow (\varepsilon_2 - 3\varepsilon_1)$, $m \leftarrow \lceil 6kd/\alpha \rceil$, $t \leftarrow \lceil 3d(k+1)/\alpha \ln m + \ln 100 \rceil$
 - 2: Define D to be the distribution over $[m]^d \times \{0, 1\}$ such that $D(x, b) = \Pr_{y \in \mathcal{B}^{-1}(x)} [f(y) = b]$.
 - 3: $\triangleright \mathcal{A}_D(\tau, f)$ denotes the output of an agnostic learner of k -monotone functions with respect to D , with excess error τ and probability of failure $1/10$
 - 4: $h: [m]^d \rightarrow \mathbb{R} \leftarrow \mathcal{A}_D(\alpha/12, f)$.
 - 5: Estimate $\Pr_{(x,b) \sim D} [h(x) \neq b]$ to within $\pm\alpha/7$ with probability of failure $1/10$, using $O(1/\alpha^2)$ queries.
 - 6: **if** the estimate is more than $\varepsilon_1 + \frac{5\alpha}{12}$ **then return REJECT**
 - 7: **end if**
 - 8: **if** $\text{dist}(h, \ell) = \Pr_{x \in [m]^d} [h(x) \neq \ell(x)] \leq 2\varepsilon_1 + \frac{5\alpha}{12}$ for some k -monotone m -block function ℓ **then return ACCEPT**
 - 9: **else return REJECT**
 - 10: **end if**
-

To prove its correctness we will rely on tools from Fourier analysis. For this reason, it will be convenient in this section to view the range as $\{-1, 1\}$ instead of $\{0, 1\}$.

The idea of the algorithm is to first attempt to learn the function, before checking whether what was learned is indeed close to both the unknown function and the set of k -monotone functions. To perform the first step efficiently enough, we need to design an agnostic learning algorithm for k -monotone functions, which is what most of this section will be dedicated to. In more detail, we will rely on a theorem of [\[KKMS08\]](#), which very broadly speaking states that any class of functions \mathcal{C} that can be well-approximated by linear combinations of a small number of functions admits a good agnostic learner (see [Theorem 7.10](#) for the formal statement). Given this, our goal is to establish the existence of a small set of functions whose linear combinations “cover” the class of k -monotone functions. To do so, we first generalize a theorem of Blais et al. [\[BCO⁺15\]](#), to show that k -monotone functions on the hypergrid have small influence ([Lemma 7.6](#)). We then use this in

¹⁰Recall that an *agnostic learner with excess error τ* for some class of functions \mathcal{C} is an algorithm that, given an unknown distribution D , an unknown arbitrary function f , and access to random labelled samples $\langle x, f(x) \rangle$ where $x \sim D$, satisfies the following. It outputs a hypothesis function \hat{h} such that $\Pr_{x \sim D} [f(x) \neq \hat{h}(x)] \leq \text{OPT}_D + \tau$ with probability at least $2/3$, where $\text{OPT}_D = \min_{h \in \mathcal{C}} \Pr_{x \sim D} [f(x) \neq h(x)]$ (i.e., it performs “almost as well as the best function in \mathcal{C} ”).

Lemma 7.8 to get that k -monotone functions have Fourier concentration on relatively small degree coefficients, which we leverage in **Lemma 7.9** to obtain our small set of functions.

7.3 Preliminary Fourier analysis tools

Definition 7.4. For a Boolean function $f: [r]^d \rightarrow \{-1, 1\}$, we define

$$\mathbf{Inf}_i[f] = 2 \Pr \left[f(x) \neq f(x^{(i)}) \right]$$

where $x = (x_1, x_2, \dots, x_d)$ is a uniformly random string over $[r]^d$, and $x^{(i)} = (x_1, x_2, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_d)$ for x' drawn independently and uniformly from $[r]$. We also define $\mathbf{Inf}[f] = \sum_{i=1}^d \mathbf{Inf}_i[f]$.

We first generalize the following result, due to Blais *et al.*, for more general domains:

Proposition 7.5 ([BCO⁺15]). *Let $f: \{0, 1\}^d \rightarrow \{-1, 1\}$ be a k -monotone function. Then $\mathbf{Inf}[f] \leq k\sqrt{d}$.*

Lemma 7.6 (Generalization). *Let $f: [r]^d \rightarrow \{-1, 1\}$ be a k -monotone function. Then $\mathbf{Inf}[f] \leq k\sqrt{d}$.*

Proof. For any two strings $y^0, y^1 \in [r]^d$, let $f_{y^0, y^1}: \{0, 1\}^d \rightarrow \{-1, 1\}$ be the function obtained by setting $f_{y^0, y^1}(x) = f(y^x)$, where $y^x \in [r]^d$ is defined as

$$y_i^x = \begin{cases} \min\{y_i^0, y_i^1\} & \text{if } x_i = 0 \\ \max\{y_i^0, y_i^1\} & \text{if } x_i = 1 \end{cases}$$

Since f was a k -monotone function, so is f_{y^0, y^1} . Thus $\mathbf{Inf}[f_{y^0, y^1}] \leq k\sqrt{d}$ for every choice of y^0 and y^1 . It is not hard to see that for any fixed $i \in [d]$ the following two processes yield the same distribution over $[r]^d \times [r]^d$:

- Draw $z \in [r]^d, z'_i \in [r]$ independently and uniformly at random, set $z' \stackrel{\text{def}}{=} (z_1, \dots, z_{i-1}, z'_i, z_{i+1}, \dots, z_d)$, and output (z, z') ;
- Draw $y^0, y^1 \in [r]^d, x \in \{0, 1\}^d$ independently and uniformly at random, and output $(y^x, y^{x^{(i)}})$.

This implies that

$$\begin{aligned} \mathbf{Inf}[f] &= \sum_{i=1}^d \mathbf{Inf}_i[f] = \sum_{i=1}^d 2 \Pr_{z \in [r]^d} [f(z) \neq f(z^{(i)})] = \sum_{i=1}^d 2 \mathbb{E}_{y^0, y^1 \in [r]^d} \left[\Pr_{x \in \{0, 1\}^d} [f(y^x) \neq f(y^{x^{(i)}})] \right] \\ &= \mathbb{E}_{y^0, y^1 \in [r]^d} \left[\sum_{i=1}^d 2 \Pr_{x \in \{0, 1\}^d} [f(y^x) \neq f(y^{x^{(i)}})] \right] = \mathbb{E}_{y^0, y^1 \in [r]^d} \left[\sum_{i=1}^d 2 \Pr_{x \in \{0, 1\}^d} [f_{y^0, y^1}(x) \neq f_{y^0, y^1}(x^{(i)})] \right] \\ &= \mathbb{E}_{y^0, y^1} [\mathbf{Inf}[f_{y^0, y^1}]] \leq \mathbb{E}_{y^0, y^1} [k\sqrt{d}] = k\sqrt{d}. \end{aligned}$$

□

For two functions $f, g: [r]^d \rightarrow \mathbb{R}$, we define the inner product $\langle f, g \rangle = \mathbb{E}_x[f(x)g(x)]$, where the expectation is taken with respect to the uniform distribution. It is known that for functions $f: [r]^d \rightarrow \mathbb{R}$, there is a “Fourier basis” of orthonormal functions f . To construct such a basis, we can take any orthonormal basis $\{\phi_0 \equiv 1, \phi_1, \dots, \phi_{|r|-1}\}$ for functions $f: [r] \rightarrow \mathbb{R}$. Given such a basis,

a Fourier basis is the collection of functions ϕ_α , where $\alpha \in [r]^d$, and $\phi_\alpha(x) = \prod_{i=1}^d \phi_{\alpha_i}(x_i)$. Then every $f: [r]^d \rightarrow \mathbb{R}$ has a unique representation $f = \sum_{\alpha \in [r]^d} \hat{f}(\alpha) \phi_\alpha$, where $\hat{f}(\alpha) = \langle f, \phi_\alpha \rangle \in \mathbb{R}$.

Many Fourier formulæ hold in arbitrary Fourier bases, an important example being Parseval's Identity: $\sum_{\alpha \in [r]^d} \hat{f}(\alpha)^2 = 1$. We will use the following property:

Lemma 7.7 ([O'D14, Proposition 8.23]). *For $\alpha \in [r]^d$, let $|\alpha|$ denote the number of nonzero coordinates in α . Then we have*

$$\mathbf{Inf}[f] = \sum_{\alpha \in [r]^d} |\alpha| \hat{f}(\alpha)^2.$$

Lemma 7.8. *If $\mathbf{Inf}[f] \leq k$, then $\sum_{\alpha: |\alpha| > k/\varepsilon} \hat{f}(\alpha)^2 \leq \varepsilon$.*

Proof. If not, then $\mathbf{Inf}[f] = \sum_{\alpha} |\alpha| \hat{f}(\alpha)^2 \geq \sum_{\alpha: |\alpha| > k/\varepsilon} |\alpha| \hat{f}(\alpha)^2 \geq \frac{k}{\varepsilon} \sum_{\alpha: |\alpha| > k/\varepsilon} \hat{f}(\alpha)^2 > \frac{k}{\varepsilon} \cdot \varepsilon = k$, a contradiction. \square

Lemma 7.9. *Let p be the function $\sum_{\alpha: |\alpha| \leq t} \hat{f}(\alpha) \phi_\alpha$. Then*

- (i) $\|p - f\|_2^2 = \mathbb{E}_{x \in [r]^d} [(p(x) - f(x))^2] = \sum_{\alpha: |\alpha| > t} \hat{f}(\alpha)^2$;
- (ii) p is expressible as a linear combination of real-valued functions over $[r]^d$, each of which only depends on at most t coordinates;
- (iii) p is expressible as a degree- t polynomial over the rd indicator functions $\mathbb{1}_{\{x_i=j\}}$ for $1 \leq i \leq d$ and $j \in [r]$.

Proof. The lemma follows immediately from Parseval's identity and from the definitions. \square

Theorem 7.10 ([KKMS08, Theorem 5]). *Let \mathcal{C} be a class of Boolean functions over \mathcal{X} and \mathcal{S} a collection of real-valued functions over \mathcal{X} such that for every $f: \mathcal{X} \rightarrow \{-1, 1\}$ in \mathcal{C} , there exists a function $p: \mathcal{X} \rightarrow \mathbb{R}$ such that p is expressible as a linear combination of functions from \mathcal{S} and $\|p - f\|_2^2 \leq \tau^2$. Then there is an agnostic learning algorithm for \mathcal{C} achieving excess error τ which has sample complexity $\text{poly}(|\mathcal{S}|, 1/\tau)$.*

Importantly, this algorithm is still successful with inconsistent labelled samples (examples), as long as they come from a distribution on $\mathcal{X} \times \{-1, 1\}$, where the marginal distribution on \mathcal{X} is uniform.

7.4 Finishing the proof of Proposition 7.3

Now we put all the pieces together. To agnostically learn a k -monotone function, we simply perform the agnostic learning algorithm of [KKMS08] on the distribution D over $[m]^d \times \{-1, 1\}$ defined by

$$D(x, b) = \Pr_{y \in \mathcal{B}^{-1}(x)} [f(y) = b].$$

To generate a sample (x, b) from D , we draw a uniformly random string in $x \in [m]^d$, and b is the result of a query for the value of $f(y)$ for a uniformly random $y \in \mathcal{B}^{-1}(x)$. From Lemma 7.9, we can take \mathcal{S} to be the set of $(k\sqrt{d}/\tau^2)$ -way products of rd indicator functions. It follows that $|\mathcal{S}| = \binom{rd}{k\sqrt{d}/\tau^2} = \exp(\tilde{O}(k\sqrt{d}/\tau^2))$ (this holds because $r = O\left(\frac{k}{\varepsilon_2 - 3\varepsilon_1}\right)$)

Proposition 7.11. *Algorithm 2 accepts all functions ε_1 -close to k -monotone functions, and rejects all functions ε_2 -far from k -monotone, when $\varepsilon_2 > 3\varepsilon_1$ (with probability at least $2/3$). Its query complexity is $\exp(\tilde{O}(k\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2))$.*

Proof. By a union bound, we have that with probability at least $8/10$ both Step 5 and Step 4 succeed. We hereafter condition on this.

Completeness. Suppose f is ε_1 -close to k -monotone. Lemma 7.1 and the triangle inequality imply that there is a k -monotone m -block function g^* such that $\text{dist}(f, g^*) \leq \varepsilon_1 + \alpha/6$. The agnostic learning algorithm thus returns a hypothesis h such that $\text{dist}(f, h) \leq \varepsilon_1 + \alpha/4$. The algorithm estimates this closeness to within $\alpha/7$, so the estimate obtained in Step 5 is at most $\varepsilon_1 + \alpha/4 + \alpha/7 < \varepsilon_1 + 5\alpha/12$ and the algorithm does not reject in this step. By the triangle inequality, h is $(2\varepsilon_1 + 5\alpha/12)$ -close to k -monotone, and the algorithm will accept. There is no estimation error here, since no queries to f are required.

Soundness. Now suppose f is ε_2 -far from k -monotone, where $\varepsilon_2 = 3\varepsilon_1 + \alpha$ for some $\alpha > 0$. Suppose the algorithm does not reject when estimating $\text{dist}(f, h)$, where h is the hypothesis returned by the agnostic learning algorithm. Then $\text{dist}(f, h) \leq \varepsilon_1 + 5\alpha/12 + \alpha/7 < \varepsilon_1 + 7\alpha/12$. By the triangle inequality, if t is a k -monotone function, $\text{dist}(h, t) \geq \text{dist}(f, t) - \text{dist}(f, h) > \varepsilon_2 - (\varepsilon_1 + 7\alpha/12) = 2\varepsilon_1 + 5\alpha/12$. The algorithm will thus reject in the final step.

Query complexity. The query complexity of the algorithm is dominated by the query complexity of the agnostic learning algorithm, which is $\exp(\tilde{O}(k\sqrt{d}/\alpha^2)) = \exp(\tilde{O}(k\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2))$. \square

8 Tolerant testing and applications to L_1 -testing

We now show how our techniques can be applied to solve an open problem on L_1 tolerant testing of monotonicity, asked at the Sublinear Algorithms Workshop 2016 [Sub16].

We start by describing a reduction lemma from L_1 distance to monotonicity of functions in $[0, 1]^{\mathcal{X}}$ to Hamming distance to monotonicity of functions in $\{0, 1\}^{\mathcal{X} \times [0, 1]}$ (that is, “trading the range for a dimension”). We note that this idea appears in Berman *et al.* [BRY14, Lemmata 2.1 and 2.3], although formulated in a slightly different way. For convenience and completeness, we state and prove here the version we shall use.

In what follows, we let \mathcal{X} be a discrete partially ordered domain equipped with a measure μ ,¹¹ that is a tuple $(\mathcal{X}, \preceq, \mu)$; and for a set $\mathcal{Y} \subseteq \mathbb{R}$ we denote by $\mathcal{M}^{(\mathcal{X} \rightarrow \mathcal{Y})} \subseteq \mathcal{Y}^{\mathcal{X}}$ the set of monotone functions from \mathcal{X} to \mathcal{Y} .

Definition 8.1 (Analogue of [BRY14, Definition 2.1]). For a function $f: \mathcal{X} \rightarrow [0, 1]$, the *threshold function* $T \circ f: \mathcal{X} \times [0, 1] \rightarrow \{0, 1\}$ is defined by

$$T \circ f(x, t) = \mathbb{1}_{\{f(x) \geq 1-t\}} = \begin{cases} 1 & \text{if } f(x) \geq 1-t \\ 0 & \text{otherwise.} \end{cases}$$

The next fact is immediate from this definition:

¹¹We will only require that (\mathcal{X}, μ) be a measurable space with finite measure, that is $\mu(\mathcal{X}) < \infty$, and shall only hereafter concern ourselves with *measurable* functions.

Fact 8.2. For any $f: \mathcal{X} \rightarrow [0, 1]$, it is the case that for every $x \in \mathcal{X}$

$$f(x) = \int_0^1 T \circ f(x, t) dt.$$

Moreover, $f \in \mathcal{M}(\mathcal{X} \rightarrow [0, 1])$ if, and only if, $T \circ f \in \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\})$.

We begin by the following characterization, which is immediately obtained from a corresponding theorem of Berman *et al.*; before stating a slightly modified version that we shall rely upon. For completeness, we give the proof of the former below.

Proposition 8.3 (Analogue of [BRY14, Lemma 2.1]). For any $f: \mathcal{X} \rightarrow [0, 1]$,

$$L_1(f, \mathcal{M}(\mathcal{X} \rightarrow [0, 1])) = L_1(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\})) = \text{dist}(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\}))$$

(In particular, the RHS is well-defined, i.e. the mapping $t \in [0, 1] \mapsto L_1(T \circ f(\cdot, t), \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\}))$, where $T \circ f(\cdot, t) \in \{0, 1\}^{\mathcal{X}}$ is the function that maps x to $T \circ f(x, t)$, is measurable.)

Proof. We write $\nu \stackrel{\text{def}}{=} \mu \times \text{Leb}_{[0, 1]}$ for the product measure on $\mathcal{X} \times [0, 1]$ induced by μ and the Lebesgue measure on $[0, 1]$; so that $\nu(\mathcal{X} \times [0, 1]) = \mu(\mathcal{X}) \cdot 1 = \mu(\mathcal{X})$.

For any fixed $t \in [0, 1]$, let $g_t \in \mathcal{M}(\mathcal{X} \rightarrow \{0, 1\})$ be any function achieving $L_1(T \circ f(\cdot, t), g_t) = L_1(T \circ f(\cdot, t), \mathcal{M}(\mathcal{X} \rightarrow \{0, 1\}))$, and define $g \in [0, 1]^{\mathcal{X}}$ by $g'(x) = \int_0^1 dt g_t(x)$ for all $x \in \mathcal{X}$: note that g is then monotone by construction.¹² Moreover, choose $h \in \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\})$ as a function achieving $L_1(T \circ f, h) = L_1(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\}))$. Then we have

$$\begin{aligned} L_1(f, \mathcal{M}(\mathcal{X} \rightarrow [0, 1])) &\leq L_1(f, g') = \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) \left| \int_0^1 dt (T \circ f(x, t) - g_t(x)) \right| \\ &\leq \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) \int_0^1 dt |T \circ f(x, t) - g_t(x)| \\ &= \int_0^1 dt \left(\frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) |T \circ f(x, t) - g_t(x)| \right) = \int_0^1 dt L_1(T \circ f(\cdot, t), g_t) \\ &\leq \int_0^1 dt L_1(T \circ f(\cdot, t), h(\cdot, t)) = \int_0^1 dt \left(\frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) |T \circ f(x, t) - h(x, t)| \right) \\ &= \frac{1}{\nu(\mathcal{X} \times [0, 1])} \int_{\mathcal{X} \times [0, 1]} \nu(dx, dt) |T \circ f(x, t) - h(x, t)| \\ &= L_1(T \circ f, h) = L_1(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\})) \end{aligned}$$

where we applied **Fact 8.2** (and the definition of $g' = \int_0^1 g_t$) for the first equality, and for the third inequality the fact that h induces (for every fixed $t \in [0, 1]$) a monotone function $h(\cdot, t) \in \mathcal{M}(\mathcal{X} \rightarrow \{0, 1\})$: so that $L_1(T \circ f(\cdot, t), g_t) \leq L_1(T \circ f(\cdot, t), h(\cdot, t))$ for all t .

¹²Additionally, since we restrict ourselves to finite \mathcal{X} , there are only finitely many distinct functions $T \circ f(\cdot, t)$ (for $t \in [0, 1]$, and therefore only finitely many distinct functions g_t .)

For the other direction of the inequality, fix any $f: \mathcal{X} \rightarrow [0, 1]$, and let $g \in \mathcal{M}(\mathcal{X} \rightarrow [0, 1])$ be (any) function achieving $L_1(f, g) = L_1(f, \mathcal{M}(\mathcal{X} \rightarrow [0, 1]))$. We can write, unrolling the definitions,

$$\begin{aligned}
L_1(f, \mathcal{M}(\mathcal{X} \rightarrow [0, 1])) &= \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) |f(x) - g(x)| \\
&= \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) \left| \int_0^1 dt (T \circ f(x, t) - T \circ g(x, t)) \right| \\
&= \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) \left| \int_0^1 dt (T \circ f(x, t) - T \circ g(x, t)) \right| \\
&= \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) \left(\int_0^1 dt (T \circ f(x, t) - T \circ g(x, t)) \mathbb{1}_{\{f(x) > g(x)\}} \right. \\
&\quad \left. + (T \circ g(x, t) - T \circ f(x, t)) \mathbb{1}_{\{g(x) > f(x)\}} \right) \\
&= \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \int_0^1 dt \mu(dx) \left((T \circ f(x, t) - T \circ g(x, t)) \mathbb{1}_{\{f(x) > g(x)\}} \right. \\
&\quad \left. + (T \circ g(x, t) - T \circ f(x, t)) \mathbb{1}_{\{g(x) > f(x)\}} \right) \\
&= \frac{1}{\nu(\mathcal{X} \times [0, 1])} \int_{\mathcal{X} \times [0, 1]} \nu(dx, dt) |T \circ f(x, t) - T \circ g(x, t)| = L_1(T \circ f, T \circ g) \\
&\geq L_1(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\}))
\end{aligned}$$

where we applied [Fact 8.2](#) for the second equality, the definition of L_1 distance for the second-to-last; and to handle the absolute values we used the fact that $|a - b| = (a - b) \mathbb{1}_{\{a > b\}} + (b - a) \mathbb{1}_{\{a > b\}}$, along with the observation that $T \circ f(x, t) > T \circ g(x, t)$ can only hold if $f(x) > g(x)$. Finally, we have $L_1(T \circ f, T \circ g) \geq L_1(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\}))$ since $T \circ g \in \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\})$, yielding the desired claim.

Finally, the fact that $L_1(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\})) = \text{dist}(T \circ f, \mathcal{M}(\mathcal{X} \times [0, 1] \rightarrow \{0, 1\}))$ is immediate from the Boolean range, as $|a - b| = \mathbb{1}_{\{a \neq b\}}$ for any $a, b \in \{0, 1\}$. \square

Proposition 8.4 (Rounding and Range-Dimension Tradeoff). *For any $f: \mathcal{X} \rightarrow [0, 1]$ and parameter $m \geq 1$, let $R_m \stackrel{\text{def}}{=} \{\frac{1}{m}, \frac{2}{m}, \dots, 1\}$. We define the m -rounding of f as $\Phi_m \circ f: \mathcal{X} \rightarrow R_m$ by*

$$\Phi_m \circ f(x) = \frac{\lceil mf(x) \rceil}{m}, \quad x \in \mathcal{X}.$$

Then we have

- (i) $\left| L_1(f, \mathcal{M}(\mathcal{X} \rightarrow [0, 1])) - L_1(\Phi_m \circ f, \mathcal{M}(\mathcal{X} \rightarrow R_m)) \right| \leq \frac{1}{m}$;
- (ii) $L_1(\Phi_m \circ f, \mathcal{M}(\mathcal{X} \rightarrow R_m)) = \text{dist}(T \circ \Phi_m \circ f, \mathcal{M}(\mathcal{X} \times R_m \rightarrow \{0, 1\}))$.

Proof of Proposition 8.4. Fix any $m \geq 1$. We start the proof of [item \(i\)](#) by the simple observation that if $f \in \mathcal{M}(\mathcal{X} \rightarrow [0, 1])$, then $\Phi_m \circ f \in \mathcal{M}(\mathcal{X} \rightarrow R_m) \subseteq \mathcal{M}(\mathcal{X} \rightarrow [0, 1])$, that is rounding preserves monotonicity; and that $\Phi_m \circ g = g$ for all $g: \mathcal{X} \rightarrow R_m$. This, along with the fact that for all $f: \mathcal{X} \rightarrow [0, 1]$

$$L_1(f, \Phi_m \circ f) = \frac{1}{\mu(\mathcal{X})} \int_{\mathcal{X}} \mu(dx) \underbrace{|\Phi_m \circ f(x) - f(x)|}_{\leq 1/m} \leq \frac{1}{m}$$

implies by the triangle inequality, for any $g \in \mathcal{M}(\mathcal{X} \rightarrow R_m) \subseteq \mathcal{M}(\mathcal{X} \rightarrow [0,1])$, that

$$L_1\left(f, \mathcal{M}(\mathcal{X} \rightarrow [0,1])\right) \leq L_1(f, g) \leq \frac{1}{m} + L_1(g, \Phi_m \circ g) + L_1(\Phi_m \circ f, \Phi_m \circ g) = \frac{1}{m} + 0 + L_1(\Phi_m \circ f, g).$$

Taking $g \in \mathcal{M}(\mathcal{X} \rightarrow R_m)$ that achieves $L_1(\Phi_m \circ f, g) = L_1(\Phi_m \circ f, \mathcal{M}(\mathcal{X} \rightarrow R_m))$, we get

$$L_1\left(f, \mathcal{M}(\mathcal{X} \rightarrow [0,1])\right) \leq \frac{1}{m} + L_1(\Phi_m \circ f, \mathcal{M}(\mathcal{X} \rightarrow R_m)).$$

For the other direction, we first note that for any two functions $f, g: \mathcal{X} \rightarrow [0, 1]$, it is the case that $L_1(f, g) \geq L_1(\Phi_m \circ f, \Phi_m \circ g) - 1/m$ (which is immediate from the definition of the rounding operator), and taking g to be the closest monotone function to f this readily yields

$$L_1\left(f, \mathcal{M}(\mathcal{X} \rightarrow [0,1])\right) \geq L_1(\Phi_m \circ f, \Phi_m \circ g) - \frac{1}{m} \geq L_1(\Phi_m \circ f, \mathcal{M}(\mathcal{X} \rightarrow R_m)) - \frac{1}{m}.$$

Finally, the proof of the second part, **item (ii)**, is identical to that of **Proposition 8.3**, replacing the Lebesgue measure on $[0, 1]$ by the counting measure on R_m . (So that integrals over $[0, 1]$ become sums over R_m , normalized by $|R_m| = m$.) \square

Given **Proposition 8.4**, it is now easy to apply the results of **Section 7** to obtain a tolerant L_1 tester for monotonicity of functions $f: [n]^d \rightarrow [0, 1]$. Indeed, given parameters $0 < \varepsilon_1 < \varepsilon_2$, one can set the rounding parameter m to $\lceil 4/(\varepsilon_2 - \varepsilon_1) \rceil$; and from query access to $f: [n]^d \rightarrow [0, 1]$, simulate query access to $\Phi_m \circ f$ and therefore to $g \stackrel{\text{def}}{=} T \circ \Phi_m \circ f: [n]^d \times R_m \rightarrow \{0, 1\}$. By **Proposition 8.4** and our choice of m , in order to distinguish

$$L_1\left(f, \mathcal{M}(\mathcal{X} \rightarrow [0,1])\right) \leq \varepsilon_1 \quad \text{vs.} \quad L_1\left(f, \mathcal{M}(\mathcal{X} \rightarrow [0,1])\right) \geq \varepsilon_2$$

it is enough to distinguish

$$\text{dist}\left(g, \mathcal{M}(\mathcal{X} \times R_m \rightarrow \{0,1\})\right) \leq \varepsilon_1 + \frac{1}{m} \quad \text{vs.} \quad \text{dist}\left(g, \mathcal{M}(\mathcal{X} \times R_m \rightarrow \{0,1\})\right) \geq \varepsilon_2 - \frac{1}{m}.$$

By our choice of m , we also have $(\varepsilon_2 - \frac{1}{m}) - (\varepsilon_1 + \frac{1}{m}) \geq \frac{\varepsilon_2 - \varepsilon_1}{2}$.

The last step is to observe that one can view equivalently g as a function $g: [n]^d \times [m] \rightarrow \{0, 1\}$; by **Proposition 8.4** and our choice of m , so that the algorithms of **Theorem 1.7** apply.

Theorem 1.9. *There exists a non-adaptive tolerant L_1 -tester for monotonicity of functions $f: [n]^d \rightarrow [0, 1]$ with query complexity*

- $\tilde{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \left(\frac{5d}{\varepsilon_2 - \varepsilon_1}\right)^d\right)$, for any $0 \leq \varepsilon_1 < \varepsilon_2 \leq 1$;
- $2^{\tilde{O}(\sqrt{d}/(\varepsilon_2 - 3\varepsilon_1)^2)}$, for any $0 \leq 3\varepsilon_1 < \varepsilon_2 \leq 1$.

Acknowledgments. We would like to thank Eric Blais for helpful remarks on an earlier version of this paper, and an anonymous reviewer for very detailed and insightful comments.

References

- [AC06] Nir Ailon and Bernard Chazelle. Information theory in property testing and monotonicity testing in higher dimension. *Inf. Comput.*, 204(11):1704–1717, 2006. [1.1.2](#)
- [ACCL07] Nir Ailon, Bernard Chazelle, Seshadhri Comandur, and Ding Liu. Estimating the distance to a monotone function. *Random Struct. Algorithms*, 31(3):371–383, 2007. [3](#)
- [AM05] Kazuyuki Amano and Akira Maruoka. A superpolynomial lower bound for a circuit computing the Clique function with at most $(1/6) \log \log n$ negation gates. *SIAM Journal on Computing*, 35(1):201–216, 2005. [1.4](#)
- [Ang87] Dana Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1987. [1](#)
- [BB16] Aleksandrs Belovs and Eric Blais. A polynomial lower bound for testing monotonicity. In *STOC*, pages 1021–1032. ACM, 2016. [1](#), [1.1.1](#), [4.2](#), [4.2.2](#), [A](#)
- [BBBY12] Maria-Florina Balcan, Eric Blais, Avrim Blum, and Liu Yang. Active property testing. In *FOCS*, pages 21–30. IEEE Computer Society, 2012. [1.1.2](#)
- [BBM12] Eric Blais, Joshua Brody, and Kevin Matulef. Property testing lower bounds via communication complexity. *Computational Complexity*, 21(2):311–358, 2012. [1.1.2](#)
- [BCGM12] Jop Briët, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Monotonicity testing and shortest-path routing on the cube. *Combinatorica*, 32(1):35–53, 2012. [1](#), [1.1.1](#)
- [BCO⁺15] Eric Blais, Clément L. Canonne, Igor Carboni Oliveira, Rocco A. Servedio, and Li-Yang Tan. Learning circuits with few negations. In *APPROX-RANDOM*, volume 40 of *LIPICs*, pages 512–527. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. [1](#), [1.1.1](#), [1.2](#), [1.4](#), [7.2](#), [7.5](#)
- [BGJ⁺09] Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, and David P. Woodruff. Transitive-closure spanners. In *SODA*, pages 932–941. SIAM, 2009. [A](#)
- [BRW05] Tüçkan Batu, Ronitt Rubinfeld, and Patrick White. Fast approximate PCPs for multidimensional bin-packing problems. *Inf. Comput.*, 196(1):42–56, 2005. [1.1.2](#)
- [BRY14] Piotr Berman, Sofya Raskhodnikova, and Grigory Yaroslavtsev. L_p -testing. In *STOC*, pages 164–173. ACM, 2014. [1.1.2](#), [1.1.2](#), [1.1.2](#), [1.3](#), [2](#), [6.1](#), [1](#), [6.1](#), [6.2](#), [8](#), [8.1](#), [8.3](#), [8](#), [A](#), [A](#)
- [BT96] Nader H. Bshouty and Christino Tamon. On the Fourier spectrum of monotone functions. *J. ACM*, 43(4):747–770, 1996. [1](#)
- [CDST15] Xi Chen, Anindya De, Rocco A. Servedio, and Li-Yang Tan. Boolean function monotonicity testing requires (almost) $n^{1/2}$ non-adaptive queries. In *STOC*, pages 519–528. ACM, 2015. [1](#), [1.1.1](#), [4.2](#), [4.2.2](#), [A](#)
- [CR14] Clément L. Canonne and Ronitt Rubinfeld. Testing probability distributions underlying aggregated data. In *ICALP (1)*, volume 8572 of *Lecture Notes in Computer Science*, pages 283–295. Springer, 2014. [1.2](#), [5.2.1](#), [5.2.1](#), [5.4](#), [5.2.1](#), [5.2.1](#)

- [CS13a] Deeparnab Chakrabarty and C. Seshadhri. An $o(n)$ monotonicity tester for boolean functions over the hypercube. In *STOC*, pages 411–418. ACM, 2013. Journal version as [CS16]. 1, 1.1.1, 1.1.2, 1.3, A
- [CS13b] Deeparnab Chakrabarty and C. Seshadhri. Optimal bounds for monotonicity and Lipschitz testing over hypercubes and hypergrids. In *STOC*, pages 419–428, 2013. 1, 1.1.2, A
- [CS14] Deeparnab Chakrabarty and C. Seshadhri. An optimal lower bound for monotonicity testing over hypergrids. *Theory of Computing*, 10:453–464, 2014. 1, 1.1.2, A
- [CS16] Deeparnab Chakrabarty and C. Seshadhri. An $o(n)$ Monotonicity Tester for Boolean Functions over the Hypercube. *SIAM J. Comput.*, 45(2):461–472, 2016. 8
- [CST14] Xi Chen, Rocco A. Servedio, and Li-Yang Tan. New algorithms and lower bounds for monotonicity testing. In *FOCS*, pages 286–295. IEEE Computer Society, 2014. 1.1.1, 1.3, A
- [DGL⁺99] Yevgeniy Dodis, Oded Goldreich, Eric Lehman, Sofya Raskhodnikova, Dana Ron, and Alex Samorodnitsky. Improved testing algorithms for monotonicity. In *RANDOM-APPROX*, volume 1671 of *Lecture Notes in Computer Science*, pages 97–108. Springer, 1999. 1, 1.1.1, 1.2
- [EKK⁺00] Funda Ergün, Sampath Kannan, Ravi Kumar, Ronitt Rubinfeld, and Mahesh Viswanathan. Spot-checkers. *J. Comput. Syst. Sci.*, 60(3):717–751, 2000. 1.1.2, 3
- [Fis04] Eldar Fischer. On the strength of comparisons in property testing. *Inf. Comput.*, 189(1):107–116, 2004. 1.1.2
- [FLN⁺02] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 474–483, 2002. 1, 1.1.1, 1.1.1, 3, 3, 3, 3, 4.2, 4.2.1, A, A
- [FR10] Shahar Fattal and Dana Ron. Approximating the distance to monotonicity in high dimensions. *ACM Trans. Algorithms*, 6(3), 2010. 1.1.2, 1.2, A
- [GGL⁺00] Oded Goldreich, Shafi Goldwasser, Eric Lehman, Dana Ron, and Alex Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000. 1, 1.1.1, 1.1.2, 1.3, 4.2.2, A
- [GK15] Siyao Guo and Ilan Komargodski. Negation-limited formulas. In *APPROX-RANDOM*, volume 40 of *LIPICs*, pages 850–866. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. 1, 1.4
- [GMOR15] Siyao Guo, Tal Malkin, Igor Carboni Oliveira, and Alon Rosen. The power of negations in cryptography. In *TCC (1)*, volume 9014 of *Lecture Notes in Computer Science*, pages 36–65. Springer, 2015. 1, 1.4
- [HK08] Shirley Halevy and Eyal Kushilevitz. Testing monotonicity over graph products. *Random Struct. Algorithms*, 33(1):44–67, 2008. 1.1.2, 3

- [Juk12] Stasys Jukna. *Boolean Function Complexity*. Springer, 2012. 1, 1.4
- [KKMS08] Adam Tauman Kalai, Adam R. Klivans, Yishay Mansour, and Rocco A. Servedio. Agnostically learning halfspaces. *SIAM J. Comput.*, 37(6):1777–1805, 2008. 1.2, 7.2, 7.10, 7.4
- [KMS15] Subhash Khot, Dor Minzer, and Muli Safra. On monotonicity testing and Boolean isoperimetric type theorems. In *FOCS*, pages 52–58. IEEE Computer Society, 2015. 1, 1.1.1, 1.1.1, 1.1.1, 1.3, 4.2, A
- [KNOW14] Pravesh Kothari, Amir Nayyeri, Ryan O’Donnell, and Chenggang Wu. Testing surface area. In *SODA*, pages 1204–1214. SIAM, 2014. 1.1.2
- [KR00] Michael J. Kearns and Dana Ron. Testing problems with sublearning sample complexity. *J. Comput. Syst. Sci.*, 61(3):428–456, 2000. 1.1.2
- [KV94] Michael J. Kearns and Leslie G. Valiant. Cryptographic limitations on learning Boolean formulae and finite automata. *J. ACM*, 41(1):67–95, 1994. 1
- [LZ16] Chengyu Lin and Shengyu Zhang. Sensitivity conjecture and log-rank conjecture for functions with small alternating numbers. *CoRR*, abs/1602.06627, 2016. 1, 1.4
- [Mar57] A. A. Markov. On the inversion complexity of systems of functions. *Doklady Akademii Nauk SSSR*, 116:917–919, 1957. English translation in [Mar58]. 1
- [Mar58] A. A. Markov. On the inversion complexity of a system of functions. *Journal of the ACM*, 5(4):331–334, October 1958. 8
- [Nee14] Joe Neeman. Testing surface area with arbitrary accuracy. In *STOC*, pages 393–397. ACM, 2014. 1.1.2
- [O’D14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. 7.7
- [OS07] Ryan O’Donnell and Rocco A. Servedio. Learning monotone decision trees in polynomial time. *SIAM J. Comput.*, 37(3):827–844, 2007. 1
- [OW09] Ryan O’Donnell and Karl Wimmer. KKL, Kruskal–Katona, and monotone nets. In *FOCS*, pages 725–734. IEEE Computer Society, 2009. 1
- [PRR06] Michal Parnas, Dana Ron, and Ronitt Rubinfeld. Tolerant property testing and distance approximation. *Journal of Computer and System Sciences*, 72(6):1012–1042, 2006. 1.1.2, 3
- [Raz85] Alexander A Razborov. Lower bounds on the monotone complexity of some Boolean functions. In *Doklady Akademii Nauk SSSR*, volume 281, pages 798–801, 1985. 1
- [Ros15] Benjamin Rossman. Correlation bounds against monotone NC^1 . In *Conference on Computational Complexity (CCC)*, 2015. 1, 1.4
- [RW92] Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992. 1

- [Ser04] Rocco A. Servedio. On learning monotone DNF under product distributions. *Inf. Comput.*, 193(1):57–74, 2004. [1](#)
- [Sub16] List of open problems in sublinear algorithms: Problem 70. <http://sublinear.info/70>, 2016. Originally posed in [BRY14]. [1.1.2](#), [1.3](#), [8](#)
- [Val84] Leslie G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984. [2](#)

A Previous work on monotonicity testing

In this appendix, we summarize the state-of-the-art on monotonicity testing. We observe that this question has been considered for functions over various domains (e.g. hypergrids, hypercubes and general posets) and ranges (notably Boolean range $\{0, 1\}$ and unbounded range \mathbb{N}); as hypergrids and hypercubes are arguably the domains that have received the most attention in the literature, we will in this overview restrict ourselves on work on these, and refer readers to [FLN⁺02, BGJ⁺09] for other various posets. We will also focus on the Boolean range $\{0, 1\}$, which is most relevant to our work. We also briefly mention the best known results (which are also tight) for unbounded range \mathbb{N} . At the end of this section, we include known results for *tolerant* monotonicity testing.

Before we go over those results, we recall some notation: namely, testers can make adaptive (a.) or non-adaptive (n.a.) queries and have 1-sided (1.s.) or 2-sided (2.s.) error. The best one could hope for would then be to obtain 1-sided non-adaptive upper bound, complemented with 2-sided adaptive lower bounds. We note that all testers included in below except tolerant testers are 1-sided (almost all of them are non-adaptive) algorithms.

Hypercubes with Boolean Range. The problem of monotonicity testing is introduced by Goldreich *et al.* [GGL⁺00] for functions $f: \{0, 1\}^d \rightarrow \{0, 1\}$. [GGL⁺00] present a simple “edge tester” with query complexity $O(d/\varepsilon)$. A tester with $O(d^{7/8}/\varepsilon^{3/2})$ queries, the first improvement in terms of the dependence on d and the first to “break” the linear barrier, was presented by Chakrabarty and Seshadhri [CS13a], further improved to $\tilde{O}(d^{5/6}/\varepsilon^4)$ by Chen *et al.* [CST14]. Recently, a $\tilde{O}(\sqrt{d}/\varepsilon^2)$ upper bound was established by Khot *et al.* [KMS15]. All these upper bounds are obtained for 1-sided, non-adaptive testers.

For 1-sided non-adaptive testers, Fischer *et al.* [FLN⁺02] showed an $\Omega(\sqrt{d})$ lower bound. For 2-sided non-adaptive testers, Chen *et al.* [CST14] obtained an $\tilde{\Omega}(d^{1/5})$ lower bound, further improved by Chen *et al.* [CDST15] to $\Omega(d^{1/2-c})$ (for any constant $c > 0$). All these lower bounds applying to non-adaptive testers, they only imply an $\Omega(\log d)$ lower bound for adaptive ones. Recently, Belovs and Blais [BB16] showed an $\Omega(d^{1/4})$ lower bound for 2-sided adaptive testers, i.e. an exponential improvement over the previous bounds. All mentioned lower bounds hold for constant $\varepsilon > 0$, and are summarized in Table 3.

Domain	Upper bound	Lower bound
$\{0, 1\}^d$	$O(\sqrt{d})$ 1.s.-n.a. [KMS15]	$\Omega(d^{1/2})$ 1.s.-n.a. [FLN ⁺ 02] $\Omega(d^{1/2-o(1)})$ 2.s.-n.a. [CDST15] $\tilde{\Omega}(d^{1/4})$ 2.s.-a. [BB16]

Table 3: Testing monotonicity of a function $f: \{0, 1\}^d \rightarrow \{0, 1\}$

Hypergrids with Boolean Range. We remark that most known previous upper bounds for testing monotonicity over hypergrids are for unbounded range, which we will be the focus of the next section. Instead, we only mention here the case of Boolean range, giving in each setting the current best known results. For testing monotonicity over the line with Boolean range (i.e. $d = 1$ case), both a 1-sided non-adaptive $O(1/\varepsilon)$ upper bound and a 2-sided adaptive $\Omega(1/\varepsilon)$ lower bound are known (both of them being folklore). For $d = 2$, Berman *et al.* [BRY14] showed a tight bound of $\Theta((\log 1/\varepsilon)/\varepsilon)$ for 1-sided non-adaptive testers. Interestingly, they also prove that “adaptivity”

helps in the $d = 2$ case: that is, they establish a 1-sided tight adaptive $O(1/\varepsilon)$ upper bound which beats $\Omega(\log 1/\varepsilon)/\varepsilon$ lower bound for 1-sided non-adaptive testers. For general d , Berman *et al.* [BRY14] give both a 1-sided non-adaptive tester with query complexity $O(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon})$, and a 1-sided adaptive tester with query complexity $O\left(d2^d \log^{d-1} \frac{1}{\varepsilon} + \frac{d^2 \log d}{\varepsilon}\right)$. The best known results can be found in Table 4.

Domain	Upper bound	Lower bound
$[n]$	$O\left(\frac{1}{\varepsilon}\right)$ 1.s.-n.a.	$\Omega\left(\frac{1}{\varepsilon}\right)$ 2.s.-a.
$[n]^2$	$O\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ 1.s.-n.a. [BRY14] $O\left(\frac{1}{\varepsilon}\right)$ 1.s.-a. [BRY14]	$\Omega\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ 1.s.-n.a. [BRY14] $\Omega\left(\frac{1}{\varepsilon}\right)$ 2.s.-a.
$[n]^d$	$O\left(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon}\right)$ 1.s.-n.a. [BRY14] $O\left(d2^d \log^{d-1} \frac{1}{\varepsilon} + \frac{d^2 \log d}{\varepsilon}\right)$ 1.s.-a. [BRY14]	$\Omega\left(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ 1.s.-n.a. [BRY14] $\Omega\left(\frac{1}{\varepsilon}\right)$ 2.s.-a.

Table 4: Testing monotonicity of a function $f: [n]^d \rightarrow \{0, 1\}$

Unbounded Range. For unbounded range, tight upper and lower bounds are known for both hypergrid and hypercube domains. Chakrabarty and Seshadhri [CS14] describe a 1-sided non-adaptive tester with $O(d \log n/\varepsilon)$ queries for the hypergrid $[n]^d$. Later, they show that $O(d \log n/\varepsilon)$ is essentially optimal even for 2-sided adaptive tester [CS13b]. For the hypercube, [CS14] give a 1-sided non-adaptive tester making $O(n/\varepsilon)$ queries, and a matching 2-sided adaptive lower bound is proved by Joshua Brody (mentioned as private communication in [CS13b]). We refer readers to [CS14, CS13b] for overviews on previous results for testing monotonicity over the hypercube and hypergrid with unbounded range. The best known results are summarized in Table 5.

Domain	Upper bound	Lower bound
$\{0, 1\}^d$	$O\left(\frac{d}{\varepsilon}\right)$ 1.s.-n.a. [CS14]	$\Omega\left(\frac{d}{\varepsilon}\right)$ 2.s.-a. [CS13b]
$[n]^d$	$O\left(\frac{d \log n}{\varepsilon}\right)$ 1.s.-n.a. [CS14]	$\Omega\left(\frac{d \log n}{\varepsilon} - \frac{1}{\varepsilon} \log \frac{1}{\varepsilon}\right)$ 2.s.-a. [CS13b]

Table 5: Testing monotonicity of a function $f: D \rightarrow \mathbb{N}$

Tolerant Testing. To the best of our knowledge, prior to our work tolerant testers for monotonicity for Boolean functions over the hypergrid were only known for dimension $d \in \{1, 2\}$. Specifically, an $O\left(\frac{\varepsilon_2}{(\varepsilon_2 - \varepsilon_1)^2}\right)$ -query upper bound is known for $d = 1$, while an $\tilde{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^4}\right)$ -query one is known for $d = 2$ [BRY14, FR10].

Domain	Upper bound	Lower bound
$[n]$	$O\left(\frac{\varepsilon_2}{(\varepsilon_2 - \varepsilon_1)^2}\right)$ [BRY14] 2.s.-a.	?
$[n]^2$	$\tilde{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^4}\right)$ [FR10] 2.s.-n.a.	?

Table 6: Tolerant testing monotonicity of a function $f: [n]^d \rightarrow \{0, 1\}$