

A Noisy-Influence Regularity Lemma for Boolean Functions

Chris Jones

Abstract

We present a regularity lemma for Boolean functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ based on noisy influence, a measure of how locally correlated f is with each input bit. We provide an application of the regularity lemma to weaken the conditions on the Majority is Stablest Theorem. We also prove a “homogenized” version stating that there is a small set of input bits so that most restrictions of f on those bits have small noisy influences. These results were sketched out by [OSTW10], but never published. With their permission, we present the full details here.

1 Introduction

A recent theme in discrete mathematics has been the development of regularity lemmas, tools which break down large-scale combinatorial objects into a constant number of easy-to-understand pieces. The goal of this note is to prove such a regularity lemma for Boolean functions $f : \{-1, 1\}^n \rightarrow \mathbb{R}$.

There are two items to address. First, how does one “break down” a Boolean function? We use a decision tree which queries individual bits at internal nodes and places a subfunction at each leaf. Furthermore, the depth of this decision tree will be independent of n . Second, what does it mean for a subfunction to be “easy-to-understand”? One notion of “easy-to-understand” that arises in other contexts is “pseudorandom”: possessing structure that is likely to arise if the object is chosen randomly. It is likely that a randomly-chosen Boolean function will not locally behave like any single input bit, and this is the notion of “easy-to-understand” we adopt here. We will precisely define these concepts in Section 2, where we prove the main statement of the regularity lemma:

Theorem 1. *For every $\delta, \gamma \in (0, 1], \epsilon > 0$ and $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that $\mathbb{E}[f^2] \leq 1$, there is a decision tree \mathcal{D} of depth at most $1/\epsilon\delta\gamma$ and functions $f_L : \{-1, 1\}^n \rightarrow \{-1, 1\}$ indexed by leaves L of \mathcal{D} such that*

(i) $f(x) = f_{\mathcal{D}(x)}(x)$

(ii) *All but at most a γ fraction of the f_L have (ϵ, δ) -small noisy influences.*

We also prove a “homogenized” version of the theorem. In this version, the decision tree must query the same bit on every level. The subfunctions of this tree correspond to the restrictions of f on the queried inputs. Thus the theorem states that there are a constant number of bits so that most restrictions of f have small influences:

Theorem 2. *For every $\delta, \gamma \in (0, 1], \epsilon > 0, 1/\epsilon\delta\gamma > 1$, and $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that $\mathbb{E}[f^2] \leq 1$, there is $J \subseteq [n]$ of size at most $2 \uparrow 1/\epsilon\delta\gamma$ such that all but at most a γ fraction of restrictions of f on J have (ϵ, δ) -small noisy influences.*

The class of low-influence functions is one way of representing real-world functions where each input contributes a small piece of the output. On the mathematical side the study of low noisy-influence functions has been driven by the development of invariance principles and connections between functions on the product probability space $\{-1, 1\}^n$ and Gaussian space $\mathcal{N}(0, 1)^m$ [MOO05]. As an application of the regularity lemma to this study, we slightly weaken the conditions on the Majority is Stablest Theorem proven in [MOO05]. The proof of Theorem 1 and the application to the Majority is Stablest Theorem are based on sketches from [OSTW10]. The complete proofs here are presented with the permission of the authors.

1.1 Previous Regularity Lemmas

Regularity lemmas and decomposition results among different classes of Boolean functions are not new [DSTW09], [Gre], [TTV09]. Work of Ben Green from 2004 [Gre] established a regularity lemma for general abelian groups, which specializes to the Boolean case:

Theorem 3. Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that $\mathbb{E}[f^2] \leq 1$, and let $\gamma \in (0, 1], \epsilon > 0$. Then there is a “generalized decision tree” \mathcal{D} of height at most $1/\gamma\epsilon^2$ and a subfunction for each leaf $f_L : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that

(i) $f(x) = f_{\mathcal{D}(x)}(x)$

(ii) All but at most a γ fraction of leaves L are ϵ -regular: $|\widehat{f}_L(S)| \leq \epsilon$ for every $S \neq \emptyset$.

A “generalized decision tree” is permitted to split on parities of arbitrary subsets of bits, rather than parity of a single bit. In the original paper, Green furthermore obtained a generalized decision tree in which all nodes at the same level query the parity of the same set of bits. One can think of this as a sort of “homogenized” version of the above theorem.

One interpretation of such a “homogenized” tree is that each leaf restricts the input to a different coset of a subspace (this subspace is the annihilator of all parity functions, with codimension the height of the tree), hence the following “arithmetic regularity lemma”:

Theorem 4. Let $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$ such that $\mathbb{E}[f^2] \leq 1$, and let $\epsilon > 0$. Then there is a subspace H of codimension at most $2 \uparrow \uparrow 1/\gamma\epsilon^2$ such that at least $1 - \gamma$ fraction of cosets L of H ensure that f is ϵ -regular on L .

Note the superexponential bound on the depth. Tower-type lower bounds have also been shown for Theorem 4 by Hosseini et al in [HLMS14].

A key philosophical difference between Green’s work and ours is that special preference is given here to individual bits, and none to strings with Hamming weight greater than 1. In comparison, the arithmetic regularity lemma considers all nonzero elements as interchangeable vectors from \mathbb{F}_2^n .

2 Proof of Regularity Lemma

This section is concerned with proving Theorems 1 and 2. First we establish some definitions. The terminology used is in alignment with [O’D14].

2.1 Definitions

Consider a Boolean function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$.

Definition 2.1. The noise stability of f at $\rho \in [0, 1]$, written $\mathbf{Stab}_\rho[f]$, is

$$\mathbf{Stab}_\rho[f] = \mathbb{E}_{\substack{(x,y) \\ \rho\text{-correlated}}} [f(x)f(y)]$$

where ρ -correlated strings (x, y) are formed by picking x uniformly from $\{-1, 1\}^n$, and y by taking each bit y_i to have correlation ρ with x_i .

There is a Fourier formula for noise stability:

$$\mathbf{Stab}_\rho[f] = \sum_{S \subseteq [n]} \rho^{|S|} \widehat{f}(S)^2$$

From this we see, when $\rho \geq 0$, $\mathbf{Stab}_\rho[f] \geq 0$, and since $\rho \leq 1$, $\mathbf{Stab}_\rho[f] \leq \mathbb{E}[f^2]$.

Definition 2.2. Define the i^{th} directional derivative operator D_i by

$$(D_i f)(x) = \frac{f(x^{(i \rightarrow 1)}) - f(x^{(i \rightarrow -1)})}{2}$$

Definition 2.3. For $\delta \in [0, 1]$, the $(1 - \delta)$ -noisy influence of x_i on f , written $\mathbf{Inf}_i^{(1-\delta)}[f]$, is

$$\mathbf{Inf}_i^{(1-\delta)}[f] = \mathbf{Stab}_{1-\delta}[D_i f]$$

$\mathbf{Stab}_{1-\delta}[f]$ is a measure of how locally constant f is, and thus $\mathbf{Inf}_i^{(1-\delta)}[f]$ measures how much f is locally correlated with the i -th input bit. Finally, low-influence functions are those that have all local correlations small:

Definition 2.4. We say that f has (ϵ, δ) -small noisy influences if $\mathbf{Inf}_i^{(1-\delta)}[f] \leq \epsilon$ for every i .

A decision tree is a particular representation of a Boolean function that computes input x by querying for a particular input bit of x , and then proceeding to the left or right child depending on that value. In our decision trees, a leaf can contain a subfunction which will be evaluated on any inputs that evaluate to the leaf. For example, this decision tree computes $(x \text{ OR } y) \text{ AND } z$ (thinking of -1 as true, $+1$ as false):




Figure 1: An example decision tree.

Here and later we write $\mathcal{D}(x)$ for the leaf output by a decision tree.

2.2 Proof of Theorems

For clarity we restate the theorem to prove here.

Theorem 1. For every $\delta, \gamma \in (0, 1], \epsilon > 0$ and $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that $\mathbb{E}[f^2] \leq 1$, there is a decision tree \mathcal{D} of depth at most $1/\epsilon\delta\gamma$ and functions $f_L : \{-1, 1\}^n \rightarrow \{-1, 1\}$ indexed by leaves L of \mathcal{D} such that

(i) $f(x) = f_{\mathcal{D}(x)}(x)$

(ii) All but at most a γ fraction of the f_L have (ϵ, δ) -small noisy influences.

When we say “ γ fraction of leaves”, we mean that making random decisions from the root of the tree leads to a leaf where the desired property holds with probability at least $1 - \gamma$.

The proof is constructive, and follows the energy increment technique used to prove other regularity lemmas. That is, if we have a decision tree \mathcal{D} that computes f , we define the energy $\varphi(\mathcal{D})$ by

$$\varphi(\mathcal{D}) = \mathbb{E}_L[\mathbf{Stab}_{1-\delta}[f_L]]$$

We will show $0 \leq \varphi(\mathcal{D}) \leq 1$. If we at any time violate the goals of Theorem 1 by having too many leaves with large noisy influences, we split each leaf on the bit with large influence to “stabilize” those leaves. The proof strategy is to show that this increases the energy of \mathcal{D} by a constant amount, and hence we won’t have to repeat this splitting operation too many times.

The following equality will be used to show the energy change as we split on a bit with high influence, and hence is the crux of the argument:

Lemma 2.1. Fix $i \in [n]$. Define $f_{-1}(x) := f(x^{i \rightarrow -1})$ and $f_1(x) := f(x^{i \rightarrow 1})$. Then

$$\frac{1}{2} \mathbf{Stab}_{1-\delta}[f_{-1}] + \frac{1}{2} \mathbf{Stab}_{1-\delta}[f_1] = \mathbf{Stab}_{1-\delta}[f] + \delta \mathbf{Inf}_i^{(1-\delta)}[f]$$

Proof. The left hand side is $\mathbb{E}_{\substack{(x,y) \\ (1-\delta)\text{-correlated}}} [f(x)f(y) \mid x_i = y_i]$. On the other hand, we can generate the same distribution by picking $(1 - \delta)$ -correlated strings (x, y) and computing $X + Y$, where $X = f(x)f(y)$ and Y is a correction factor that takes value $f(x)f(y^{i \rightarrow x_i}) - f(x)f(y)$. Note that Y is only nonzero when $x_i \neq y_i$. Taking expectations and applying linearity,

$$\begin{aligned} \mathbb{E}[X] &= \mathbf{Stab}_{1-\delta}[f] \\ \mathbb{E}[Y] &= \Pr[x_i \neq y_i] \left(\frac{1}{2} \mathbb{E}[Y \mid x_i = -1, x_i \neq y_i] + \frac{1}{2} \mathbb{E}[Y \mid x_i = 1, x_i \neq y_i] \right) \\ &= \frac{\delta}{2} \left(\frac{1}{2} \mathbb{E}[f(x^{i \rightarrow -1})f(y^{i \rightarrow -1}) - f(x^{i \rightarrow -1})f(y^{i \rightarrow 1})] \right. \\ &\quad \left. + \frac{1}{2} \mathbb{E}[f(x^{i \rightarrow 1})f(y^{i \rightarrow 1}) - f(x^{i \rightarrow 1})f(y^{i \rightarrow -1})] \right) \\ &= \delta \mathbb{E} \left[\left(\frac{f(x^{i \rightarrow 1}) - f(x^{i \rightarrow -1})}{2} \right) \left(\frac{f(y^{i \rightarrow 1}) - f(y^{i \rightarrow -1})}{2} \right) \right] \\ &= \delta \mathbb{E}[D_i f(x) D_i f(y)] = \delta \mathbf{Inf}_i^{(1-\delta)}[f] \end{aligned}$$

□

Now we’re ready to fill in the details of Theorem 1.

Proof of Theorem 1. We construct a decision tree \mathcal{D} with the desired properties. Start \mathcal{D} out as a single-leaf decision tree with f itself at the leaf. We perform the following iterative splitting process on our decision tree \mathcal{D} : suppose we have a leaf L and a bit x_j such that $\mathbf{Inf}_j^{(1-\delta)}[f_L] > \epsilon$. Form decision tree \mathcal{D}' by replacing L with a query to x_j , and subfunctions $f_{L,-1}, f_{L,1}$ defined by $f_{L,-1}(x) = f_L(x^{j \rightarrow -1})$ and $f_{L,1}(x) = f_L(x^{j \rightarrow 1})$. How does the energy change? We claim it increases by at least $\epsilon \delta$ just on this leaf:

$$\begin{aligned} \varphi(\mathcal{D}') - \varphi(\mathcal{D}) &= \mathbb{E}_{L \sim \mathcal{D}'}[\mathbf{Stab}_{1-\delta}[f_L]] - \mathbb{E}_{L \sim \mathcal{D}}[\mathbf{Stab}_{1-\delta}[f_L]] \\ &= \Pr[\text{select leaf } L] \left(\frac{1}{2} \mathbf{Stab}_{1-\delta}[f_{L,-1}] + \frac{1}{2} \mathbf{Stab}_{1-\delta}[f_{L,1}] - \mathbf{Stab}_{1-\delta}[f_L] \right) \end{aligned}$$

By Lemma 2.1 the latter quantity is

$$\Pr[\text{select leaf } L] \delta \mathbf{Inf}_j^{(1-\delta)}[f_L] > \Pr[\text{select leaf } L] \delta \epsilon$$

If there are at most γ fraction of leaves that don't have (ϵ, δ) -small noisy influences, we are done. If not, performing the above replacement on each leaf, we replace \mathcal{D} with our new decision tree \mathcal{D}' such that $\varphi(\mathcal{D}') \geq \varphi(\mathcal{D}) + \epsilon\delta\gamma$ and \mathcal{D}' has depth at most one greater than that of \mathcal{D} . Our next goal is to show termination by showing the energy is bounded.

Recall the definition of $\varphi(\mathcal{D})$,

$$\varphi(\mathcal{D}) = \mathbb{E}_L[\mathbf{Stab}_{1-\delta}[f_L]]$$

We have $\mathbf{Stab}_{1-\delta}[f] \geq 0$ (see Definition 2.1), so averaging maintains $\varphi(\mathcal{D}) \geq 0$. We will show that $\varphi(\mathcal{D}) \leq 1$ to provide an upper bound on the energy.

Even if we hadn't chosen a bit with $\mathbf{Inf}_j^{(1-\delta)}[f_L] > \epsilon$, since $\mathbf{Inf}_j^{(1-\delta)}[f] \geq 0$, we still know that the energy does not decrease if we split on any leaf. During the iteration, a subfunction f_L at depth k fixes k bits of f , no two subfunctions fix the same bits to the same values (they differ at their least common ancestor), and no root-to-leaf path splits on the same variable twice (once a variable x_i has been split on, any subfunctions f_L in that subtree are constant with respect to x_i , and hence have $\mathbf{Inf}_i^{(1-\delta)}[f_L] = 0$). From these three properties, we can extend our tree \mathcal{D} via this splitting operation to a complete binary tree \mathcal{T} of depth n , where each subfunction of \mathcal{T} is constant and the subfunctions take on values $f(x)$ for each $x \in \{-1, 1\}^n$. Nondecreasing energy upon splitting implies we can bound the energy of \mathcal{D} by the energy of \mathcal{T} , which is

$$\varphi(\mathcal{T}) = \mathbb{E}_L[\mathbf{Stab}_{1-\delta}[f_L]] = \mathbb{E}_{x \sim \{-1, 1\}^n}[\mathbf{Stab}_{1-\delta}[f(x)]] = \mathbb{E}_{x \sim \{-1, 1\}^n}[f(x)^2] \leq 1$$

Since $\varphi(\mathcal{D}) \leq 1$ at all times yet increases by $\epsilon\delta\gamma$, we iterate at most $1/\epsilon\delta\gamma$ times. This yields a bound on the depth of \mathcal{D} of at most $1/\epsilon\delta\gamma$. The final decision tree \mathcal{D} computes f , and has no leaf with large noisy influences. \square

We can further enforce that the outcome be a ‘‘homogeneous’’ decision tree where every level of the tree queries the same bit. The relation between the previous theorem and the next is the same as the relation between Green's Theorem 3 and Theorem 4, and the proof adopts the same technique as in [Gre].

Theorem 2. *For every $\delta, \gamma \in (0, 1], \epsilon > 0, 1/\epsilon\delta\gamma > 1$, and $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ such that $\mathbb{E}[f^2] \leq 1$, there is $J \subseteq [n]$ of size at most $2 \uparrow\uparrow 1/\epsilon\delta\gamma$ such that all but at most a γ fraction of restrictions of f on J have (ϵ, δ) -small noisy influences.*

Proof. A homogeneous decision tree produced from Theorem 1 will yield the desired J by letting J be all coordinates split on all levels. To produce such a decision tree we can perform the exact same iteration as in the proof of Theorem 1, with the following modification: suppose at level k we have over γ fraction of leaves that need to be split on high-influence variables x_1, x_2, \dots, x_K . Instead of splitting each leaf on one variable, we split every leaf on every variable from x_1, x_2, \dots, x_K , exactly in that order. Ignoring repeats, we may assume that the x_1, x_2, \dots, x_K are all distinct. The energy difference due to a particular leaf L is

$$\Pr[\text{select leaf } L] \left(\sum_y 2^{-K} \mathbf{Stab}_{1-\delta}[f_{L,y}] - \mathbf{Stab}_{1-\delta}[f_L] \right)$$

where the $f_{L,y}$ range over all possible restrictions y of the variables x_1, x_2, \dots, x_K . Notice that this is independent of the order in which the variables x_1, x_2, \dots, x_K were split. Hence the energy difference is equal to that in which we first split L on the x_i with high influence, and then on the remaining x_1, \dots, x_K . This shows that the energy increase is again at least $\epsilon\delta\gamma$, though now every level queries the same variable.

Suppose the total depth of our decision tree after k iterations is $d(k)$. We can have at most $K \leq 2^{d(k)}$ splits on the next iteration, hence the total depth of the tree satisfies

$$d(k+1) \leq d(k) + 2^{d(k)} \quad d(0) = 0$$

By induction we prove $d(k) \leq 2 \uparrow\uparrow k - 1$. Checking the induction step,

$$\begin{aligned} d(k+1) &\leq 2 \uparrow\uparrow k - 1 + \frac{1}{2} \cdot 2 \uparrow\uparrow (k+1) \leq \frac{1}{2} \cdot 2 \uparrow\uparrow (k+1) - 1 + \frac{1}{2} \cdot 2 \uparrow\uparrow (k+1) \\ &= 2 \uparrow\uparrow (k+1) - 1 \end{aligned}$$

\square

3 Application: Quasirandom Functions and Majority is Stablest

The (mean 0) Majority is Stablest Theorem, first proven in [MOO05], says that among Boolean functions with mean 0, Maj_n has asymptotically the highest noise stability. In its full generality, the Majority is Stablest Theorem bounds the noise stability of an arbitrary Boolean function f by a function dependent on $\mathbb{E}[f]$.

Definition 3.1. Fix $\rho \in [0, 1]$. The Gaussian quadrant probability $\Lambda_\rho(\mu) : [0, 1] \rightarrow [0, 1]$ is defined by

$$\Lambda_\rho(\mu) = \Pr[z_1 \leq t \cap z_2 \leq t]$$

where z_1, z_2 are standard Gaussians with correlation ρ , and t is the inverse of the standard Gaussian CDF at μ i.e. t is such that the area under the standard Gaussian and to the left of t is μ .

Theorem 5. (General-Volume Majority Is Stablest Theorem) Let $f : \{-1, 1\}^n \rightarrow [0, 1]$ such that f has $(\epsilon, \frac{1}{\log(1/\epsilon)})$ -small noisy influences. Then for any $0 \leq \rho < 1$,

$$\text{Stab}_\rho[f] \leq \Lambda_\rho(\mathbb{E}[f]) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1 - \rho}$$

Question: can we get by with a weaker notion of pseudorandomness than small noisy influences? We prove that the Majority is Stablest Theorem still holds if we replace “small noisy influences” with “small low-degree Fourier coefficients”.

Definition 3.2. We say that $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ is (ϵ, δ) -quasirandom if $|\widehat{f}(S)| \leq \epsilon$ for $0 < |S| \leq 1/\delta$.

Informally, a quasirandom function is one which has small low-degree Fourier coefficients. It is strictly a weaker condition than having small (noisy) influences:

Lemma 3.1. Suppose $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ has (ϵ^2, δ) -small noisy influences for some $\delta \in (0, 1/2)$, $\epsilon > 0$. Then f is $(\epsilon, O(\delta))$ -quasirandom.

Proof. We go by contrapositive. Suppose there is a Fourier coefficient $|\widehat{f}(S)| > \epsilon$ for some small S , $|S| \leq 1/\delta$. For any $i \in S$, evaluating the Fourier formula for noisy influences on $D_i f$,

$$\mathbf{Inf}_i^{(1-O(\delta))}[f] = \sum_{R \ni i} (1 - O(\delta))^{|R|-1} \widehat{f}(R)^2$$

for an appropriate linear factor $O(\delta)$ to be chosen later.

$$\sum_{R \ni i} (1 - O(\delta))^{|R|-1} \widehat{f}(R)^2 \geq (1 - O(\delta))^{|S|-1} \widehat{f}(S)^2 > (1 - O(\delta))^{1/\delta} \epsilon^2 \geq \epsilon^2$$

We choose the constant in $O(\delta)$ so that the last inequality holds. Divide all δ in the proof by the appropriate constant to prove the stated claim. \square

We will prove that the Majority is Stablest Theorem holds under the assumption that f is $(o(1), o(1))$ -quasirandom. That is, there are quasirandomness parameters that tend to 0 so that any f satisfying those parameters also satisfy the Majority is Stablest inequality. Here is the generalization we prove:

Theorem 6. For $f : \{-1, 1\}^n \rightarrow [0, 1]$ such that f is $(\epsilon 2^{-\frac{\log^2(1/\epsilon)}{\epsilon \log \log(1/\epsilon)}}, \epsilon \frac{\log \log(1/\epsilon)}{\log^2(1/\epsilon)})$ -quasirandom, and for any $0 \leq \rho < 1$,

$$\text{Stab}_\rho[f] \leq \Lambda_\rho(\mathbb{E}[f]) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1 - \rho}$$

The following lemma is philosophically the reason for using the regularity lemma with quasirandom functions. Informally, quasirandom functions don't change their means too much on restriction. This is Proposition 6.12 of [O'D14].

Lemma 3.2. *Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ and $\epsilon \geq 0, \delta > 0$.*

- (1) *If f is (ϵ, δ) -quasirandom then any restriction of at most $1/\delta$ coordinates changes f 's mean by at most $2^{1/\delta}\epsilon$.*
- (2) *If f is not (ϵ, δ) -quasirandom then some restriction to at most $1/\delta$ coordinates changes f 's mean by more than ϵ .*

3.1 Proof of Theorem 6

Proof. Suppose $f : \{0, 1\}^n \in [0, 1]$ satisfies the conditions of Theorem 6. By Theorem 1 find a decision tree \mathcal{D} computing f of height at most $\frac{\log^2(1/\epsilon)}{\epsilon \log \log(1/\epsilon)}$ so that, for all but at most a $\gamma = \frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}$ fraction of leaves, subfunctions f_L have $(\epsilon, \frac{1}{\log(1/\epsilon)})$ -small noisy influences.

By the General-Volume Majority is Stablest Theorem (Theorem 5), for all but at most γ fraction of subfunctions f_L we have

$$\mathbf{Stab}_\rho[f_L] \leq \Lambda_\rho(\mathbb{E}[f_L]) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho}$$

The functions f_L are restrictions of f to at most $\frac{\log^2(1/\epsilon)}{\epsilon \log \log(1/\epsilon)}$ coordinates (the height of \mathcal{D}). By Lemma 3.2, $\mathbb{E}[f_L]$ is at most

$$2^{\frac{\log^2(1/\epsilon)}{\epsilon \log \log(1/\epsilon)}} \left(\epsilon 2^{-\frac{\log^2(1/\epsilon)}{\epsilon \log \log(1/\epsilon)}} \right) = \epsilon$$

away from $\mathbb{E}[f]$. Indeed, we could have picked any quasirandomness parameters in the statement of this theorem that ensured this bound.

Λ_ρ is 2-Lipschitz (Exercise 11.19 of [O'D14]), thus

$$\begin{aligned} \mathbf{Stab}_\rho[f_L] &\leq \Lambda_\rho(\mathbb{E}[f]) + 2|\mathbb{E}[f] - \mathbb{E}[f_L]| + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho} \\ &\leq \Lambda_\rho(\mathbb{E}[f]) + O(\epsilon) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho} \end{aligned}$$

This bound is independent of L . By looking at the initial energy compared with the final energy in Theorem 1,

$$\mathbf{Stab}_\rho[f] \leq \mathbb{E}_L[\mathbf{Stab}_\rho[f_L]]$$

For γ fraction of leaves we can do no better than upper bound $\mathbf{Stab}_\rho[f_L] \leq \mathbb{E}[f_L^2] \leq 1$. For the rest of the leaves, we drop the $(1-\gamma)$ term and use our derived bound on $\mathbf{Stab}_\rho[f_L]$,

$$\begin{aligned} \mathbf{Stab}_\rho[f] &\leq \gamma + \Lambda_\rho(\mathbb{E}[f]) + O(\epsilon) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho} \\ &= \frac{\log \log(1/\epsilon)}{\log(1/\epsilon)} + \Lambda_\rho(\mathbb{E}[f]) + O(\epsilon) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho} \\ &\leq \frac{\log \log(1/\epsilon)}{\log(1/\epsilon)} \cdot \frac{1}{1-\rho} + \Lambda_\rho(\mathbb{E}[f]) + O(\epsilon) \cdot \frac{1}{1-\rho} + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho} \\ &= \Lambda_\rho(\mathbb{E}[f]) + O\left(\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho} \end{aligned}$$

where the last line follows because $\epsilon \rightarrow 0$ (indeed, any polynomial in ϵ) much faster than $\frac{\log \log(1/\epsilon)}{\log(1/\epsilon)}$. \square

4 References

- [DSTW09] Ilias Diakonikolas, Rocco A. Servedio, Li-Yang Tan, and Andrew Wan. A regularity lemma, and low-weight approximators, for low-degree polynomial threshold functions. 09 2009.
- [Gre] Ben Green. A szemerédi-type regularity lemma in abelian groups, with applications.
- [HLMS14] Kaave Hosseini, Shachar Lovett, Guy Moshkovitz, and Asaf Shapira. An improved lower bound for arithmetic regularity. 05 2014.
- [MOO05] E. Mossel, R. O’Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. In *Proc. 46th Ann. Foundations of Computer Science*, pages 21–30. IEEE Computer Society, Los Alamitos, CA, 2005.
- [O’D14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [OSTW10] Ryan O’Donnell, Rocco Servedio, Li-Yang Tan, and Andrew Wan. A regularity lemma for low noisy-influences. October 2010.
- [TTV09] Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *Proceedings of the 2009 24th Annual IEEE Conference on Computational Complexity, CCC ’09*, pages 126–136, Washington, DC, USA, 2009. IEEE Computer Society.