Advice Lower Bounds for the Dense Model Theorem

Thomas Watson*

December 14, 2012

Abstract

We prove a lower bound on the amount of nonuniform advice needed by black-box reductions for the Dense Model Theorem of Green, Tao, and Ziegler, and of Reingold, Trevisan, Tulsiani, and Vadhan. The latter theorem roughly says that for every distribution D that is δ-dense in a distribution that is ϵ'-indistinguishable from uniform, there exists a "dense model" for D, that is, a distribution that is δ-dense in the uniform distribution and is ϵ-indistinguishable from D. This ϵ-indistinguishability is with respect to an arbitrary small class of functions F. For the natural case where $\epsilon' \geq \Omega(\epsilon\delta)$ and $\epsilon \geq \delta^{O(1)}$, our lower bound implies that $\Omega\big(\sqrt{(1/\epsilon)\log(1/\delta)} \cdot \log|F|\big)$ advice bits are necessary. There is only a polynomial gap between our lower bound and the best upper bound for this case (due to Zhang), which is $O\big((1/\epsilon^2)\log(1/\delta) \cdot \log|F|\big)$. Our lower bound can be viewed as an analog of list size lower bounds for list-decoding of error-correcting codes, but for "dense model decoding" instead.

1 Introduction

The question of whether the prime numbers contain arbitrarily long arithmetic progressions was a long-standing and famous open problem until Green and Tao [GT08] answered the question in the affirmative in a breakthrough paper in 2004. A key ingredient in their proof is a certain transference principle which, very roughly, states the following. Let U denote the set of positive integers. Then for every $D \subseteq U$, if there exists an $R \subseteq U$ such that D is dense in R and R is "indistinguishable" from U, then there exists an $M \subseteq U$ such that M is dense in U and D is "indistinguishable" from M. Tao and Ziegler [TZ08] proved a much more general version of the transference principle, which has come to be known as the Dense Model Theorem (since M is a dense "model" for D).

Reingold, Trevisan, Tulsiani, and Vadhan [RTTV08] demonstrated the relevance of the Dense Model Theorem to computer science, and they gave a new proof which is much simpler and achieves better parameters than the proof of Green, Tao, and Ziegler. Gowers [Gow10] independently came up with a similar proof. In addition to the original application of showing that the primes contain arbitrarily long arithmetic progressions, the Dense Model Theorem has found applications in differential privacy [MPRV09], pseudoentropy and leakage-resilient cryptography [BSW03, RTTV08, DP08], and graph decompositions [RTTV08], as well as further applications in additive combinatorics [GW11b, GW12]. Subsequent variants of the Dense Model Theorem have found applications in cryptography [GW11a] and pseudorandomness [TTV09].

*Computer Science Division, University of California, Berkeley. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-0946797 and by the National Science Foundation under Grant No. CCF-1017403.

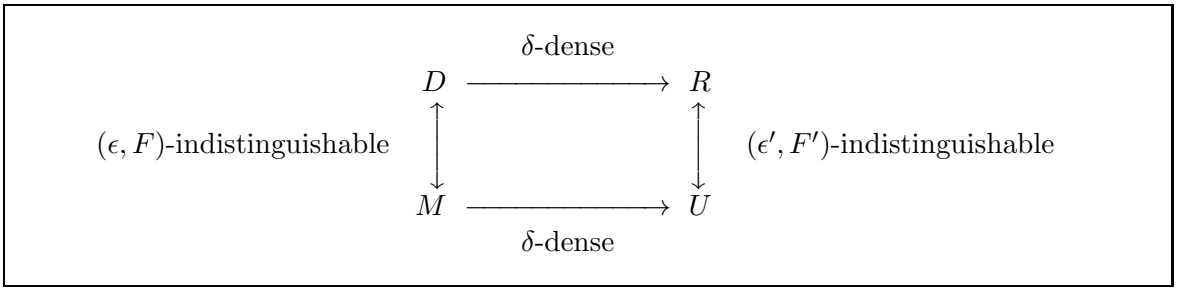

Figure 1: Relations among distributions in the Dense Model Theorem

To formally state the Dense Model Theorem, we first need some definitions. We identify $\{0,1\}^{2^n}$ with the set of functions from $\{0,1\}^n$ to $\{0,1\}$. We use $\mathcal{D}_n$ to denote the set of all distributions on $\{0,1\}^n$. The domain $\{0,1\}^n$ could be replaced by any finite set of size 2^n; we use the domain $\{0,1\}^n$ for concreteness.

Definition 1. *We say $D_1 \in \mathcal{D}_n$ is* δ-dense *in $D_2 \in \mathcal{D}_n$ if for all $x \in \{0,1\}^n$, $\Pr_{D_1}[x] \leq \frac{1}{\delta}\Pr_{D_2}[x]$.*

Definition 2. *We say $f \in \{0,1\}^{2^n}$ ϵ-*distinguishes *$D_1, D_2 \in \mathcal{D}_n$ if $\left|\mathrm{E}_{D_1}[f] - \mathrm{E}_{D_2}[f]\right| > \epsilon$.*

Definition 3. *For $F \subseteq \{0,1\}^{2^n}$, we say $D_1, D_2 \in \mathcal{D}_n$ are* (ϵ, F)-indistinguishable *if there is no $f \in F$ that ϵ-distinguishes D_1 and D_2.*

The following is quantitatively the best known version of the theorem, due to Zhang [Zha11] (building on [RTTV08, BHK09]).

Theorem 1 (Dense Model Theorem). *For every $F \subseteq \{0,1\}^{2^n}$ and every $D \in \mathcal{D}_n$, if there exists an $R \in \mathcal{D}_n$ such that D is δ-dense in R and (R, U) are (ϵ', F')-indistinguishable where $U \in \mathcal{D}_n$ is the uniform distribution, then there exists an $M \in \mathcal{D}_n$ such that M is δ-dense in U and (D, M) are (ϵ, F)-indistinguishable, where $\epsilon' \geq \Omega(\epsilon\delta)$ and F' consists of all linear threshold functions with ± 1 coefficients applied to $O\big((1/\epsilon^2)\log(1/\delta)\big)$ functions from F.*

The relations among the four distributions in Theorem 1 are illustrated in Figure 1. We remark that the theorem also holds when we allow $[0,1]$-valued functions f rather than just $\{0,1\}$-valued functions f. The proof of [RTTV08] gives the same result but where $O\big((1/\epsilon^2)\log(1/\epsilon\delta)\big)$ functions from F are combined to get a function from F'. The original proof of [TZ08] achieves an F' which is qualitatively simpler, namely all products of $\mathrm{poly}(1/\epsilon, 1/\delta)$ functions from F, but it only achieves $\epsilon' \geq \exp(-\mathrm{poly}(1/\epsilon, 1/\delta))$.[1] We note that the dependence $\epsilon' \geq \Omega(\epsilon\delta)$ is tight in two senses.

- The Dense Model Theorem is actually false when $\epsilon' > \epsilon\delta$, even if $F' = \{0,1\}^{2^n}$. See [Zha11] for the simple argument.
- The following converse to the Dense Model Theorem holds: If there exists an $M \in \mathcal{D}_n$ such that M is δ-dense in U and (D, M) are (ϵ, F)-indistinguishable, then there exists an $R \in \mathcal{D}_n$ such that D is δ-dense in R and (R, U) are (ϵ', F')-indistinguishable, where $\epsilon' = \epsilon\delta$ and $F' = F$. To see this, note that $U = \delta M + (1-\delta)\widehat{M}$ for some $\widehat{M} \in \mathcal{D}_n$, so we can let $R = \delta D + (1-\delta)\widehat{M}$; then D is δ-dense in R, and for every $f \in \{0,1\}^{2^n}$ we have $\mathrm{E}_R[f] - \mathrm{E}_U[f] = \delta\big(\mathrm{E}_D[f] - \mathrm{E}_M[f]\big)$ and thus if $\left|\mathrm{E}_R[f] - \mathrm{E}_U[f]\right| > \epsilon'$ then $\left|\mathrm{E}_D[f] - \mathrm{E}_M[f]\right| > \epsilon$.

[1] Another proof that also achieves this is given in [RTTV08].

The Dense Model Theorem has an undesirable feature: The class F' is more complex than the class F. Thus, if we wish to conclude that D and M are indistinguishable for a class F, we need to assume that R and U are indistinguishable for a more complex class F'. The less complex F' is, the stronger the theorem is. The reason for this loss in complexity is because the theorem is proved using a *black-box reduction*. In other words, the contrapositive is proved: We assume that for every M δ-dense in U there exists a function from F that ϵ-distinguishes D and M, and we show that some of these functions can be plugged into the reduction to get a function that ϵ'-distinguishes R and U. Thus the resulting function is necessarily more complex than the functions that get plugged into the reduction. There are three notions of complexity that are interesting to address in this context.

1. *Computational complexity.* If F consists of functions computed by small constant-depth circuits (AC^0), then can we let F' consist of functions computed by (slightly larger) constant-depth circuits? This is not known to be true when $\epsilon' \geq \Omega(\epsilon\delta)$, because the reductions of [RTTV08, Zha11] involve a linear threshold function, which cannot be computed by small constant-depth circuits. Is it necessary that the reduction computes a linear threshold function? The original result of [TZ08] shows that this is *not* necessary if ϵ' is inverse exponentially small.

2. *Query complexity.* If F consists of functions computed by circuits of size s, then F' will need to consist of functions computed by circuits of a larger size s' — but how much larger? If the reduction makes q queries to functions from F, then plugging in size-s circuits for these functions yields a circuit of size $\geq q \cdot s$, and thus we must have $s' \geq q \cdot s$. Hence it is desirable to minimize q. Can we do better than $q \leq O\big((1/\epsilon^2)\log(1/\delta)\big)$ as in Theorem 1?

3. *Advice complexity.* Suppose F consists of functions computed by uniform algorithms running in time t (that is, a single algorithm computes a sequence of functions, one for each $n = 1, 2, 3, \ldots,$). Then can we let F' consist of functions computed by uniform algorithms running in some (slightly larger) time t'? (Here, the distributions D, M, R, U would need to be sequences of distributions, and a distinguisher would only be required to succeed for infinitely many n.) The proofs of [RTTV08, Zha11] do not yield this, because the reductions need a nonuniform advice string to provide some extra information about the nth distribution D.[2] How many bits of advice are needed?

Before proceeding we draw attention to the fact that, as we just alluded to, the advice strings used by the reductions of [RTTV08, Zha11] depend on D *but do not depend on* R. Hence something a little stronger than Theorem 1 actually holds: Although the statement of Theorem 1 says we need to assume that for some R in which D is δ-dense, there is no function in F' that ϵ'-distinguishes R and U, we actually only need to assume that there is no function in F' that simultaneously ϵ'-distinguishes U from every R in which D is δ-dense (the quantifiers are swapped). We are interested in proving lower bounds on the complexity of this type of black-box reduction for the Dense Model Theorem, where the advice does not depend on R.

The *query complexity* was considered by Zhang [Zha11], who showed that for a certain type of nonadaptive black-box reduction, $\Omega\big((1/\epsilon^2)\log(1/\delta)\big)$ queries are necessary when $\epsilon' \geq \Omega(\epsilon\delta)$ and $\epsilon \geq \delta^{O(1)}$, matching the upper bound of $O\big((1/\epsilon^2)\log(1/\delta)\big)$ for this case. In this paper we consider

[2] The paper [RTTV08] also proves a version of the Dense Model Theorem that satisfies a certain technical relaxation of being "uniform", where ϵ figures into the quantification over all algorithms and all n.

the *advice complexity*. We show that for arbitrary black-box reductions, $\Omega\big(\sqrt{(1/\epsilon)\log(1/\delta)}\cdot\log|F|\big)$ advice bits are necessary when $\epsilon' \geq \Omega(\epsilon\delta)$ and $\epsilon \geq \delta^{O(1)}$, which comes close to matching the upper bound of $O\big((1/\epsilon^2)\log(1/\delta)\cdot\log|F|\big)$ for this case. Our result also holds for much more general settings of the parameters (with some degradation in the lower bound). Proving lower bounds on the *computational complexity* remains open.

Let us formally state what we mean by a black-box reduction. Recall the standard notation $[k] = \{1,\ldots,k\}$.

Definition 4. *An* $(n,\epsilon,\delta,\epsilon',k,\alpha)$-reduction *(for the Dense Model Theorem) is a function*

$$\mathrm{Dec} : \big(\{0,1\}^{2^n}\big)^k \times \{0,1\}^\alpha \to \{0,1\}^{2^n}$$

such that for all $f_1,\ldots,f_k \in \{0,1\}^{2^n}$ *and all* $D \in \mathcal{D}_n$*, if for every* $M \in \mathcal{D}_n$ *that is* δ*-dense in the uniform distribution* $U \in \mathcal{D}_n$ *there exists an* $i \in [k]$ *such that* f_i ϵ*-distinguishes* D *and* M*, then there exists an advice string* $a \in \{0,1\}^\alpha$ *such that for every* $R \in \mathcal{D}_n$ *in which* D *is* δ*-dense,* $\mathrm{Dec}(f_1,\ldots,f_k,a)$ ϵ'*-distinguishes* R *and* U*.*

The proofs of [RTTV08, Zha11] work by exhibiting such reductions. The functions $\{f_1,\ldots,f_k\}$ correspond to the class F (which, if we were considering uniform algorithms, would be the restrictions of all the algorithms in the class to a particular input length n). We now state our theorem.

Theorem 2 (Main). *If there exists an* $(n,\epsilon,\delta,\epsilon',k,\alpha)$*-reduction for the Dense Model Theorem, and if* $w > 1$ *is an integer such that* $2^{w+2}\cdot\delta^{w/160} \leq \epsilon'$*, then*

$$\alpha \;\geq\; \big\lfloor \tfrac{1}{160w}\sqrt{(1/\epsilon)\log_2(1/\delta)}\big\rfloor \cdot \log_2 k - \log_2 w - 1$$

provided $2^n \geq \frac{w\log_2 k}{\epsilon\delta^2(\epsilon')^2}$*,* $\epsilon \leq 1/64\log_2(1/\delta)$*, and* $k \geq 1/16\epsilon^4$*.*

For the case where $\epsilon' \geq \Omega(\epsilon\delta)$ and $\epsilon \geq \delta^{O(1)}$ (which is reasonable), the condition $2^{w+2}\cdot\delta^{w/160} \leq \epsilon'$ is met provided w is a sufficiently large constant and δ is less than a sufficiently small constant,[3] and thus we get a lower bound $\alpha \geq \Omega\big(\sqrt{(1/\epsilon)\log(1/\delta)}\cdot\log k\big)$. Note that the three conditions at the end of the statement of Theorem 2 are very generous.[4]

We point out that the black-box reductions we consider can be viewed as a kind of analog of list-decoders for error-correcting codes. In the setting of list-decoding, the decoder is given a received word and is required to output a list containing all the codewords that are within a certain Hamming distance of the received word. In the setting of "dense model decoding", the distribution D is analogous to the correct codeword, and the decoder is only provided with some "corrupted" information about D, namely the functions $f_1,\ldots,f_k$. Note that these functions contain some information about D, since D can be distinguished from any M δ-dense in U using one of the f_i's. The decoder must output a list (corresponding to all possible values of the advice string $a \in \{0,1\}^\alpha$), but the goal is less ambitious than finding D itself; the list just needs to contain a function that distinguishes U from every R in which D is δ-dense. Thus advice lower bounds

[3] The statement of Theorem 2 requires $\delta < 2^{-160}$. This constant can be drastically improved; we chose 2^{-160} since it is convenient for the proof.

[4] The bound $2^n \geq \frac{w\log_2 k}{\epsilon\delta^2(\epsilon')^2}$ can be relaxed somewhat. We chose this bound since it is reasonable and is convenient to state.

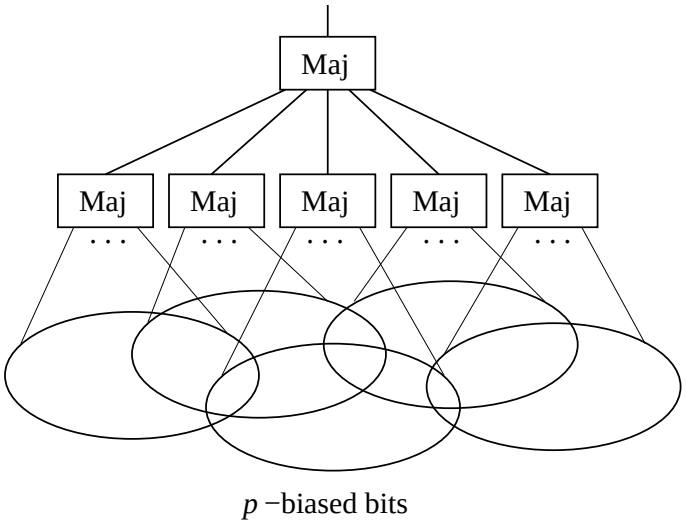

Figure 2: The majority of majorities

for the Dense Model Theorem are a kind of analog of list size lower bounds for list-decoding. See Section 1.1 for previous work on list size lower bounds (as well as other previous work that is relevant to the topic of this paper). For the case of approximate local list-decoding (also known as hardness amplification), getting the number of advice bits from $\text{poly}(1/\epsilon)$ down to the lower bound of $\Omega(\log(1/\epsilon))$ proved to be quite a challenge [IJK09, IJKW10]. In contrast, we show that in the setting of the Dense Model Theorem, the known advice lengths are already in the right regime, namely $\text{poly}(1/\epsilon)$.

The rest of this paper is devoted to proving Theorem 2. In Section 2 we give some intuition for the proof, and then in Section 3 we give the formal proof. We now give a quick preview of some of the ingredients that go into the proof. We use the probabilistic method to find a class of functions $f_1, \ldots, f_k$ for which many advice strings are needed to "cover" all the distributions D that do not have dense models. The key technical ingredients in the analysis include (1) a combinatorial argument identifying when several distributions D cannot share the same advice string, and (2) an analysis of a majority of majorities applied to overlapping sets of p-biased bits, where the sets form an almost-disjoint family (see Figure 2). The latter analysis makes use of extremely tight lower bounds on the tail probabilities of the binomial distribution, which we also prove.

1.1 Related Work

Lu, Tsai, and Wu [LTW11] proved lower bounds on the computational complexity, query complexity, and advice complexity of black-box reductions for the Hardcore Lemma (which was introduced by Impagliazzo [Imp95]). Our proof bears some similarity to their advice lower bound proof, but it diverges significantly. Zhang [Zha11] proved tight lower bounds on the query complexity of nonadaptive black-box reductions for the Dense Model Theorem, again with a proof somewhat reminiscent of the corresponding argument in [LTW11].

There has been extensive work on lower bounds for black-box hardness amplification and list-decoding of error-correcting codes. Regarding *advice complexity*, Guruswami and Vadhan [GV10] and Blinovsky [Bli86] proved a tight $\Omega(1/\epsilon^2)$ list size lower bound for decoding arbitrary binary

error-correcting codes up to radius $1/2 - \epsilon$. Lu, Tsai, and Wu [LTW08] proved similar advice lower bounds for black-box hardness amplification. List size lower bounds for decoding from erasures can be found in [Gur03, Wat11]. Regarding *query complexity*, Shaltiel and Viola [SV10] proved lower bounds for nonadaptive black-box hardness amplification, matching the upper bounds known for the XOR Lemma [Lev87, Imp95, GNW11, KS03]. Artemenko and Shaltiel [AS11] proved lower bounds for adaptive black-box hardness amplification, which are not tight in general but are tight in some settings. Regarding *computational complexity*, Shaltiel and Viola [SV10] showed that decoders for black-box hardness amplification must implicitly compute majority on $\Theta(1/\epsilon)$ bits and hence such decoders cannot be implemented with small constant-depth circuits when ϵ is small. Gutfreund and Rothblum [GR08] showed that adaptive local list-decoders for binary codes must implicitly compute majority on $\Theta(1/\epsilon)$ bits to handle radius $1/2 - \epsilon$, under a restriction on the list size.

2 Intuition

According to Definition 4, for Dec to succeed as a reduction, it must be the case that for all $f_1, \ldots, f_k \in \{0,1\}^{2^n}$ and all $D \in \mathcal{D}_n$, if D has no "dense model" then there is some advice string a such that $\mathrm{Dec}(f_1, \ldots, f_k, a)$ "covers" D in a certain sense. To show that Dec needs many advice strings in order to succeed, we find functions $f_1, \ldots, f_k \in \{0,1\}^{2^n}$ and a large family of distributions in $\mathcal{D}_n$ such that

(i) each distribution in the family has no dense model (with respect to $f_1, \ldots, f_k$), and

(ii) each function $f \in \{0,1\}^{2^n}$ covers few of the distributions in the family.

So (i) implies that each distribution in the family needs to get covered, while (ii) implies that for each advice string a, $\mathrm{Dec}(f_1, \ldots, f_k, a)$ does not cover very many of them. Since the family is large, many advice strings are needed.

First we describe a technique for achieving (i), then we describe a technique for achieving (ii), and then we show how to consolidate the techniques to achieve both properties simultaneously. When we say D has no "dense model" we mean that for every $M \in \mathcal{D}_n$ that is δ-dense in U there exists an $i \in [k]$ such that f_i ϵ-distinguishes D and M. When we say a function "covers" D we mean that it ϵ'-distinguishes R and U for every $R \in \mathcal{D}_n$ in which D is δ-dense. The only distributions D we need to consider are uniform distributions over subsets of $\{0,1\}^n$.

Given $f_1, \ldots, f_k \in \{0,1\}^{2^n}$, what is an example of a distribution with no dense model? Suppose we pick any $I \subseteq [k]$ of size $1/4\epsilon$ and we let X_I be the set of all $x \in \{0,1\}^n$ such that $f_i(x) = 1$ for the majority of $i \in I$. Suppose we take D_I to be the uniform distribution over X_I. Then we have $\Pr_{x \sim D_I,\ i \sim I}[f_i(x) = 1] \geq 1/2 + 2\epsilon$ where $i \sim I$ means picking $i \in I$ uniformly at random. If X_I is roughly a $\delta/2$ fraction of $\{0,1\}^n$, then every distribution M that is δ-dense in U has at least half its mass outside of X_I, on strings x where $\Pr_{i \sim I}[f_i(x) = 1] \leq 1/2 - 2\epsilon$. It is possible to show that $\Pr_{x \sim M,\ i \sim I}[f_i(x) = 1] < \Pr_{x \sim D_I,\ i \sim I}[f_i(x) = 1] - \epsilon$ and thus there exists an $i \in I$ (depending on M) such that f_i ϵ-distinguishes D_I and M. So if $|X_I| \approx (\delta/2)2^n$ then D_I has no dense model. This is the technique we use for finding distributions without dense models.

Now, what is an example of a pair of distributions such that no function can cover both simultaneously? If we can show that every pair of distributions in the family is like this, then we will have achieved (ii). Because of an issue described below, we actually need to consider small collections of distributions rather than just pairs, but for now we consider pairs. Suppose D is uniform over

some $X \subseteq \{0,1\}^n$ of size roughly $(\delta/2)2^n$, and similarly D' is uniform over some $X' \subseteq \{0,1\}^n$ of size roughly $(\delta/2)2^n$. If $X \cap X' = \emptyset$, then it can be shown that no function covers both D and D'.[5] Furthermore, if $|X \cap X'|$ is at most roughly $\epsilon' 2^n$ then this property still holds.

To consolidate the two techniques, we would like to find a large family of sets $I \subseteq [k]$ each of size $1/4\epsilon$, such that

(A) $|X_I| \approx (\delta/2)2^n$ for each I in the family, and

(B) the pairwise intersections of the X_I's (for I in the family) all have size at most roughly $\epsilon' 2^n$.

This would imply that the corresponding distributions D_I (for I in the family) have no dense models, and no function would cover more than one of them, so (i) and (ii) would be achieved.

We choose the functions $f_1, \ldots, f_k \in \{0,1\}^{2^n}$ randomly in some way, and we argue that for an appropriate family of sets I, properties (A) and (B) both hold with high probability. Property (A) suggests that we should choose p so that the probability a majority of $1/4\epsilon$ independent coins each with expectation p come up 1 is exactly $\delta/2$. Then we can set $f_i(x) = 1$ with probability p independently for each $i \in [k]$ and each $x \in \{0,1\}^n$, so for each I of size $1/4\epsilon$, $\Pr[x \in X_I] = \delta/2$. Then by concentration, $|X_I| \approx (\delta/2)2^n$ with high probability over $f_1, \ldots, f_k$.

If we choose $f_1, \ldots, f_k$ randomly in this way, how big will $|X_I \cap X_{I'}|$ be, for I and I' in the family? By concentration, we would have that with high probability over $f_1, \ldots, f_k$, $|X_I \cap X_{I'}|$ is roughly 2^n times $\Pr[x \in X_I \cap X_{I'}]$ (which is the same for all $x \in \{0,1\}^n$), so we would like the latter probability to be $\leq \epsilon'$. So what is the probability that the conjunction of two majorities of p-biased bits is 1? The best case is if $I \cap I' = \emptyset$, in which case the probability is exactly $(\delta/2)^2$. There are two problems with this.

(1) We cannot get a very large family of sets I if we require them to be pairwise disjoint.

(2) This requires $\epsilon' \geq (\delta/2)^2$. In a typical setting where $\epsilon' \geq \Omega(\epsilon\delta)$, this would require $\epsilon > \delta$, which is an odd and somewhat severe restriction.

To solve problem (1), we use the natural idea to allow the sets I to be pairwise almost-disjoint, rather than disjoint (which allows us to get a much larger family). So if $|I \cap I'|$ is at most some value b, how small does b have to be to ensure that the probability both majorities are 1 is not much more than $(\delta/2)^2$? We analyze this using the following trick: If both majorities are 1, then the fraction of coins that are 1 among $I \cup I'$ is at least q, where $q = 1/2 - 2\epsilon b = \frac{1/4\epsilon - b}{1/2\epsilon} \leq \frac{|I|/2 + |I'|/2 - b}{|I \cup I'|}$. Using an extremely tight characterization of the tail probabilities of the binomial distribution (which we prove using known techniques but which we could not find in the literature), we can show that $p \approx 1/2 - \sqrt{\epsilon \log(1/\delta)}$ and the probability of getting $\geq q$ fraction of 1's among the $|I \cup I'|$ coins is not much more than $(\delta/2)^2$ provided q is at least a constant factor closer to $1/2$ than p is, say $q \approx 1/2 - \sqrt{\epsilon \log(1/\delta)}/4$. Thus it suffices to have $b \approx \sqrt{\epsilon \log(1/\delta)}/8\epsilon \geq \Omega\big(\sqrt{(1/\epsilon)\log(1/\delta)}\big)$. Since the family of sets I needs to be in the universe $[k]$, there exists such a family of roughly k^b many sets with pairwise intersections bounded in size by b. Since each function can cover D_I for only one I in the family, roughly k^b advice strings are needed, which gives an advice lower bound of roughly $\log(k^b) \geq \Omega\big(\sqrt{(1/\epsilon)\log(1/\delta)} \cdot \log k\big)$.

Problem (2) is solved in the formal proof by considering small collections of sets from the family, rather than pairs. The parameter w in Theorem 2 is used to determine how big these collections

[5] Actually, there is an issue having to do with the absolute value signs in the definition of distinguishing; this is dealt with in the formal proof.

should be. Then instead of requiring that the conjunction of two majorities accepts with small probability, we need that the majority of several majorities accepts with small probability, which explains where Figure 2 comes from.

3 Formal Proof

In Section 3.1, Section 3.2, and Section 3.3 we give preliminary lemmas, definitions, and notation. Then in Section 3.4 we give the proof of Theorem 2.

3.1 Binomial Distribution Tail

We let $\text{Tail}(m, p, q)$ denote the probability that when m independent coins are flipped each with probability p of heads, at least a q fraction of the coins are heads (in other words, the probability the (m, p) binomial distribution is at least qm). For our proof of Theorem 2 we need extremely tight upper and lower bounds on the value of $\text{Tail}(m, p, q)$. Such bounds can be given in terms of the fundamental quantity

$$\text{RE}(q\|p) \;=\; q\log_2(\tfrac{q}{p}) + (1-q)\log_2(\tfrac{1-q}{1-p})$$

which is known by a variety of names such as relative entropy, information divergence, and Kullback-Leibler distance.[6]

We need the following fact, which can be seen using derivatives.

Fact 1. *For all $1/4 \le p \le q \le 3/4$, we have $2(q-p)^2 \le \text{RE}(q\|p) \le 4(q-p)^2$.*

We also need the following standard and well-known form of the Chernoff-Hoeffding bound.

Lemma 1. *For all $m \ge 1$ and all $0 \le p \le q \le 1$, we have $\text{Tail}(m, p, q) \le 2^{-\text{RE}(q\|p)m}$.*

Lemma 1 is very tight, as shown by the following lemma, which we prove for completeness.

Lemma 2. *For all $m \ge 1$ and all $1/4 \le p \le q \le 1$, we have $\text{Tail}(m, p, q) \ge \frac{1}{48\sqrt{m}} \cdot 2^{-\text{RE}(q\|p)m}$.*

Proof. First, assume that qm is an integer. Then lower bounding $\text{Tail}(m, p, q)$ by the first term of the sum, we have

$$\begin{aligned}
2^{\text{RE}(q\|p)m} \cdot \text{Tail}(m,p,q) &\ge 2^{\text{RE}(q\|p)m} \cdot \tbinom{m}{qm} p^{qm}(1-p)^{(1-q)m} \\
&= \tfrac{q^{qm}(1-q)^{(1-q)m}}{p^{qm}(1-p)^{(1-q)m}} \cdot \tbinom{m}{qm} p^{qm}(1-p)^{(1-q)m} \\
&= q^{qm}(1-q)^{(1-q)m} \cdot \tbinom{m}{qm} \\
&\ge q^{qm}(1-q)^{(1-q)m} \cdot \tfrac{1}{3\sqrt{qm}} \cdot \tfrac{1}{q^{qm}(1-q)^{(1-q)m}} \\
&\ge \tfrac{1}{3\sqrt{m}}
\end{aligned}$$

where the fourth line follows by Stirling approximations. Now suppose qm is not an integer, and let $q' = \lceil qm \rceil / m$. Then we have $\text{Tail}(m, p, q) = \text{Tail}(m, p, q') \ge \frac{1}{3\sqrt{m}} \cdot 2^{-\text{RE}(q'\|p)m}$. We claim

[6] RE is more often notated D or D_{KL}, but we use RE to avoid confusion with our distributions D.

that $\mathrm{RE}(q'\|p) - \mathrm{RE}(q\|p) \le 4/m$, from which it follows that $2^{-\mathrm{RE}(q'\|p)m}/2^{-\mathrm{RE}(q\|p)m} \ge 1/16$ and thus $\mathrm{Tail}(m,p,q) \ge \frac{1}{48\sqrt{m}} \cdot 2^{-\mathrm{RE}(q\|p)m}$. We now argue the claim. Since $q' \le q + 1/m$, we have $q'\log_2(\frac{q'}{p}) - q\log_2(\frac{q}{p}) \le (1/m)\log_2(\frac{1}{p}) + q\log_2(\frac{q+1/m}{q})$. We have $(1/m)\log_2(\frac{1}{p}) \le 2/m$ since $p \ge 1/4$, and we have $q\log_2(\frac{q+1/m}{q}) \le q \cdot 2/qm = 2/m$. Thus $q'\log_2(\frac{q'}{p}) - q\log_2(\frac{q}{p}) \le 4/m$. Since $q' \ge q$, we have $(1-q')\log_2(\frac{1-q'}{1-p}) - (1-q)\log_2(\frac{1-q}{1-p}) \le 0$. Summing gives the claim. □

Although Lemma 2 is very simple and general, for our purpose we can only use it for a limited range of parameters, namely when $\epsilon \gg \delta$. This is because $\mathrm{RE}(q\|p)$ could be so close to 0 that $\frac{1}{48\sqrt{m}}$ completely swamps $2^{-\mathrm{RE}(q\|p)m}$, in which case Lemma 2 is not very tight. To handle the full range of ϵ and δ, we use the following stronger lower bound for the case $q = 1/2$.

Lemma 3. *For all $m \ge 9$ and all $1/4 \le p < 1/2$, we have*

$$\mathrm{Tail}(m,p,1/2) \;\ge\; \min\left(\tfrac{1}{256},\; \tfrac{1}{128\sqrt{m}(1/2-p)}\right) \cdot 2^{-\mathrm{RE}(1/2\|p)m}.$$

Proof. Let $q = \lceil m/2 \rceil / m$. Let $h = \lfloor \sqrt{m}/3 \rfloor$, and note that $1 \le h \le (1-q)m$. We have

$$\begin{aligned}
2^{\mathrm{RE}(q\|p)m} \cdot \mathrm{Tail}(m,p,q) &= \tfrac{q^{qm}(1-q)^{(1-q)m}}{p^{qm}(1-p)^{(1-q)m}} \cdot \textstyle\sum_{i=0}^{(1-q)m} \binom{m}{qm+i} p^{qm+i}(1-p)^{(1-q)m-i} \\
&= q^{qm}(1-q)^{(1-q)m} \cdot \textstyle\sum_{i=0}^{(1-q)m} \binom{m}{qm+i} \left(\tfrac{p}{1-p}\right)^i \\
&= q^{qm}(1-q)^{(1-q)m} \cdot \textstyle\sum_{i=0}^{(1-q)m} \binom{m}{qm} \left(\tfrac{p}{1-p}\right)^i \prod_{j=1}^{i} \left(1 + \tfrac{m-2qm-i}{qm+j}\right) \\
&\ge \tfrac{1}{3\sqrt{qm}} \cdot \textstyle\sum_{i=0}^{(1-q)m} \left(\tfrac{p}{1-p}\right)^i \prod_{j=1}^{i} \left(1 + \tfrac{m-2qm-i}{qm+j}\right) \\
&\ge \tfrac{1}{3\sqrt{qm}} \cdot \textstyle\sum_{i=0}^{(1-q)m} \left(\tfrac{p}{1-p}\right)^i \left(1 - \tfrac{2(i+1)}{m}\right)^i \\
&\ge \tfrac{1}{3\sqrt{qm}} \cdot \left(1 - \tfrac{2(h+1)}{m}\right)^h \cdot \textstyle\sum_{i=0}^{h} \left(\tfrac{p}{1-p}\right)^i \\
&\ge \tfrac{1}{3\sqrt{qm}} \cdot \left(1 - \tfrac{2h(h+1)}{m}\right) \cdot \textstyle\sum_{i=0}^{h} \left(\tfrac{p}{1-p}\right)^i \\
&\ge \tfrac{1}{6\sqrt{qm}} \cdot \textstyle\sum_{i=0}^{h} \left(\tfrac{p}{1-p}\right)^i \\
&= \tfrac{1}{6\sqrt{qm}} \cdot \tfrac{1-(\frac{p}{1-p})^{h+1}}{1-\frac{p}{1-p}} \\
&\ge \tfrac{1}{6\sqrt{qm}} \cdot \tfrac{1-e^{-(1-\frac{p}{1-p})h}}{1-\frac{p}{1-p}}
\end{aligned}$$

where the fourth line follows by $\binom{m}{qm} \ge \frac{1}{3\sqrt{qm}} \cdot \frac{1}{q^{qm}(1-q)^{(1-q)m}}$ which holds by Stirling approximations, the fifth line follows by $1/2 \le q \le 1/2 + 1/2m$, and the eighth line follows by the definition of h and $m \ge 9$. If $(1 - \frac{p}{1-p})h < 1$ then the expression is at least

$$\tfrac{1}{6\sqrt{qm}} \cdot \tfrac{1-\left(1-\frac{1}{2}(1-\frac{p}{1-p})h\right)}{1-\frac{p}{1-p}} \;=\; \tfrac{h}{12\sqrt{qm}} \;\ge\; \tfrac{1}{64}.$$

If $(1 - \frac{p}{1-p})h \ge 1$ then the expression is at least

$$\tfrac{1}{6\sqrt{qm}} \cdot \tfrac{1-1/e}{1-\frac{p}{1-p}} \;\ge\; \tfrac{1}{10\sqrt{qm}} \cdot \tfrac{1}{1-\frac{p}{1-p}} \;\ge\; \tfrac{1}{10\sqrt{qm}} \cdot \tfrac{1}{4\cdot(1/2-p)} \;\ge\; \tfrac{1}{32\sqrt{m}(1/2-p)}$$

where the last inequality uses $m \geq 9$. Thus we have shown that in either case,

$$\text{Tail}(m,p,1/2) \;=\; \text{Tail}(m,p,q) \;\geq\; \min\left(\tfrac{1}{64},\; \tfrac{1}{32\sqrt{m(1/2-p)}}\right) \cdot 2^{-\,\text{RE}(q\|p)m}.$$

To finish the proof, we just need to show that $2^{-\,\text{RE}(q\|p)m}/2^{-\,\text{RE}(1/2\|p)m} \geq 1/4$. For this, it suffices to show that $\text{RE}(q\|p) - \text{RE}(1/2\|p) \leq 2/m$. Since $q \leq 1/2+1/2m$, we have $q\log_2(\frac{q}{p}) - (1/2)\log_2(\frac{1/2}{p}) \leq (1/2m)\log_2(\frac{1}{p}) + (1/2)\log_2(\frac{1/2+1/2m}{1/2})$. We have $(1/2m)\log_2(\frac{1}{p}) \leq 1/m$ since $p \geq 1/4$, and we have $(1/2)\log_2(\frac{1/2+1/2m}{1/2}) \leq (1/2)\cdot 2/m = 1/m$. Thus $q\log_2(\frac{q}{p}) - (1/2)\log_2(\frac{1/2}{p}) \;\leq\; 2/m$. Since $q \geq 1/2$, we have $(1-q)\log_2(\frac{1-q}{1-p}) - (1-1/2)\log_2(\frac{1-1/2}{1-p}) \leq 0$. Summing yields $\text{RE}(q\|p) - \text{RE}(1/2\|p) \leq 2/m$. □

3.2 Combinatorial Designs

For our proof of Theorem 2 we need the existence of large families of almost-disjoint subsets of a finite set. Such combinatorial designs have numerous applications in theoretical computer science, one of the more famous being in the pseudorandom generator construction of Nisan and Wigderson [NW94].

Definition 5. *An* (ℓ, k, s, b)-design *is a family of sets* $I_1, \ldots, I_\ell \subseteq [k]$ *all of size* s *such that* $|I_j \cap I_{j'}| \leq b$ *for every* $j \neq j'$.

Lemma 4. *For every* k, s, b *there exists an* (ℓ, k, s, b)*-design with* $\ell \geq k^{b/8}$*, provided* $k \geq 16s^4$.

There is nothing very novel about this lemma, and this precise version follows from a result in [EFF85], but we provide a simple, self-contained proof here. The proof uses the probabilistic method with a simple concentration bound for the hypergeometric distribution.

Proof. If $b = 0$ then the lemma holds trivially, so assume $b \geq 1$. Let $\ell = \lceil k^{b/8} \rceil$. We pick sets $I_1, \ldots, I_\ell$ independently and uniformly at random from all subsets of $[k]$ of size s, and we argue that with positive probability $I_1, \ldots, I_\ell$ forms an (ℓ, k, s, b)-design. We claim that for every $j, j' \in [\ell]$ with $j \neq j'$, $\Pr\left[|I_j \cap I_{j'}| > b\right] \leq 2k^{-b/2}$. From this it follows by a union bound that

$$\Pr\left[I_1, \ldots, I_\ell \text{ does not form an } (\ell,k,s,b)\text{-design}\right] \;\leq\; \tbinom{\ell}{2} \cdot 2k^{-b/2} \;<\; \ell^2 \cdot k^{-b/2} \;\leq\; 1$$

where the final inequality $\lceil k^{b/8} \rceil \leq k^{b/4}$ follows by $k \geq 16$ and $b \geq 1$. To prove the claim, consider any $j \neq j'$ and fix any particular choice of I_j. Now consider picking points $i_1, \ldots, i_s \in [k]$ independently and uniformly at random (with replacement). Since the expected number of points that land in I_j is s^2/k, a standard relative-error form of the Chernoff bound tells us that

$$\Pr_{i_1,\ldots,i_s}\left[|\{h \;:\; i_h \in I_j\}| > b\right] \;\leq\; \left(\tfrac{es^2}{bk}\right)^b \;\leq\; k^{-b/2}$$

using $es^2 \leq k^{1/2}$ and $b \geq 1$ (where e is the base of the natural logarithm). The probability that $i_1, \ldots, i_s$ are all distinct is at least $1 - \binom{s}{2}/k \geq 1/2$ by a union bound and using $k \geq s^2$. Conditioning on the event that $i_1, \ldots, i_s$ are all distinct is equivalent to sampling $I_{j'} = \{i_1, \ldots, i_s\}$ of size s uniformly at random, and probabilities increase by at most a factor of 2 conditioned on this event. Thus $\Pr\left[|I_j \cap I_{j'}| > b\right] \leq 2k^{-b/2}$, as claimed. □

3.3 Notational Preliminaries

The parameters $n, \epsilon, \delta, \epsilon', k$, and w are fixed as in the statement of Theorem 2, and we always use D, M, R, U (possibly subscripted) to denote distributions in $\mathcal{D}_n$, in their respective roles as in Definition 4.

We let Maj denote the majority function on bit strings, and for even length strings we break ties by returning 1. We let And denote the and function on bit strings. We let Maj^t denote the function that takes t bit strings and returns their majorities as a length-t bit string. We use $\circ$ for function composition.

We also adhere to the following notational conventions. We use x for elements of $\{0,1\}^n$ and X for subsets of $\{0,1\}^n$. We use f for elements of $\{0,1\}^{2^n}$ (identified with functions from $\{0,1\}^n$ to $\{0,1\}$) and F for subsets of $\{0,1\}^{2^n}$. We use $[k]$ to index functions f, and we use i for elements of $[k]$ and I for subsets of $[k]$. We use $[\ell]$ to index subsets I (as in Definition 5), and we use j for elements of $[\ell]$ and J for subsets of $[\ell]$. Furthermore, we generally use s for the size of I, and t for the size of J.

The following notation is with respect to fixed $f_1, \ldots, f_k \in \{0,1\}^{2^n}$. Given $I \subseteq [k]$ we define

- f_I is the function that takes $x \in \{0,1\}^n$ and returns the length-$|I|$ bit string $(f_i(x))_{i \in I}$;
- X_I is the set of $x \in \{0,1\}^n$ on which $\mathrm{Maj} \circ f_I$ returns 1;
- D_I is the uniform distribution over X_I (and if $X_I = \emptyset$ then D_I is undefined).

The following notation is with respect to fixed $f_1, \ldots, f_k \in \{0,1\}^{2^n}$ and fixed $I_1, \ldots, I_\ell \subseteq [k]$. Given $J \subseteq [\ell]$ we define

- f_{I_J} is the function that takes $x \in \{0,1\}^n$ and returns the $|J|$-tuple $(f_{I_j}(x))_{j \in J}$;
- X_{I_J} is the set of $x \in \{0,1\}^n$ on which $\mathrm{Maj} \circ \mathrm{Maj}^{|J|} \circ f_{I_J}$ returns 1.

We use $\sim$ to denote sampling from a distribution (for example $x \sim D$), and we use the convention that sampling from a set (for example $i \sim I$) means sampling from the uniform distribution over that set.

3.4 Proof of Theorem 2

Consider an arbitrary function

$$\mathrm{Dec} : \left(\{0,1\}^{2^n}\right)^k \times \{0,1\}^\alpha \to \{0,1\}^{2^n}.$$

Supposing that $\alpha < \left\lfloor \frac{1}{160w} \sqrt{(1/\epsilon)\log_2(1/\delta)} \right\rfloor \cdot \log_2 k - \log_2 w - 1$, we show that Dec is not an $(n, \epsilon, \delta, \epsilon', k, \alpha)$-reduction. We first introduce some terminology to make things concise. Given $f_1, \ldots, f_k \in \{0,1\}^{2^n}$, a *dense model* for $D \in \mathcal{D}_n$ is an $M \in \mathcal{D}_n$ that is δ-dense in the uniform distribution $U \in \mathcal{D}_n$ and is such that for all $i \in [k]$, f_i does not ϵ-distinguish D and M. We say a function $f \in \{0,1\}^{2^n}$ *covers* $D \in \mathcal{D}_n$ if for every $R \in \mathcal{D}_n$ in which D is δ-dense, f ϵ'-distinguishes R and U.

Thus to show that Dec is not an $(n, \epsilon, \delta, \epsilon', k, \alpha)$-reduction, we need to find $f_1, \ldots, f_k \in \{0,1\}^{2^n}$ such that some D has no dense model but is not covered by $\mathrm{Dec}(f_1, \ldots, f_k, a)$ for any advice string $a \in \{0,1\}^\alpha$.

3.4.1 Distributions Without Dense Models

The following claim is our tool for finding distributions that have no dense models.

Claim 1. *For every $f_1, \ldots, f_k \in \{0,1\}^{2^n}$ and every $I \subseteq [k]$ of size $0 < s \leq 1/4\epsilon$ (for some s), if $0 < |X_I| \leq (2\delta/3)2^n$ then D_I has no dense model.*

Proof. We only consider the case when s is odd (essentially the same argument works when s is even). Suppose we pick $i \in I$ uniformly at random. Then for each $x \in X_I$ we have $\Pr_{i\sim I}[f_i(x) = 1] \geq 1/2 + 1/2s \geq 1/2 + 2\epsilon$, and for each $x \in \{0,1\}^n \backslash X_I$ we have $\Pr_{i\sim I}[f_i(x) = 1] \leq 1/2 - 1/2s \leq 1/2 - 2\epsilon$. Thus we have $\Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] \geq 1/2 + 2\epsilon$. Now consider an arbitrary M that is δ-dense in U. We have

$$\begin{aligned}
\Pr_{x\sim M,\, i\sim I}[f_i(x) = 1] &\leq \Pr_{x\sim M}[x \notin X_I] \cdot (1/2 - 2\epsilon) + \textstyle\sum_{x^* \in X_I} \Pr_{x\sim M}[x = x^*] \cdot \Pr_{i\sim I}[f_i(x^*) = 1] \\
&\leq \big(1 - |X_I|/\delta 2^n\big) \cdot (1/2 - 2\epsilon) + \textstyle\sum_{x^* \in X_I} (1/\delta 2^n) \cdot \Pr_{i\sim I}[f_i(x^*) = 1] \\
&= \big(1 - |X_I|/\delta 2^n\big) \cdot (1/2 - 2\epsilon) + (|X_I|/\delta 2^n) \cdot \Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] \\
&\leq (1/3) \cdot (1/2 - 2\epsilon) + (2/3) \cdot \Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] \\
&= \Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] - (1/3) \cdot \big(\Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] - (1/2 - 2\epsilon)\big) \\
&\leq \Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] - (1/3) \cdot 4\epsilon.
\end{aligned}$$

Here the second line follows because $\Pr_{i\sim I}[f_i(x^*) = 1] > 1/2 - 2\epsilon$ holds for all $x^* \in X_I$ and thus the whole expression only gets larger by shifting probability mass from $\{0,1\}^n \backslash X_I$ to X_I. Similarly, the fourth line follows because the third line is a convex combination of $1/2 - 2\epsilon$ and $\Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1]$, so the whole expression gets larger by shifting weight to the larger of the two.

Since $\Pr_{x\sim D_I,\, i\sim I}[f_i(x) = 1] - \Pr_{x\sim M,\, i\sim I}[f_i(x) = 1] > \epsilon$, there must exist an $i \in I$ such that $\mathrm{E}_{D_I}[f_i] - \mathrm{E}_M[f_i] > \epsilon$ and thus f_i ϵ-distinguishes D_I and M. Hence M is not a dense model for D_I. This finishes the proof of Claim 1. □

3.4.2 Distributions That Cannot Be Covered

We say a function $f \in \{0,1\}^{2^n}$ *positively covers* $D \in \mathcal{D}_n$ if for every $R \in \mathcal{D}_n$ in which D is δ-dense, $\mathrm{E}_R[f] - \mathrm{E}_U[f] > \epsilon'$ (note the absence of absolute value signs). Observe that if $f \in \{0,1\}^{2^n}$ covers D then either f or its complement positively covers D. This is because if there existed $R_1, R_2 \in \mathcal{D}_n$ in which D is δ-dense and such that $\mathrm{E}_{R_1}[f] < \mathrm{E}_U[f] < \mathrm{E}_{R_2}[f]$, then some convex combination R_3 of R_1 and R_2 would have $\mathrm{E}_{R_3}[f] = \mathrm{E}_U[f]$. However, D would be δ-dense in R_3 since the set of R in which D is δ-dense is convex, so f would not cover D.

Claim 2. *For every $f_1, \ldots, f_k \in \{0,1\}^{2^n}$, every $I_1, \ldots, I_\ell \subseteq [k]$ (for some ℓ), and every $J \subseteq [\ell]$ of size $t > 1$ (for some t), if $|X_{I_J}| \leq (\epsilon'/2)2^n$ and $|X_{I_j}| \geq (\delta/2 - \epsilon'/4)2^n$ for all $j \in J$ then there is no function that simultaneously positively covers D_{I_j} for all $j \in J$.*

Proof. Assume that $|X_{I_J}| \leq (\epsilon'/2)2^n$ and $|X_{I_j}| \geq (\delta/2 - \epsilon'/4)2^n$ for all $j \in J$. Consider an arbitrary $f \in \{0,1\}^{2^n}$ and let X be the set of $x \in \{0,1\}^n$ such that $f(x) = 1$. For $\tau \in \{0, 1, \ldots, t\}$ let $X^{(\tau)}$ be the set of $x \in \{0,1\}^n$ such that there are exactly τ values of $j \in J$ for which $x \in X_{I_j}$ (in other words, $(\mathrm{Maj}^t \circ f_{I_J})(x)$ has Hamming weight τ). Note that $X_{I_J} = \bigcup_{\tau=t'}^{t} X^{(\tau)}$ where $t' = \lceil t/2 \rceil$. Let

$\pi = \min_{j\in J}\left[\mathrm{E}_{D_{I_j}}[f]\right]$. Then for every $j \in J$ we have $|X \cap X_{I_j}| \geq \pi \cdot |X_{I_j}| \geq \pi \cdot (\delta/2 - \epsilon'/4)2^n$. We have

$$\begin{aligned}(t/2)\cdot\left(|X| + |X_{I_J}|\right) &\geq (t/2)\cdot|X \cap \overline{X_{I_J}}| + t\cdot|X \cap X_{I_J}| \\ &\geq \textstyle\sum_{\tau=0}^{t} \tau\cdot|X \cap X^{(\tau)}| \\ &= \textstyle\sum_{j\in J}|X \cap X_{I_j}| \\ &\geq t\cdot\pi\cdot(\delta/2 - \epsilon'/4)2^n\end{aligned}$$

which implies that

$$|X| \;\geq\; \pi\cdot(\delta - \epsilon'/2)2^n - |X_{I_J}| \;\geq\; \pi\delta 2^n - \epsilon' 2^n \;=\; (\pi - \epsilon'/\delta)\cdot\delta 2^n$$

since $\pi \leq 1$ and $|X_{I_J}| \leq (\epsilon'/2)2^n$. We might have $\pi - \epsilon'/\delta < 0$, but this is not problematic. Let M be a distribution δ-dense in U that maximizes $\mathrm{E}_M[f]$, and observe that

$$\mathrm{E}_M[f] \;=\; \min\left(|X|/\delta 2^n, 1\right) \;\geq\; \pi - \epsilon'/\delta.$$

We have $U = \delta M + (1-\delta)\widehat{M}$ for some $\widehat{M} \in \mathcal{D}_n$. Let $j \in J$ be such that $\mathrm{E}_{D_{I_j}}[f] = \pi$, and define the distribution $R = \delta D_{I_j} + (1-\delta)\widehat{M}$ so that D_{I_j} is δ-dense in R. Then we have

$$\mathrm{E}_R[f] \;=\; \delta\pi + (1-\delta)\,\mathrm{E}_{\widehat{M}}[f]$$

and

$$\mathrm{E}_U[f] \;=\; \delta\,\mathrm{E}_M[f] + (1-\delta)\,\mathrm{E}_{\widehat{M}}[f] \;\geq\; \delta\pi - \epsilon' + (1-\delta)\,\mathrm{E}_{\widehat{M}}[f] \;=\; \mathrm{E}_R[f] - \epsilon'$$

so f does not positively cover D_{I_j}. This finishes the proof of Claim 2. □

3.4.3 Setting the Parameters

Define $s = \lfloor 1/4\epsilon \rfloor$ and $t = w$ and $b = \left\lfloor \frac{1}{20t}\sqrt{(1/\epsilon)\log_2(1/\delta)} \right\rfloor$. By Lemma 4 there exists an (ℓ, k, s, b)-design $I_1, \ldots, I_\ell$ with $\ell = \lceil k^{b/8} \rceil$ (note that we do have $k \geq 16s^4$). Define p to be such that $\mathrm{Tail}(s, p, 1/2) = \delta/2$.

Claim 3. $\frac{1}{2}\sqrt{\epsilon\log_2(1/\delta)} \;\leq\; 1/2 - p \;\leq\; 2\sqrt{\epsilon\log_2(1/\delta)} \;\leq\; 1/4.$

Proof. The bound $2\sqrt{\epsilon\log_2(1/\delta)} \leq 1/4$ holds by our assumption $\epsilon \leq 1/64\log_2(1/\delta)$. To prove the upper bound on $1/2 - p$, define $p' = 1/2 - 2\sqrt{\epsilon\log_2(1/\delta)} \geq 1/4$. Then we have

$$\mathrm{Tail}\left(s, p', 1/2\right) \;\leq\; 2^{-\,\mathrm{RE}(1/2\|p')s} \;\leq\; 2^{-2(1/2-p')^2 s} \;=\; \delta^{8\epsilon s} \;\leq\; \delta^{8/5} \;<\; \delta/2$$

by Lemma 1 and Fact 1, and where the penultimate inequality uses $\epsilon \leq 1/20$. Thus $p \geq p'$. To prove the lower bound on $1/2 - p$, assume it does not hold. Then we would have the contradiction

$$\begin{aligned}\delta/2 &\geq \min\left(\tfrac{1}{256},\; \tfrac{1}{128\sqrt{s}(1/2-p)}\right)\cdot 2^{-\,\mathrm{RE}(1/2\|p)s} \\ &\geq \min\left(\tfrac{1}{256},\; \tfrac{1}{32\sqrt{\log_2(1/\delta)}}\right)\cdot 2^{-\,\mathrm{RE}(1/2\|p)s} \\ &\geq \delta^{1/2}\cdot 2^{-\,\mathrm{RE}(1/2\|p)s}\end{aligned}$$

$$\begin{aligned}
&\geq \delta^{1/2} \cdot 2^{-4(1/2-p)^2 s} \\
&\geq \delta^{1/2} \cdot 2^{-(1/2-p)^2/\epsilon} \\
&\geq \delta^{1/2} \cdot \delta^{1/4}
\end{aligned}$$

where the first line follows by Lemma 3 (note that we do have $s \geq 9$), the third line follows by[7] $\delta \leq 2^{-16}$, and the fourth line follows by Fact 1. This finishes the proof of Claim 3. □

3.4.4 The Majority of Majorities

We choose $f_1, \ldots, f_k$ randomly by setting $f_i(x) = 1$ with probability p independently for each $i \in [k]$ and each $x \in \{0,1\}^n$.

Claim 4. *For every $J \subseteq [\ell]$ of size t and every $x \in \{0,1\}^n$, we have $\Pr_{f_1,\ldots,f_k}[x \in X_{I_J}] \leq \epsilon'/4$.*

Proof. Define $t' = \lceil t/2 \rceil$. Note that if $(\mathrm{Maj} \circ \mathrm{Maj}^t \circ f_{I_J})(x) = 1$ then there exists a subset $J' \subseteq J$ of size t' such that $(\mathrm{And} \circ \mathrm{Maj}^{t'} \circ f_{I_{J'}})(x) = 1$. Thus we have

$$\Pr_{f_1,\ldots,f_k}\left[(\mathrm{Maj} \circ \mathrm{Maj}^t \circ f_{I_J})(x) = 1\right] \;\leq\; 2^t \cdot \max_{J' \subseteq J \,:\, |J'|=t'} \Pr_{f_1,\ldots,f_k}\left[(\mathrm{And} \circ \mathrm{Maj}^{t'} \circ f_{I_{J'}})(x) = 1\right].$$

Consider an arbitrary $J' \subseteq J$ of size t'. Define $m = \left|\bigcup_{j \in J'} I_j\right|$ and notice that since $I_1, \ldots, I_\ell$ is an (ℓ, k, s, b)-design, by inclusion-exclusion we have

$$t's - \tbinom{t'}{2} b \;\leq\; m \;\leq\; t's. \tag{1}$$

Define $s' = \lceil s/2 \rceil$ and $q = 1/2 - t'b/2s$. If $(\mathrm{And} \circ \mathrm{Maj}^{t'} \circ f_{I_{J'}})(x) = 1$ then for each $j \in J'$ we have $\sum_{i \in I_j} f_i(x) \geq s'$ and so by inclusion-exclusion we have

$$\textstyle\sum_{i \in \bigcup_{j \in J'} I_j} f_i(x) \;\geq\; \left(\sum_{j \in J'} \sum_{i \in I_j} f_i(x)\right) - \binom{t'}{2} b \;\geq\; t's' - \binom{t'}{2} b \;\geq\; qt's \;\geq\; qm.$$

It follows that

$$\begin{aligned}
\Pr_{f_1,\ldots,f_k}\left[(\mathrm{And} \circ \mathrm{Maj}^{t'} \circ f_{I_{J'}})(x) = 1\right] \;&\leq\; \Pr_{f_1,\ldots,f_k}\left[\textstyle\sum_{i \in \bigcup_{j \in J'} I_j} f_i(x) \geq qm\right] \\
&=\; \mathrm{Tail}(m, p, q) \\
&\leq\; 2^{-\mathrm{RE}(q\|p)m} \\
&=\; \left(2^{-\mathrm{RE}(1/2\|p)s}\right)^{(m/s)\cdot(\mathrm{RE}(q\|p)/\,\mathrm{RE}(1/2\|p))} \\
&\leq\; \left(\delta^{1/10}\right)^{(m/s)\cdot(\mathrm{RE}(q\|p)/\,\mathrm{RE}(1/2\|p))}
\end{aligned}$$

where the third line follows by Lemma 1 and the fifth line follows by nonnegativity of RE and

$$2^{-\mathrm{RE}(1/2\|p)s} \;\leq\; 2^{-2(1/2-p)^2 s} \;\leq\; \delta^{\epsilon s/2} \;\leq\; \delta^{1/10}$$

which holds by Fact 1, Claim 3, and $\epsilon \leq 1/20$. We have

$$m/s \;\geq\; t' - (t')^2 b/2s \;\geq\; t'/2 \;\geq\; t/4 \tag{2}$$

[7] The existence of w in the statement of Theorem 2 actually implies $\delta \leq 2^{-160}$.

by (1) and $b \le s/t'$ (which can be shown using the final inequality in Claim 3). We also have $t'b/2s \le \frac{1}{8}\sqrt{\epsilon \log_2(1/\delta)}$ and thus $q - p \ge \frac{3}{4}(1/2 - p)$ by Claim 3. Hence by Fact 1 we have

$$\mathrm{RE}(q\|p)/\,\mathrm{RE}(1/2\|p) \;\ge\; \tfrac{(q-p)^2}{2(1/2-p)^2} \;\ge\; \tfrac{(\frac{3}{4}(1/2-p))^2}{2(1/2-p)^2} \;\ge\; 1/4. \tag{3}$$

Using (2) and (3) we get

$$\Pr_{f_1,\ldots,f_k}\left[(\mathrm{And} \circ \mathrm{Maj}^{t'} \circ f_{I_{J'}})(x) = 1\right] \;\le\; \left(\delta^{1/10}\right)^{(t/4)\cdot(1/4)} \;=\; \delta^{t/160}.$$

We conclude that

$$\Pr_{f_1,\ldots,f_k}[x \in X_{I_J}] \;\le\; 2^t \cdot \delta^{t/160} \;\le\; \epsilon'/4.$$

This finishes the proof of Claim 4. □

3.4.5 Putting It All Together

For every $j \in [\ell]$ and every $x \in \{0,1\}^n$, we have $\Pr_{f_1,\ldots,f_k}[x \in X_{I_j}] = \mathrm{Tail}(s, p, 1/2) = \delta/2$. Standard relative-error forms of the Chernoff bound give

$$\begin{aligned}
\Pr_{f_1,\ldots,f_k}\left[|X_{I_j}| < (\delta/2 - \epsilon'/4)2^n\right] &\le e^{-2^n(\epsilon')^2/16\delta} \\
\Pr_{f_1,\ldots,f_k}\left[|X_{I_j}| > (2\delta/3)2^n\right] &\le e^{-2^n\delta/54} \\
\Pr_{f_1,\ldots,f_k}\left[|X_{I_J}| > (\epsilon'/2)2^n\right] &\le e^{-2^n\epsilon'/12}
\end{aligned}$$

where the latter holds for each $J \subseteq [\ell]$ of size t, using Claim 4. Thus by a union bound we have

$$\begin{aligned}
&\Pr_{f_1,\ldots,f_k}\left[\begin{array}{l}(\delta/2 - \epsilon'/4)2^n \le |X_{I_j}| \le (2\delta/3)2^n \text{ for all } j \in [\ell] \text{ and} \\ |X_{I_J}| \le (\epsilon'/2)2^n \text{ for all } J \subseteq [\ell] \text{ of size } t\end{array}\right] \\
&\ge\; 1 - \ell \cdot e^{-2^n(\epsilon')^2/16\delta} - \ell \cdot e^{-2^n\delta/54} - \tbinom{\ell}{t} \cdot e^{-2^n\epsilon'/12} \\
&>\; 0
\end{aligned}$$

since $2^n \ge \frac{t \log_2 k}{\epsilon\delta^2(\epsilon')^2}$. Fix a choice of $f_1, \ldots, f_k$ such that the above event occurs.

For every $J^* \subseteq [\ell]$ of size $2t-1$, there is no $a \in \{0,1\}^\alpha$ such that $\mathrm{Dec}(f_1, \ldots, f_k, a)$ simultaneously covers D_{I_j} for all $j \in J^*$, because otherwise for some $J \subseteq J^*$ of size t, either $\mathrm{Dec}(f_1, \ldots, f_k, a)$ or its complement would simultaneously positively cover D_{I_j} for all $j \in J$, which would contradict Claim 2.

Thus for each $a \in \{0,1\}^\alpha$, the number of $j \in [\ell]$ such that D_{I_j} is covered by $\mathrm{Dec}(f_1, \ldots, f_k, a)$ is at most $2t - 2$. This implies that the number of $j \in [\ell]$ for which there exists an $a \in \{0,1\}^\alpha$ such that $\mathrm{Dec}(f_1, \ldots, f_k, a)$ covers D_{I_j} is at most $2^\alpha \cdot (2t-2) < k^{b/8} \le \ell$ since $\alpha \le (b/8)\log_2 k - \log_2 t - 1$. Thus there exists a $j \in [\ell]$ such that D_{I_j} is not covered by $\mathrm{Dec}(f_1, \ldots, f_k, a)$ for any $a \in \{0,1\}^\alpha$. By Claim 1, D_{I_j} has no dense model, so Dec is not an $(n, \epsilon, \delta, \epsilon', k, \alpha)$-reduction. This finishes the proof of Theorem 2.

4 Open Problems

The first open problem is to quantitatively improve our lower bound (or give an improved upper bound) on the advice complexity of black-box reductions for the Dense Model Theorem. It is also open to give any nontrivial lower bound on the computational complexity.

Another open problem is to give any lower bound (or improved upper bound) on the complexity of black-box reductions for the original formulation of the Dense Model Theorem, where the advice is allowed to depend on R.

Gentry and Wichs [GW11a] proved a lemma that is similar in spirit to the Dense Model Theorem, but where the notion of D_1 being δ-dense in D_2 is replaced by D_1 being an *extension* of D_2, that is, D_1 is a distribution on $(n+\ell)$-bit strings whose marginal on the first n bits is D_2. It would be interesting to give lower bounds (or improved upper bounds) on the computational complexity, query complexity, or advice complexity of black-box reductions for this lemma. Also, the proof of Gentry and Wichs goes through the Minimax Theorem (as in [Imp95, RTTV08]); it would be interesting to give a boosting-style proof (as in [Imp95, KS03, BHK09, TTV09, Zha11]).

Acknowledgments

I thank Anand Bhaskar, Siu Man Chan, and Siu On Chan for helpful discussions, and anonymous reviewers for useful comments.

References

[AS11] Sergei Artemenko and Ronen Shaltiel. Lower bounds on the query complexity of non-uniform and adaptive reductions showing hardness amplification. In *Proceedings of the 15th International Workshop on Randomization and Computation*, pages 377–388, 2011.

[BHK09] Boaz Barak, Moritz Hardt, and Satyen Kale. The uniform hardcore lemma via approximate Bregman projections. In *Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1200, 2009.

[Bli86] Volodia Blinovsky. Bounds for codes in the case of list decoding of finite volume. *Problems of Information Transmission*, 22(1):7–19, 1986.

[BSW03] Boaz Barak, Ronen Shaltiel, and Avi Wigderson. Computational analogues of entropy. In *Proceedings of the 7th International Workshop on Randomization and Computation*, pages 200–215, 2003.

[DP08] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science*, pages 293–302, 2008.

[EFF85] Paul Erdős, Péter Frankl, and Zoltán Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51(1-2):79–89, 1985.

[GNW11] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao's XOR-Lemma. *Studies in Complexity and Cryptography*, pages 273–301, 2011.

[Gow10] Timothy Gowers. Decompositions, approximate structure, transference, and the Hahn-Banach Theorem. *Bulletin of the London Mathematical Society*, 42(4):573–606, 2010.

[GR08] Dan Gutfreund and Guy Rothblum. The complexity of local list decoding. In *Proceedings of the 12th International Workshop on Randomization and Computation*, pages 455–468, 2008.

[GT08] Ben Green and Terence Tao. The primes contain arbitrarily long arithmetic progressions. *Annals of Mathematics*, 167(2):481–547, 2008.

[Gur03] Venkatesan Guruswami. List decoding from erasures: Bounds and code constructions. *IEEE Transactions on Information Theory*, 49(11):2826–2833, 2003.

[GV10] Venkatesan Guruswami and Salil Vadhan. A lower bound on list size for list decoding. *IEEE Transactions on Information Theory*, 56(11):5681–5688, 2010.

[GW11a] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 99–108, 2011.

[GW11b] Timothy Gowers and Julia Wolf. Linear forms and higher-degree uniformity for functions on $\mathbb{F}_p^n$. *Geometric and Functional Analysis*, 21(1):36–69, 2011.

[GW12] Timothy Gowers and Julia Wolf. Linear forms and quadratic uniformity for functions on $\mathbb{F}_p^n$. *Mathematika*, 57(2):215–237, 2012.

[IJK09] Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Approximate list-decoding of direct product codes and uniform hardness amplification. *SIAM Journal on Computing*, 39(2):564–605, 2009.

[IJKW10] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. Uniform direct product theorems: Simplified, optimized, and derandomized. *SIAM Journal on Computing*, 39(4):1637–1665, 2010.

[Imp95] Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 538–545, 1995.

[KS03] Adam Klivans and Rocco Servedio. Boosting and hard-core sets. *Machine Learning*, 53(3):217–238, 2003.

[Lev87] Leonid Levin. One-way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987.

[LTW08] Chi-Jen Lu, Shi-Chun Tsai, and Hsin-Lung Wu. On the complexity of hardness amplification. *IEEE Transactions on Information Theory*, 54(10):4575–4586, 2008.

[LTW11] Chi-Jen Lu, Shi-Chun Tsai, and Hsin-Lung Wu. Complexity of hard-core set proofs. *Computational Complexity*, 20(1):145–171, 2011.

[MPRV09] Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil Vadhan. Computational differential privacy. In *Proceedings of the 29th International Cryptology Conference*, pages 126–142, 2009.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs. randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994.

[RTTV08] Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Dense subsets of pseudorandom sets. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science*, pages 76–85, 2008.

[SV10] Ronen Shaltiel and Emanuele Viola. Hardness amplification proofs require majority. *SIAM Journal on Computing*, 39(7):3122–3154, 2010.

[TTV09] Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *Proceedings of the 24th IEEE Conference on Computational Complexity*, pages 126–136, 2009.

[TZ08] Terence Tao and Tamar Ziegler. The primes contain arbitrarily long polynomial progressions. *Acta Mathematica*, 201:213–305, 2008.

[Wat11] Thomas Watson. Query complexity in errorless hardness amplification. In *Proceedings of the 15th International Workshop on Randomization and Computation*, pages 688–699, 2011.

[Zha11] Jiapeng Zhang. On the query complexity for showing dense model. Technical Report TR11-038, Electronic Colloquium on Computational Complexity, 2011.