

Sharp bounds for population recovery

Anindya De
Northwestern University
de.anindya@gmail.com

Ryan O’Donnell*
Carnegie Mellon University
odonnell@cs.cmu.edu

Rocco A. Servedio†
Columbia University
rocco@cs.columbia.edu

March 3, 2017

Abstract

The *population recovery problem* is a basic problem in noisy unsupervised learning that has attracted significant research attention in recent years [WY12, DRWY12, MS13, BIMP13, LZ15, DST16]. A number of different variants of this problem have been studied, often under assumptions on the unknown distribution (such as that it has restricted support size). In this work we study the sample complexity and algorithmic complexity of the most general version of the problem, under both bit-flip noise and erasure noise model. We give essentially matching upper and lower sample complexity bounds for both noise models, and efficient algorithms matching these sample complexity bounds up to polynomial factors.

1 Introduction

1.1 The erasure noise and bit-flip noise population recovery problems

The *noisy population recovery (NPR)* problem is to learn an unknown probability distribution \mathcal{D} on $\{0, 1\}^n$, under ν -noise, to ℓ_∞ -accuracy ϵ .¹ In this problem the learner gets access to independent *samples* \mathbf{y} , each distributed as follows: First $\mathbf{x} \sim \mathcal{D}$, and then $\mathbf{y} \sim \text{Noise}_\nu(\mathbf{x})$, where $\text{Noise}_\nu(\cdot)$ denotes either the application of bit-flip noise or erasure noise (described below). The learner’s task is to output an estimate $\hat{\mathcal{D}}$ of \mathcal{D} satisfying $\|\hat{\mathcal{D}} - \mathcal{D}\|_\infty \leq \epsilon$ (with high probability). For the sake of a compact representation, we assume the learner only outputs the nonzero values of $\hat{\mathcal{D}}$; this means that a successful learner need only output $O(1/\epsilon)$ nonzero values. We are interested in minimizing both the *sample complexity* and the *running time* of learning algorithms.

A simpler variation of the NPR problem is the *estimation* task. Here the algorithm doesn’t need to output a complete $\hat{\mathcal{D}}$; it only needs to output an ϵ -accurate estimate of $\mathcal{D}(u)$ for a given input $u \in \{0, 1\}^n$. Certainly the estimation task is no harder than full NPR; conversely, it is known and not hard (see Section 2.1) that given the ability to do estimation, one can do full NPR with just a $\text{poly}(n, 1/\epsilon)$ factor slowdown. Hence we mainly focus on estimation in this paper.

As mentioned above, we consider two different models of noise. Each involves a parameter $0 < \nu < 1$; smaller values of ν correspond to more noise, so ν may be better thought of as a “correlation” parameter.

*Supported by NSF grant CCF-1618679.

†Supported by NSF grants CCF-1420349 and CCF-1563155.

¹With high probability. Because we are not concerned with logarithmic factors in our time/sample complexity, we will for simplicity omit discussion of the standard tricks (independent repetition, taking the median of estimators) used to boost success probabilities. We will also always assume, without loss of generality, that ϵ is at most some sufficiently small absolute constant.

Erasure noise. For $x \in \{0, 1\}^n$ we define $\text{Erase}_{1-\nu}(x)$ to be the distribution on $\{0, 1, ?\}^n$ given by independently replacing each coordinate of x with the symbol ‘?’ with probability $1 - \nu$. Thus ν is the retention probability for each coordinate.

Bit-flip noise. For $x \in \{0, 1\}^n$ we define $\text{Flip}_{\frac{1-\nu}{2}}(x)$ to be the distribution on $\{0, 1\}^n$ given by independently flipping each coordinate of x with probability $\frac{1-\nu}{2}$. Equivalently, each coordinate of x is retained with probability ν (as in erasure noise), and is otherwise replaced with a uniformly random bit. This is also the model of noise associated to the so-called “Bonami–Beckner noise operator” T_ν .

1.2 Our results

For the bit-flip noise population recovery problem, our main result is a lower bound on the sample complexity of estimation, as well as a full NPR algorithm whose running time (hence also sample complexity) matches it up to polynomial factors:

Theorem 1.1. *Let $\epsilon > 0$ be sufficiently small and let $n \in \mathbb{N}$. Then any estimation algorithm for NPR with bit-flip noise must use at least the following number of samples:*

$$\begin{cases} \exp\left(\Theta\left(n^{1/3} \cdot \ln^{2/3}(1/\epsilon)/\nu^{2/3}\right)\right) & \text{if } \frac{\ln(1/\epsilon)}{n} \leq \nu \leq 1/2, \\ \exp\left(\Theta\left(n^{1/3} \cdot \ln^{2/3}(1/\epsilon) \cdot (1-\nu)^{1/3}\right)\right) & \text{if } 1/2 \leq \nu \leq 1 - \frac{\ln(1/\epsilon)}{n}. \end{cases}$$

Furthermore, there is an algorithm for the full NPR problem with bit-flip noise having running time and samples equal to the above times $\text{poly}(n, 1/\epsilon)$.

Prior to this work and the very recent and independent work of [PSW17], no nontrivial upper or lower bounds were known even for the sample complexity of the general bit-flip noise population recovery problem. (See [WY12, LZ15, DST16] for earlier works that gave upper bounds and algorithms under the additional assumption that the unknown distribution \mathcal{D} is guaranteed to be supported on at most k strings.)

For the erasure noise population recovery problem, our main results are also essentially matching upper and lower bounds on sample complexity and an algorithm whose running time is polynomial in n and the sample complexity:

Theorem 1.2. *Let $\epsilon > 0$ be sufficiently small and let $n \in \mathbb{N}$.*

- Assume that $\sqrt{16 \ln(1/\epsilon)/n} \leq \nu \leq 1/160$. Then any estimation algorithm for NPR with erasure noise must use at least $1/\epsilon^{\Omega(1/\nu)}$ samples.
- There is an algorithm for the full NPR problem with erasure noise using time and samples at most $\text{poly}(n, 1/\epsilon^{1/\nu})$.

For this problem, in earlier work [MS13] gave an algorithm with sample complexity and running time $(n/\epsilon)^{O(\log(1/\nu)/\nu)}$.

Finally, we note that in very recent and independent work, [PSW17] have obtained very similar results to Theorems 1.1 and 1.2 for the population recovery problem.

1.3 Our techniques

Our approach is similar in spirit to, and shares some technical similarities with, the recent work of [DOS16, NP16] on the trace reconstruction problem. We take an analytic view on the combinatorial process defined by the bit-flip and erasure noise operators, and convert the sample complexity questions for these population recovery problems to questions about the extrema of real-coefficient polynomials satisfying certain conditions on various circles in the complex plane; we then obtain our sample complexity bounds by analyzing these extremal polynomial questions. The main algorithmic ingredient in our results is linear programming.

2 Preliminaries

2.1 Well-known preliminary reductions

Estimation, enumeration, and recovery. Variants of the NPR problem with relaxed goals have been studied in the literature. One is the aforementioned *estimation* problem. Another (complementary) variant is called *enumeration*: in the enumeration problem, the learning algorithm is only required to output a list of strings x_1, \dots, x_m that is guaranteed (with high probability) to include all strings that have probability at least ϵ under \mathcal{D} ; such strings are sometimes referred to as “heavy hitters.” Batman et al. [BIMP13] gave a range of results for the enumeration problem.

It is easy to see that a solution to the estimation problem can be efficiently bootstrapped to full NPR given the ability to solve the enumeration problem (simply run estimation, with a sufficiently boosted success probability, on each of the m strings in the list obtained from enumeration). In turn, it is also well known that an estimation algorithm can be efficiently transformed into an enumeration algorithm via a “branch-and-prune” approach. Roughly speaking, such an approach maintains a not-too-large (size at most $O(1/\epsilon)$) set of i -bit prefixes that is known to contain all the “heavy hitters”; to construct the set of $(i + 1)$ -bit prefixes, the approach first “branches” to extend each i -bit prefix x to both $x0$ and $x1$, and then “prunes” any element of $\{x0, x1\}$ that is determined, using the estimation procedure, not to be a heavy hitter. (Note that since only heavy hitters are maintained it will again be the case that the set of $(i + 1)$ -bit prefixes has size at most $O(1/\epsilon)$.) As [BIMP13] observe, an early example of such a branch-and-prune routine that performs enumeration given an oracle for estimation is the Goldreich–Levin algorithm [GL89] for list-decoding the Hadamard code. Both Dvir et al. [DRWY12] and Batman et al. [BIMP13] give fairly detailed analyses of the above-described reduction from enumeration to estimation; we omit the details here and refer the interested reader to Section 6.1 of [DRWY12] and Section 2 of [BIMP13] respectively.

Summarizing the reductions discussed above, we have that NPR is (up to polynomial factors) no harder than the estimation problem, and it is also clearly no easier than estimation (since estimation is a subproblem of general NPR). Thus in the rest of this paper we restrict our attention to the estimation problem.

Symmetrization. We further recall some well-known tricks that have been used in past papers on NPR. First, in the estimation problem, we may assume without loss of generality that the string u whose probability is to be estimated is $u = (0, \dots, 0)$. This is because the learner can easily convert samples from \mathcal{D} to samples from “ $\mathcal{D} \oplus u$ ”.

Next, for the problem of estimating $\mathcal{D}(0, \dots, 0)$, we may assume without loss of generality that \mathcal{D} is *symmetric*, meaning that it gives equal probability mass to all strings at the same Hamming weight. In other words, \mathcal{D} is effectively given by a probability distribution \mathcal{D}^{sym} on $[0..n]$, with

$\mathcal{D}(x) = \mathcal{D}^{\text{sym}}(|x|)/\binom{n}{|x|}$. On one hand, if $\mathcal{D}(0, \dots, 0)$ can be estimated in the general case, it can certainly be estimated in the symmetric case. On the other hand, given a general distribution \mathcal{D} , the learner can randomly permute the coordinates of each sample, effectively obtaining access to samples from a symmetric distribution \mathcal{D}^{sym} , with $\mathcal{D}^{\text{sym}}(0, \dots, 0) = \mathcal{D}(0, \dots, 0)$. Thus it suffices for the learner to be able to estimate in the symmetric case.

In this symmetric case, we will write the unknown \mathcal{D}^{sym} more simply as a probability (row) vector $[p_0 \ p_1 \ \dots \ p_n]$. Although the learner observes full strings, it may as well only consider the Hamming weights of the strings it receives. Thus we may think of it as obtaining samples from the probability (row) vector $[q_0 \ q_1 \ \dots \ q_n]$, where

$$q = pA, \quad A_{ij} = \Pr[\text{a weight } i \text{ string becomes a weight } j \text{ string under } \nu \text{ noise}]. \quad (1)$$

It is not hard to write down the entries of A in either noise model. We remark that, after symmetrization, the bit-flip model becomes equivalent to running the well-known *Ehrenfest urn model* for continuous time tn , where $e^{-t} = \nu$. It is easy to write down the known generating function for that model:

Proposition 2.1 ([Sie47, BH51]). *For A associated to the Flip $_{\frac{1-\nu}{2}}$ noise model, and z an indeterminate,*

$$\sum_{j=0}^n A_{ij} z^j = \left(\frac{1-\nu}{2} + \frac{1+\nu}{2} z \right)^i \left(\frac{1+\nu}{2} + \frac{1-\nu}{2} z \right)^{n-i}.$$

For the erasure model, the generating function is even simpler. The following is easily verified:

Proposition 2.2. *For A associated to the Erase $_{1-\nu}$ noise model, and z an indeterminate,*

$$\sum_{j=0}^n A_{ij} z^j = ((1-\nu) + \nu z)^i.$$

To recap, in the estimation problem the learner's task is to estimate p_0 to accuracy ϵ , given samples from q . We recall the well-known fact that, by taking the empirical distribution of $O(n/\delta^2)$ samples, the learner may obtain an estimate \hat{q} of q satisfying $\|\hat{q} - q\|_1 \leq \delta$ (with high probability). Although $q = pA$, as noted in previous works one unfortunately cannot effectively estimate p_0 simply as the first coordinate of $\hat{q}A^{-1}$, because A is very poorly conditioned. Instead one needs a more sophisticated approach.

3 Reduction to an analytic problem

It is not hard to characterize the optimal sample complexity for the estimation problem. Define

$$\eta(\epsilon, \nu) = \min_{\substack{\text{probability vectors } p, p' \\ |p_0 - p'_0| > 2\epsilon}} \|pA - p'A\|_1$$

(where the parameter ν implicitly appears within A). If two probability vectors p and p' have $|p_0 - p'_0| > 2\epsilon$, then a successful estimation algorithm must be able to distinguish the two cases. But if $q = pA$, $q' = p'A$ are close, in the sense that $\|q - q'\|_1 \leq \delta$, then a learning algorithm will need $\Omega(1/\delta)$ samples to distinguish them with high probability. We conclude:

Proposition 3.1. *The sample complexity of any population recovery algorithm — indeed, any estimation algorithm — is $\Omega(1/\eta(\epsilon, \nu))$.*

On the other hand, suppose the lower bound $\eta(\epsilon, \nu) \geq \delta$ holds. Consider an estimation algorithm that first produces an empirical estimate \hat{q} with $\|\hat{q} - q\|_1 < \delta$ using $O(n/\delta^2)$ samples, and then exactly solves the following optimization problem using linear programming:

$$\min_{\text{probability vectors } p'} \|\hat{q} - p'A\|_1.$$

(This can be efficiently written as an LP with $O(n)$ variables and constraints and with rational numbers of $\text{poly}(n)$ bit-complexity.²) We claim that any optimal solution p' will have $|p_0 - p'_0| \leq 2\epsilon$. Otherwise, by definition $\|pA - p'A\|_1 > \eta(\epsilon, \nu) \geq \delta$; but $\|\hat{q} - pA\|_1 < \delta$, a contradiction. Thus we get an efficient solution to the estimation problem (except with accuracy only 2ϵ). In conclusion, we have established the following:

Proposition 3.2. *The estimation problem can be solved with $\text{poly}(n, 1/\eta(\epsilon/2, \nu))$ time and samples.*

Thus we see that, up to polynomial factors, both the sample complexity and runtime complexity of the estimation problem is effectively controlled by the parameter $\eta(\epsilon, \nu)$.

We now further simplify the definition of $\eta(\epsilon, \nu)$, similar to what was done in [DOS16]. The difference of two probability vectors is precisely any vector in the set

$$\Delta = \{[c_0 \ c_1 \ \cdots \ c_n] : \sum_i c_i = 0, \sum_i |c_i| \leq 2\}.$$

Thus we have that

$$\eta(\epsilon, \nu) = \min_{\substack{c \in \Delta \\ c_0 > 2\epsilon}} \|cA\|_1.$$

Furthermore, elementary complex analysis (see e.g. Proposition 3.5 of [DOS16]) shows that, for complex z and column vector $\mathbf{j} = (1, z, z^2, \dots, z^n)$,

$$\max_{|z|=1} |cA\mathbf{j}| \leq \|cA\|_1 \leq \sqrt{n+1} \cdot \max_{|z|=1} |cA\mathbf{j}|.$$

Note also that $cA\mathbf{j}$ is a polynomial in z that is easily calculated from the generating function of the noise process (see Propositions 2.1, 2.2). We obtain:

Theorem 3.3. *Up to a factor of $\sqrt{n+1}$ on the outside, and up to a factor of 2 on ϵ , we have that*

$$\eta(\epsilon, \nu) \asymp \min_{\substack{c \in \Delta \\ c_0 > 2\epsilon}} \begin{cases} \max_{|z|=1} |F_c(z)| & \text{in the Flip}_{\frac{1-\nu}{2}} \text{ noise model,} \\ \max_{|z|=1} |E_c(z)| & \text{in the Erase}_{1-\nu} \text{ noise model,} \end{cases}$$

where

$$F_c(z) = \sum_{i=0}^n c_i \left(\frac{1-\nu}{2} + \frac{1+\nu}{2} z \right)^i \left(\frac{1+\nu}{2} + \frac{1-\nu}{2} z \right)^{n-i}, \quad (2)$$

$$E_c(z) = \sum_{i=0}^n c_i ((1-\nu) + \nu z)^i. \quad (3)$$

²For simplicity in this paper we assume that ϵ and ν are rational quantities of $\text{poly}(n)$ bits known to the learning algorithm.

Given $c \in \Delta$ with $c_0 > 2\epsilon$, define the following polynomial (with real coefficients and a complex parameter):

$$Q_c(v) = \sum_{i=0}^n c_i v^i.$$

Thus the assumptions on c are equivalent to $Q_c(0) > 2\epsilon$, $Q_c(1) = 0$, and $L(Q_c) \leq 2$, where $L(Q_c)$ is the *length* of Q_c ; i.e., the sum of the absolute values of its coefficients.

In analyzing E_c above, we use that $E_c(z) = \sum_{i=0}^n c_i u^i$, where $u = (1 - \nu) + \nu z$. As z traces out the unit circle $|z| = 1$, the parameter u traces out the circle $\partial D_\nu(1 - \nu)$ of radius ν centered at the real value $1 - \nu$. Thus

$$\max_{|z|=1} |E_c(z)| = \max_{u \in \partial D_\nu(1-\nu)} |Q_c(u)|,$$

where

$$Q_c(v) = \sum_{i=0}^n c_i v^i.$$

In analyzing F_c above, we use that

$$F_c(z) = \left(\frac{\nu}{\frac{1+\nu}{2} - \frac{1-\nu}{2} w} \right)^n \sum_{i=0}^n c_i w^i, \quad \text{where } w = \frac{\frac{1-\nu}{2} + \frac{1+\nu}{2} z}{\frac{1+\nu}{2} + \frac{1-\nu}{2} z}. \quad (4)$$

The parameter w (being a Möbius transformation of z) traces out the unit circle as z does, and for $w = e^{i\theta}$ it is not hard to compute that

$$\left| \frac{\nu}{\frac{1+\nu}{2} - \frac{1-\nu}{2} w} \right|^2 = \frac{2\nu^2}{(1 - \cos \theta) + (1 + \cos \theta)\nu^2} = \frac{1}{1 + \frac{(1-\nu^2)\sin^2(\theta/2)}{\nu^2}}. \quad (5)$$

Thus

$$\max_{|z|=1} |F_c(z)| = \max_{-\pi < \theta \leq \pi} \left(\frac{1}{1 + \frac{(1-\nu^2)\sin^2(\theta/2)}{\nu^2}} \right)^{n/2} \cdot |Q_c(e^{i\theta})|. \quad (6)$$

We finally conclude:

Corollary 3.4. *Up to a factor of $\sqrt{n+1}$ on the outside, and up to a factor of 2 on ϵ , we have that*

$$\eta(\epsilon, \nu) \asymp \min_Q \begin{cases} \max_{-\pi < \theta \leq \pi} \left(\frac{1}{1 + \frac{(1-\nu^2)\sin^2(\theta/2)}{\nu^2}} \right)^{n/2} \cdot |Q(e^{i\theta})| & \text{in the Flip}_{\frac{1-\nu}{2}} \text{ noise model,} \\ \max_{u \in \partial D_\nu(1-\nu)} |Q(u)|, & \text{in the Erase}_{1-\nu} \text{ noise model,} \end{cases}$$

where the minimum is over real-coefficient polynomials Q of degree at most n satisfying $Q(0) > 2\epsilon$, $Q(1) = 0$, and $L(Q) \leq 2$.

Combining Propositions 3.1 and 3.2 with Corollary 3.4, we see that Theorems 1.1 and 1.2 follow from giving bounds on the two quantities specified in Corollary 3.4 (or in Theorem 3.3). We give such bounds in the following sections.

4 Circle bounds for erasure noise

4.1 A lower bound on $\eta(\epsilon, \nu)$ for erasure noise

Notice that $L(Q) \leq 2$ implies that $|Q(u)| \leq 2$ for all $|u| = 1$. We have the following:

Theorem 4.1. *Let Q be a complex polynomial with $|Q(0)| \geq 2\epsilon$ and $|Q(u)| \leq 2$ for $|u| = 1$. Then for $0 < \nu < 1/2$ we have $\max_{u \in \partial D_\nu(1-\nu)} |Q(u)| \geq 2\epsilon^{\frac{1-\nu}{\nu}}$. (For $1/2 \leq \nu \leq 1$ we immediately get a lower bound of 2ϵ , by the Maximum Modulus Principle.)*

Proof. Let U be the unit circle, let O be the circle of radius $1/2$ centered at $1/2$, which lies inside U , and let $C = \partial D_\nu(1-\nu)$, which lies inside O . The Möbius transformation $A(u) = 1/(1-u)$ takes these circles to vertical lines U' , O' , and C' with real parts $1/2$, 1 , and $1/2\nu$, respectively. Defining the function $f(z) = Q(A^{-1}(z))$, we have that f is bounded on the strip defined by U' and C' , and we have that $\sup_{y \in U'} |f(y)| \leq 2$, $\sup_{y \in O'} |f(y)| \geq 2\epsilon$. Writing M for the maximum modulus of f on C' , the Hadamard Three-Lines Theorem implies that

$$2^{\frac{1-2\nu}{1-\nu}} M^{\frac{\nu}{1-\nu}} \geq 2\epsilon,$$

which completes the proof after rearrangement. \square

4.2 An upper bound on $\eta(\epsilon, \nu)$ for erasure noise

In this section, we will prove the following theorem:

Theorem 4.2. *There is an absolute constant $\tau > 0$ such that for every $\nu \leq 1/10$, $0 < \epsilon < \tau$ and $\ln(1/\epsilon)/\nu^2 \leq n$, there exists a vector $c \in \Delta$ with $c_0 > 2\epsilon$ such that the polynomial $Q_c(v) = \sum_{i=0}^n c_i v^i$ satisfies*

$$\sup_{v \in \partial D_{\nu/16}(1-\nu/16)} |Q_c(v)| = \epsilon^{-\Omega(1/\nu)}.$$

In order to prove this theorem, we will collect a few facts at the beginning. Given $a, r > 0$, define the set $B_{a,r}$ as

$$B_{a,r} = \{(1-8a) + 4a(z+z^{-1}) : z \in \partial D_r(0)\}.$$

We now make a few observations about the set $B_{a,r}$ as r varies. In particular, we have the following fact:

Fact 4.3. *For $r \in \{1, 2, 4\}$, the sets $B_{a,r}$ are as follows:*

- For $r = 1$, the set $B_{a,r}$ is the line segment joining 1 and $1 - 16a$.
- For $r = 2$, the set $B_{a,r}$ is the ellipse centered at $1 - 8a$ with major axis $[1 - 8a - 10a, 1 - 8a + 10a]$ and minor axis $[1 - 8a + 6i, 1 - 8a - 6i]$.
- For $r = 4$, the set $B_{a,r}$ is the ellipse centered at $1 - 8a$ with major axis is $[1 - 8a - 17a, 1 - 8a + 17a]$ and minor axis is $[1 - 8a + 15i, 1 - 8a - 15i]$.

It is quite easy to observe that the circle $D_{4a}(1-4a)$ is contained in $B_{a,2}$. By Hadamard's three circle theorem, any holomorphic function f satisfies

$$\sup_{z \in D_{4a}(1-4a)} |f(z)| \leq \sup_{z \in B_{a,2}} |f(z)| \leq \sqrt{\sup_{z \in B_{a,1}} |f(z)|} \cdot \sqrt{\sup_{z \in B_{a,4}} |f(z)|}. \quad (7)$$

Consequently, we have the following corollary.

Corollary 4.4. *Let $c \in \Delta$ and $Q_c(v) = \sum_{i=0}^n c_i v^i$. Then,*

$$\sup_{z \in D_{4a}(1-4a)} |Q_c(z)| \leq \sqrt{\sup_{z \in B_{a,1}} |Q_c(z)|} \cdot 2\sqrt{\exp(9an)}.$$

Proof. We apply (7) to the function Q_c and then observe that

$$\sup_{z \in B_{a,4}} |Q_c(z)| \leq \sup_{z \in B_{a,4}} |z|^n \cdot \left(\sum_{j=0}^n |c_j| \right) \leq 2 \cdot (1+9a)^n \leq 2 \cdot \exp(9an),$$

which concludes the proof. \square

We next recall the following theorem from [Erd16]:

Theorem 4.5 (Lemma 3.3 of [Erd16]). *For any $L \in [0, 1/17)$ and $M \in \mathbb{N}$, there is a real-coefficient polynomial $p(z) = \sum_{j=0}^M a_j z^j$ with $|a_0| \geq L \cdot (\sum_{j=1}^M |a_j|)$ such that p has at least $T_{L,M} = \min\{\frac{2}{7}\sqrt{M \cdot (-\ln L)}, M\}$ repeated roots at 1.*

We will also use the following result from [BEK99]:

Claim 4.6 (Lemma 5.4 of [BEK99]). *Let $p : \mathbb{C} \rightarrow \mathbb{C}$ be defined as $p(x) = \sum_{j=0}^M a_j z^j$ where $|a_j| \leq 1$ for all $0 \leq j \leq n$. Further, let p have k repeated roots at 1. Let A define the interval $[1 - k/(9M), 1]$. Then*

$$\sup_{z \in A} |p(z)| \leq (M+1) \left(\frac{e}{9} \right)^k.$$

With these two results in hand, we are now ready to prove Theorem 4.2.

Proof of Theorem 4.2. Let us set $M = \lfloor \ln(1/\epsilon)/\nu^2 \rfloor$ and let $p(z) = \sum_{j=0}^M c_j z^j$ be the polynomial from Theorem 4.5 with $L = 2\epsilon$. Let us also scale the coefficients such that $|c_0| = 2\epsilon$ and thus $\sum_{j=0}^M |c_j| \leq 2$. As $\ln(1/\epsilon)/\nu^2 \leq n$, $M \leq n$ and thus our construction is well-defined. The polynomial p has at least T roots at 1, where

$$T = \min \left\{ \frac{2 \ln(1/\epsilon)}{7\nu}, \frac{\ln(1/\epsilon)}{\nu^2} \right\} = \frac{2 \ln(1/\epsilon)}{7\nu}.$$

Let us define $\theta = T/(9M) = (2/63) \cdot \nu$. By applying Claim 4.6, it follows that

$$\sup_{[1-\theta, 1]} |p(z)| \leq (M+1) \cdot \left(\frac{e}{9} \right)^T \leq \left(\frac{1}{3} \right)^T.$$

Here the last inequality uses the relation between T and M and $\epsilon \leq \tau$. Finally, set $a = \nu/63$. Then, applying Corollary 4.4, we obtain

$$\sup_{z \in D_{4a}(1-4a)} |p(z)| \leq \sqrt{\sup_{z \in B_{a,1}} |p(z)|} \cdot 2\sqrt{\exp(9aM)} \leq \sqrt{\left(\frac{1}{3} \right)^T} \cdot 4 \cdot \exp(9aM).$$

Plugging in $a = \nu/63$ and $T = (2M\nu)/7$, we obtain that

$$\sup_{z \in D_{4a}(1-4a)} |p(z)| \leq (1/\epsilon)^{\Omega(\nu)},$$

which concludes the proof. \square

5 Circle bounds for bit-flip noise

5.1 A lower bound on $\eta(\epsilon, \nu)$ for bit-flip noise

In this section we prove the following theorem:

Theorem 5.1. *For $0 < \nu, \epsilon < 1$ and $n \in \mathbb{N}$ which satisfy $\frac{2\ln(2/\epsilon)}{n} \leq \nu \leq 1 - \frac{2\ln(2/\epsilon)}{n}$, we have*

$$\eta(\epsilon, \nu) \geq \epsilon \cdot \exp\left(-O\left(\frac{\ln^{2/3}(1/\epsilon) \cdot (n(1-\nu^2))^{1/3}}{\nu^{2/3}}\right)\right).$$

Proof. Fix any vector $[c_0 \ c_1 \ \dots \ c_n] \in \Delta$ with $|c_0| > 2\epsilon$. Recalling Theorem 3.3 and (4), to prove Theorem 5.1 it suffices to show that the function $F_c(z)$ as defined in (4) satisfies

$$\max_{|z|=1} |F_c(z)| \geq \epsilon \cdot \exp\left(-O\left(\frac{\ln^{2/3}(1/\epsilon) \cdot (n(1-\nu^2))^{1/3}}{\nu^{2/3}}\right)\right). \quad (8)$$

To prove this, we recall (6) which states that

$$\max_{|z|=1} |F_c(z)| = \max_{-\pi < \theta \leq \pi} \left(\frac{1}{1 + \frac{(1-\nu^2)\sin^2(\theta/2)}{\nu^2}}\right)^{n/2} \cdot |Q_c(e^{i\theta})|.$$

Next, we observe that for $-\pi < \theta \leq \pi$, we have

$$1 - (1-\nu^2)\sin^2(\theta/2) \in \left[1 - \frac{(1-\nu^2)\theta^2}{4}, 1 - \frac{(1-\nu^2)\theta^2}{16}\right],$$

where the last inclusion uses $\theta^2/16 \leq \sin^2(\theta/2) \leq \theta^2/4$, which holds for $\theta \in [-\pi, \pi]$. Using the elementary fact $e^{-x} \leq 1/(1+x)$ for all $x \geq 0$, it follows that

$$\left(\frac{1}{1 + \frac{(1-\nu^2)\sin^2(\theta/2)}{\nu^2}}\right) \geq \exp\left(-\frac{1-\nu^2}{4\nu^2}\theta^2\right)$$

and thus, we have

$$\max_{|z|=1} |F_c(z)| \geq \max_{-\pi < \theta \leq \pi} \exp\left(-\frac{1-\nu^2}{8\nu^2}\theta^2 n\right) \cdot |Q_c(e^{i\theta})|. \quad (9)$$

To finish the proof, we recall Corollary 3.2 of [BE97]:

Theorem 5.2 (Corollary 3.2 of [BE97]). *There is a universal constant $c > 0$ such that the following holds: Let $Q(z)$ be a univariate polynomial with complex coefficients, $Q(z) = \sum_{j=0}^n b_j z^j$ with $|b_0| = 1$ and all coefficients $|b_j| \leq M$. Let A be a subarc of the unit circle with length a , where $0 < a < 2\pi$. Then there is some $w \in A$ such that*

$$|Q(w)| \geq \exp\left(\frac{-c(1 + \ln M)}{a}\right).$$

Applying this theorem to the polynomial Q_c with its “ M ” set to $1/c_0$ and its “ a ” set to θ^* and combining with (9), we obtain

$$\max_{|z|=1} |F_c(z)| \geq \max_{-\pi < \theta^* \leq \pi} \exp\left(-\frac{1-\nu^2}{8\nu^2}\theta^{*2} n\right) \cdot \exp\left(-\theta(1) \cdot \frac{1 + \ln(1/c_0)}{\theta^*}\right).$$

Finally set θ^* as

$$\theta^* = \frac{1}{10} \cdot \frac{\nu^{2/3} \cdot \ln^{1/3}(1/\epsilon)}{(n(1-\nu^2))^{1/3}}, a$$

and plug in the right hand side of the above expression (it is easy to see that the constraints on ν imply that $\theta^* \leq 1$). This finishes the proof. \square

5.2 An upper bound on $\eta(\epsilon, \nu)$ for bit-flip noise

In this section we prove the following theorem:

Theorem 5.3. *There is a universal constant $c > 0$ such that for $\nu, 0 < \epsilon < c$ and $n \in \mathbb{N}$ which satisfy $(\frac{2\ln(2/\epsilon)}{n})^{1/4} \leq \nu \leq 1 - \frac{2\ln(2/\epsilon)}{n}$, we have*

$$\eta(\epsilon, \nu) = \exp\left(-\Omega\left(\frac{\ln^{2/3}(1/\epsilon) \cdot (n(1-\nu^2))^{1/3}}{\nu^{2/3}}\right)\right).$$

Recalling (6), to prove this result we must demonstrate the existence of a vector $[c_0 c_1 \dots c_n] \in \Delta$, $|c_0| > 2\epsilon$ such that $F_c(z)$ satisfies

$$\sup_{\|z\|=1} |F_c(z)| = \exp\left(\Omega\left(-\frac{\ln^{2/3}(1/\epsilon) \cdot (n(1-\nu^2))^{1/3}}{\nu^{2/3}}\right)\right), \quad (10)$$

where we recall from Equation (2) that

$$F_c(z) = \sum_{i=0}^n c_i \left(\frac{1-\nu}{2} + \frac{1+\nu}{2}z\right)^i \left(\frac{1+\nu}{2} + \frac{1-\nu}{2}z\right)^{n-i}.$$

To prove this, we will use Theorem 4.5 and the following lemma, which relates the multiplicity of roots of a polynomial at 1 with the supremum of p on an arc centered at 1.

Lemma 5.4 (Lemma 4.7 in [BE97]). *Suppose $p : \mathbb{C} \rightarrow \mathbb{C}$ is a polynomial of the form $p(z) = \sum_{j=0}^M a_j z^j$, where $|a_j| \leq 9$ and p has k repeated roots at 1. If A denotes the arc of the unit circle that is symmetric around 1 and has length $(2k)/(9M)$, then*

$$\sup_{z \in A} |p(z)| \leq 9(M+1) \cdot \left(\frac{e}{9}\right)^k.$$

Proof of Theorem 5.3. With these results in hand we are ready to specify our construction of $[c_0, \dots, c_n]$. For this, we set M as follows:

$$M = \lfloor n^{2/3} \cdot \ln^{1/3}(1/\epsilon) \cdot (1-\nu^2)^{2/3} \cdot \nu^{-4/3} \rfloor.$$

We first make the following observations about M . (i) Since $\nu^4 \geq \frac{\ln(1/\epsilon)}{n}$, it is the case that $M \leq n$. (ii) Since $1-\nu \geq 2\ln(2/\epsilon)/n$, it is moreover the case that $M \geq \ln(1/\epsilon)$.

For M as defined above, let us rescale the polynomial in Theorem 4.5 so that $|a_0| = 2\epsilon$ and thus, $\sum_{j=1}^M |a_j| \leq 1$. We now set $c_j = a_j$ for all $1 \leq j \leq M$ and $c_j = 0$ otherwise. Note that since $M \leq n$, this is well-defined.

By construction, the polynomial $p(z)$ defined as $p(z) = \sum_{j=0}^N c_j z^j$ has at least T repeated roots at 1, where

$$T = \min \left\{ \frac{2}{7} \sqrt{M \cdot \ln(1/2\epsilon)}, M \right\} = \frac{2}{7} \sqrt{M \cdot \ln(1/2\epsilon)},$$

where the last equality uses $1 - \nu \geq 2 \ln(2/\epsilon)/n$. We note for later reference that

$$T = \Omega \left(n^{1/3} \cdot \ln^{2/3}(1/\epsilon) \cdot (1 - \nu^2)^{1/3} \cdot \nu^{-2/3} \right). \quad (11)$$

Let us define θ^* as

$$\theta^* = \frac{2T}{9M} = \frac{4}{63} \sqrt{\frac{\ln(1/2\epsilon)}{M}} \leq \frac{4}{63} \cdot \frac{\ln^{1/3}(1/\epsilon) \cdot \nu^{2/3}}{n^{1/3} \cdot (1 - \nu^2)^{1/3}}. \quad (12)$$

Observe that since $1 - \nu \geq 2 \ln(1/\epsilon)/n$, it holds that $\theta^* \leq 4/63$. Let A be the arc of the unit circle $A = \{e^{i\theta} \mid -\theta^* \leq \theta \leq \theta^*\}$. Applying Lemma 5.4 (and observing that all degree $M + 1$ and higher coefficients of p are zero), we obtain that

$$\sup_{z \in A} |p(z)| = 9 \cdot (M + 1) \cdot \left(\frac{e}{9}\right)^T \leq \left(\frac{1}{3}\right)^T. \quad (13)$$

Here the last inequality uses $T = \frac{2}{7} \sqrt{M \cdot \ln(1/2\epsilon)}$ and the fact that ϵ is at most some sufficiently small constant.

Now we turn our attention to $F_c(z)$. Recalling (4), we have that

$$\sup_{|z|=1} |F_c(z)| = \sup_{|w|=1} \left| \left(\frac{\nu}{\frac{1+\nu}{2} - \frac{1-\nu}{2}w} \right)^n \cdot \sum_{i=0}^n c_i w^i \right|. \quad (14)$$

Let us write $\Phi_c(w)$ to denote $\left(\frac{\nu}{\frac{1+\nu}{2} - \frac{1-\nu}{2}w} \right)^n \cdot \sum_{i=0}^n c_i w^i$, so we seek to upper bound $\sup_{|w|=1} |\Phi_c(w)|$.

We do this by upper bounding $|\Phi_c(w)|$ separately on the sets A and \bar{A} .

First, we bound $|\Phi_c(w)|$ in the set A as follows:

$$\sup_{w \in A} |\Phi_c(w)| \leq \sup_{w \in A} |p(w)| \leq e^{-\Omega(T)} = \exp \left(-\Omega \left(\frac{\ln^{2/3}(1/\epsilon) \cdot (n(1 - \nu^2))^{1/3}}{\nu^{2/3}} \right) \right). \quad (15)$$

Here the first inequality uses the fact that $\left| \frac{\nu}{\frac{1+\nu}{2} - \frac{1-\nu}{2}w} \right| \leq 1$, the second inequality uses (13), and the last equality uses (11).

To bound $|\Phi_c(w)|$ in \bar{A} , we will need a couple of facts. First, since $\sum_{j=0}^n |c_j| \leq 2$, it is the case that $|p(w)| \leq 2$ for all $|w| = 1$, and consequently

$$|\Phi_c(w)| \leq 2 \left| \frac{\nu}{\frac{1+\nu}{2} - \frac{1-\nu}{2}w} \right|^n.$$

Recalling (5), we have

$$\sup_{w \in \bar{A}} |\Phi_c(w)| \leq 2 \left(\frac{1}{1 + \frac{(1-\nu^2) \sin^2(\theta^*/2)}{\nu^2}} \right)^{n/2} \leq 2 \left(\frac{1}{1 + \frac{(1-\nu^2)(\theta^*)^2}{8\nu^2}} \right)^{n/2},$$

where the last inequality uses $\sin^2(\theta^*/2) \geq (\theta^*)^2/8$ which holds since $\theta^* \leq 4/63$. Finally, again using $\nu^4 \geq \ln(1/\epsilon)/n$ and recalling (12), we have $\frac{(1-\nu^2)(\theta^*)^2}{8\nu^2} \leq 4/63$ (with room to spare). Thus, we have that

$$\begin{aligned} \sup_{w \in \bar{A}} |\Phi_c(w)| &\leq 2 \left(\frac{1}{1 + \frac{(1-\nu^2)(\theta^*)^2}{8\nu^2}} \right)^{n/2} \leq \exp \left(-\Omega \left(\frac{(1-\nu^2)(\theta^*)^2 n}{\nu^2} \right) \right) \\ &\leq \exp \left(-\Omega \left(\frac{\ln^{2/3}(1/\epsilon) \cdot (n(1-\nu^2))^{1/3}}{\nu^{2/3}} \right) \right), \end{aligned}$$

where for the last inequality we used $\theta^* = \Theta(1) \cdot \frac{\ln^{1/3}(1/\epsilon) \cdot \nu^{2/3}}{n^{1/3} \cdot (1-\nu^2)^{1/3}}$, which follows from (12). Combining with (15) finishes the proof. \square

Acknowledgments

A. D. would like to thank Mike Saks for suggesting the noisy population recovery problem for unrestricted support and for many illuminating conversations about this problem. The authors would like to thank Tamas Erdélyi for several helpful email exchanges about [Erd16].

References

- [BE97] P. Borwein and T. Erdélyi. Littlewood-type problems on subarcs of the unit circle. *Indiana Univ. Math. J.*, 46:1323–1346, 1997. 5.1, 5.2, 5.4
- [BEK99] P. Borwein, T. Erdélyi, and G. Kós. Littlewood-type problems on $[0, 1]$. *Proc. London Math. Soc. (3)*, 79(1):22–46, 1999. 4.2, 4.6
- [BH51] R. Bellman and T. Harris. Recurrence times for the Ehrenfest model. *Pacific Journal of Mathematics*, 1(2):179–193, 1951. 2.1
- [BIMP13] L. Batman, R. Impagliazzo, C. Murray, and R. Paturi. Finding heavy hitters from lossy or noisy data. In *APPROX-RANDOM 2013*, pages 347–362, 2013. (document), 2.1
- [DOS16] A. De, R. O’Donnell, and R. Servedio. Optimal mean-based algorithms for trace reconstruction. Available at <https://arxiv.org/abs/1612.03148>, 2016. 1.3, 3
- [DRWY12] Z. Dvir, A. Rao, A. Wigderson, and A. Yehudayoff. Restriction access. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA*, pages 19–33, 2012. (document), 2.1
- [DST16] A. De, M. E. Saks, and S. Tang. Noisy population recovery in polynomial time. In *IEEE 57th Annual Symposium on Foundations of Computer Science*, pages 675–684, 2016. (document), 1.2
- [Erd16] T. Erdélyi. Coppersmith–Rivlin type inequalities and the order of vanishing of polynomials at 1. *Acta Arithmetica*, 172:271–284, 2016. 4.2, 4.5, 5.2
- [GL89] O. Goldreich and L. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the Twenty-First Annual Symposium on Theory of Computing*, pages 25–32, 1989. 2.1

- [LZ15] S. Lovett and J. Zhang. Improved noisy population recovery, and reverse Bonami–Beckner inequality for sparse functions. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015*, pages 137–142, 2015. ([document](#)), 1.2
- [MS13] A. Moitra and M. Saks. A polynomial time algorithm for lossy population recovery. In *54th Annual IEEE Symposium on Foundations of Computer Science*, pages 110–116, 2013. ([document](#)), 1.2
- [NP16] F. Nazarov and Y. Peres. Trace reconstruction with $\exp(O(n^{1/3}))$ samples. Available at <https://arxiv.org/abs/1612.03599>, 2016. 1.3
- [PSW17] Y. Polyanskiy, A. T. Suresh, and Y. Wu. Sample complexity of population recovery. Available at <https://arxiv.org/abs/1702.05574>, Submitted on 18 Feb 2017. 1.2, 1.2
- [Sie47] H. A. F. Siebert. Note on the Ehrenfest problem. Technical Report LADC-438, Technical Information Division, Oak Ridge Operations, Oak Ridge, Tennessee, 1947. 2.1
- [WY12] A. Wigderson and A. Yehudayoff. Population recovery and partial identification. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ*, pages 390–399, 2012. ([document](#)), 1.2