Electronic Colloquium on Computational Complexity, Report No. 26 (2016)

Noisy population recovery in polynomial time

Anindya De*
Northwestern University
Evanston, IL, USA
anindya@eecs.northwestern.edu

Michael Saks †
Department of Mathematics
Rutgers University
Piscataway, NJ, USA
saks@math.rutgers.edu

Sijian Tang ‡
Department of Mathematics
Rutgers University
Piscataway, NJ, USA
st509@math.rutgers.edu

February 19, 2016

Abstract

In the noisy population recovery problem of Dvir et al. [DRWY12], the goal is to learn an unknown distribution f on binary strings of length n from noisy samples. For some parameter $\mu \in [0,1]$, a noisy sample is generated by flipping each coordinate of a sample from f independently with probability $(1-\mu)/2$. We assume an upper bound k on the size of the support of the distribution, and the goal is to estimate the probability of any string to within some given error ε. It is known that the algorithmic complexity and sample complexity of this problem are polynomially related to each other.

We show that for $\mu > 0$, the sample complexity (and hence the algorithmic complexity) is bounded by a polynomial in k, n and $1/\varepsilon$ improving upon the previous best result of $\mathsf{poly}(k^{\log\log k}, n, 1/\varepsilon)$ due to Lovett and Zhang [LZ15].

Our proof combines ideas from [LZ15] with a *noise attenuated* version of Möbius inversion. In turn, the latter crucially uses the construction of *robust local inverse* due to Moitra and Saks [MS13].

1 Introduction

1.1 Background and Our Result

The population recovery problem is a basic problem in noisy unsupervised learning which has received significant attention in the recent past [DRWY12, WY12, MS13, LZ15]. In this problem,

*Some part of this work was done while the author was a postdoc at DIMACS, Rutgers.

†Supported by NSF grant CCF-1218711 and by Simons Foundation award 332622.
‡Supported by NSF grant CCF-1218711

ISSN 1433-8092

there is an unknown distribution f over binary strings of length n, and an error parameter $0 < \mu < 1$. Noisy samples from it are generated as:

- Choose a string x according to f.
- Flip each coordinate of x independently with probability $\frac{1-\mu}{2}$.

Given access to these noisy samples, the task of the learner is to output a set of strings S and for each string x in S, an estimate $\tilde{f}(x)$ of $f(x)$, such that $|\tilde{f}(x) - f(x)| \leq \epsilon$. And for all $x \notin S$, $f(x) \leq \epsilon$. For $\mu = 1$, the problem is trivial to solve, whereas for $\mu = 0$, the distribution f cannot be recovered with any number of samples. As μ becomes smaller, the learning problem becomes progressively harder. There is an alternate (and easier) model called the lossy model where instead of flipping bits, each bit is replaced by a '?' independently with probability $1 - \mu$ and presented to the learner.

This problem was introduced by Dvir et al. [DRWY12] who related it to the problem of learning DNF from restrictions. For the lossy model, Dvir et al. [DRWY12] gave a polynomial time algorithm for population recovery for any $\mu \gtrsim 0.365$. Their analysis was improved by Batman, et al. [BIMP13] who showed that the same algorithm works for any $\mu > 1 - 1/\sqrt{2} \approx 0.293$. Subsequently, Moitra and Saks [MS13] gave a polynomial time algorithm for population recovery in the lossy model for any $\mu > 0$.

For the noisy sample problem, algorithms are known only when the support size of f is bounded by a parameter k. Wigderson and Yehudayoff [WY12] developed a framework called "partial identification" and used this to give an algorithm that runs in time $\mathsf{poly}(k^{\log k}, n, 1/\epsilon)$ for any $\mu > 0$. They also showed that their framework cannot obtain algorithms running in time better than $\mathsf{poly}(k^{\log \log k})$.

Lovett and Zheng [LZ15] improved on this to show that the time complexity of this problem is at most $\mathsf{poly}(k^{\log \log k}, n, 1/\epsilon)$ for any $\mu > 0$. Interestingly, while their algorithm matches the lower bound in [WY12], their algorithm departs from the framework of [WY12]. This offers the possibility that one might be able to achieve better algorithms by extending the techniques of [LZ15]. Another interesting feature of this problem is that the algorithmic complexity of the problem is polynomial in the sample complexity of the problem. This seems to have been first explicitly mentioned in [LZ15] though they refer to [BIMP13, MS13]. Thus, it suffices to focus on bounding the sample complexity of the noisy population recovery problem (which is a purely information theoretic quantity).

In this paper, we improve on the results of [LZ15] and show that for any $\mu > 0$, the time complexity of noisy population recovery problem is at most $\mathsf{poly}(k, n, 1/\epsilon)$. This is the first polynomial time algorithm for any $\mu < 1$. The following is our main theorem.

Theorem 1.1. *For any $\mu > 0$, there exists an algorithm for the noisy population recovery problem, running in time $\mathsf{poly}((k/\epsilon)^{O_\mu(1)}, n)$. Here $O_\mu(1) = \tilde{O}(1/\mu^4)$.*

For the ensuing discussion, we first fix some preliminaries.

1.2 Preliminaries

In this section, we include some basic preliminaries concerning Fourier expansion and noise operators. Let $f : \{0,1\}^n \to \mathbb{R}$. Recall that any such f can be expressed uniquely as a linear combination of *characters*, where for $S \subseteq [n]$, the character $\chi_S(x)$ is equal to $\prod_{i \in S}(-1)^{x_i}$.

For $S \subseteq [n]$, the *Fourier coefficient* $\widehat{f}(S)$ is defined to be $\widehat{f}(S) = \sum_{x \in \{0,1\}^n} f(x)\chi_S(x)$. With this definition, it follows that $f(x) = \sum_{S \subseteq [n]} f_S \cdot \chi_S(x)$ where $f_S = 2^{-n} \cdot |\widehat{f}(S)|$. We define $\|f\|_1 = \sum_x |f(x)|$ and $\|\widehat{f}\|_{L_1} = 2^{-n} \sum_S |\widehat{f}(S)| = \sum_S |f_S|$. Also we define the support of the Fourier spectrum be $\mathsf{supp}(\widehat{f}) = \{S : \widehat{f}(S) \neq 0\}$.
Let $\mathcal{F}$ be the space of real-valued functions on $\{0,1\}^n$. For $S \subseteq [n]$, we define the operator $X_S : \mathcal{F} \to \mathcal{F}$ as,

$$(X_S f)(x) = f(x) \cdot \chi_S(x),$$

where $f \in \mathcal{F}$. Next, we define the Bonami-Beckner noise operator. For $\mu > 0$ and $i \in [n]$, define $T_{\mu,i} : \mathcal{F} \to \mathcal{F}$ to be the operator that only adds noise in coordinate i. In other words,

$$(T_{\mu,i} f)(x) = \frac{1+\mu}{2} \cdot f(x) + \frac{1-\mu}{2} \cdot f(x^i),$$

where x^i is the element obtain by flipping the i-th bit of x. For $S \subset [n]$ define the operator $T_{\mu,S} : \mathcal{F} \to \mathcal{F}$ to be the tensor product of $T_{\mu,i}$ for $i \in S$. In other words,

$$(T_{\mu,S} f)(x) = \sum_{T \subseteq S} \frac{(1+\mu)^{|T|} \cdot (1-\mu)^{|S \setminus T|}}{2^{|S|}} \cdot f(x^T),$$

where x^T is obtained by flipping x in the coordinates in T. We define the Bonami-Beckner operator $T_\mu : \mathcal{F} \to \mathcal{F}$ as $T_\mu = T_{\mu,[n]}$. Another way to define the action of T_μ is the following. Let D_μ be the product distribution on $\{0,1\}^n$ such that for all $i \in [n]$, $\Pr[e_i = 0] = (1+\mu)/2$, $\Pr[e_i = 1] = (1-\mu)/2$. Then,

$$(T_\mu f)(x) = \mathbb{E}_{e \sim D_\mu}[f(x+e)].$$

This definition implies $|(T_\mu f)(x)| \leq 1, \forall x$.

1.2.1 Robust local inverse for the noise matrix

Let us define the matrix $A_{\mu,n} \in \mathbb{R}^{(n+1)\times(n+1)}$ as

$$A_{\mu,n}(i,j) = \binom{i}{j} \cdot \mu^j \cdot (1-\mu)^{i-j}.$$

We index $[n+1]$ by $0 \leq i \leq n$ and $\binom{i}{j}$ is defined to be 0 if $j > i$. In a key part of the paper, we will use the following key theorem from [MS13].

Theorem 1.2. (Moitra-Saks [MS13]) *For any $\epsilon > 0$, there exists $v \in \mathbb{R}^{n+1}$ such that $\|A_{\mu,n} \cdot v - e_0\|_\infty \leq \epsilon$, $\|v\|_\infty \leq (2/\epsilon)^{(1/\mu)\cdot\log(2/\mu)}$ and the zeroth coordinate of $A_{\mu,n} \cdot v$ is 1. Here $e_0 \in \mathbb{R}^{n+1}$ denotes the unit vector with 1 at the zeroth coordinate. Further, v can be computed in time $\mathsf{poly}(n)$.*

The non-trivial aspect of the above theorem is that while $A_{\mu,n}$ has very small singular values and as a result, $\|A_{\mu,n}^{-1} \cdot e_0\|_\infty$ can be exponentially large in n, by settling for an ϵ-approximate inverse, it is possible to achieve a significantly better bound. Unfortunately, Theorem 1.2 is not exactly stated in these words in [MS13] though it follows very easily from the results there. In Appendix B, we sketch the details on how to obtain Theorem 1.2 from the results in [MS13].

1.2.2 Möbius inversion Let $(P, \preceq)$ be a poset. Let $\mathcal{F}_P$ be the space of real-valued functions on P. Define $\mu : P \times P \to \mathbb{R}$ recursively as follows:

$$\text{For } x \in P,\ \mu(x,x) = 1.$$

$$\text{For } x, y \in P,\ \mu(x,y) = \mathbf{1}_{x \preceq y} \cdot \Big(\sum_{x \preceq z \prec y} -\mu(x,z) \Big).$$

We define $\boldsymbol{\zeta} : \mathcal{F}_P \to \mathcal{F}_P$ and $\mu : \mathcal{F}_P \to \mathcal{F}_P$ as

$$(\boldsymbol{\zeta} f)(x) = \sum_{x \preceq y} f(y) \quad \text{and} \quad (\boldsymbol{\mu} f)(x) = \sum_{x \preceq y} \mu(x,y) \cdot f(y).$$

It is well known (see [Sta97]) that the transforms $\boldsymbol{\zeta}$ and $\boldsymbol{\mu}$ are inverses of each other. $\boldsymbol{\mu}$ is usually referred to as the Möbius transform of the poset P. While $\boldsymbol{\zeta}$ is always well-conditioned (i.e. $\|\boldsymbol{\zeta}\|_{1\to\infty} \leq 1$), the same is not always true for $\boldsymbol{\mu}$. In particular, the entries of matrix defined by μ can be exponentially large in $|P|$. However, in this paper, we consider a special kind of poset for which $\|\boldsymbol{\mu}\|_{1\to\infty}$ is bounded. To state the next proposition, we will require the following definition.

Definition 1. *For $x \in \mathcal{P}([n])$, define $x^{\downarrow} = \{y : y \preceq x\}$. For $C \subseteq \mathcal{P}([n])$, define $C^{\downarrow} = \cup_{x \in C} x^{\downarrow}$. We read $C^{\downarrow}$ as the "downset" generated by C. Also, if C is a set such that for any $x \in C$, the set $x^{\downarrow} \subseteq C$, we say C is "downward closed". Note that since the underlying poset is $\mathcal{P}([n])$, for $x, y \in \mathcal{P}([n])$, $x \preceq y$ is equivalent to $x \subseteq y$.*

Proposition 1.1. *Let $C \subseteq \mathcal{P}([n])$ and $C^{\downarrow} = \{y : \exists x \in C,\ y \preceq x\}$. Consider the poset defined by $C^{\downarrow}$ ordered by set inclusion. Then, the Möbius transform $\boldsymbol{\mu}$ for this poset is defined by*

$$(\boldsymbol{\mu} f)(x) = \sum_{x \preceq y} (-1)^{|y \setminus x|} \cdot f(y).$$

Proof. Since it is obvious that $\boldsymbol{\zeta}$ and $\boldsymbol{\mu}$ are invertible transforms, we will just verify that $\boldsymbol{\mu} \circ \boldsymbol{\zeta} : f \mapsto f$. To see this,

$$\begin{aligned} (\boldsymbol{\mu} \circ \boldsymbol{\zeta} f)(x) &= \sum_{x \preceq y} (-1)^{|y \setminus x|} \cdot (\boldsymbol{\zeta} f)(y), \\ &= \sum_{x \preceq y} (-1)^{|y \setminus x|} \sum_{y \preceq z} f(z) = \sum_{x \preceq y \preceq z} (-1)^{|y \setminus x|} \cdot f(z). \end{aligned}$$

Since the set $C^{\downarrow}$ is downward closed, it is easy to see that for any $z \succ x$, the sum $\sum_{x \preceq y \preceq z} (-1)^{|y \setminus x|} = 0$. This implies that $(\boldsymbol{\mu} \circ \boldsymbol{\zeta} f)(x) = f(x)$ which proves the claim. □

2 Proof Overview

We recall that the samples available to learner are obtained in the following manner: First, an element of $\{0,1\}^n$ is sampled according to f and then each coordinate is flipped independently with probability $(1-\mu)/2$. In other words, we have observations from the distribution $T_\mu f$

and we want to obtain an estimate of f. Dvir, et al. [DRWY12] gave a reduction to the case that there is a known subset X of size $2k$ that contains the support of the distribution; for convenience we rescale parameters so that $|X| = k$. Thus, from now onwards, we can assume that we know the support of f and our task is to estimate the weight assigned by f to these points.

Let us assume that the support of f is $\{x_1, \ldots, x_k\}$. To prove Theorem 1.1, it suffices to give an algorithm to compute $f(x_1)$ up to error ϵ. We first show that without loss of generality, we can assume that $x_1 = 0$ (i.e. the origin). To see this, note that the distribution $x_1 \oplus T_\mu f$ is the same as $T_\mu g$ where $g(x) = f(x \oplus x_1)$.

A very basic observation concerning $T_\mu f$ is that $\widehat{f}(S)$ can be computed efficiently from $T_\mu f$ as long as $|S|$ is small. The following claim formalizes this. For the rest of the discussion, let $\gamma(x, m, z)$ be defined as $\gamma(x, m, z) = \mathsf{poly}((1/x)^m, z)$.

Claim 2.1. *For $S \subset [n]$, $\widehat{f}(S)$ can be computed to additive accuracy ϵ with probability $1 - \kappa$ using $\gamma(\mu, |S|, \epsilon) \cdot \log(1/\kappa)$ samples from $T_\mu f$ in time $n \cdot \gamma(\mu, |S|, \epsilon) \cdot \log(1/\kappa)$.*

Proof. Observe that $\widehat{f}(S) = \mu^{-|S|} \cdot \mathbf{E}_{x \sim T_\mu f} \chi(x)$. Using the fact that $|\chi(x)| \leq 1$ and applying Chernoff bound, we get the claim. □

Thus, for any fixed $\mu > 0$, as long as $|S| = O_\mu(\log k)$, the time and sample complexity of computing $\widehat{f}(S)$ using samples from $T_\mu f$ is bounded by $\mathsf{poly}(n, k, 1/\epsilon)$. For $f : \{0,1\}^n \to \mathbb{R}$, define $\mathsf{Att}(f) \subseteq \mathcal{F}$ as

$$\mathsf{Att}(f) = \{g \in \mathcal{F} \mid E \subseteq \{0,1\}^n \text{ and } \ g : x \mapsto f(x) \cdot (T_\mu \cdot \mathbf{1}_E)(x)\}.$$

A key step in our algorithm is to generalize Claim 2.1 to show that we can compute $\widehat{g}(S)$ for $g \in \mathsf{Att}(f)$ from (sample access to) $T_\mu f$ with the same sample complexity as Claim 2.1.

Claim 2.2. *Let $g : \{0,1\}^n \to \mathbb{R}$ be defined as $g(x) = f(x) \cdot (T_\mu \mathbf{1}_E)(x)$. For $S \subset [n]$, $\widehat{g}(S)$ can be computed to additive accuracy ϵ with probability $1 - \delta$ using $\gamma(\mu, |S|, \epsilon) \cdot \log(1/\kappa)$ samples from $T_\mu f$ in time $n \cdot \gamma(\mu, |S|, \epsilon) \cdot \log(1/\kappa)$. Here, we assume that $\mathbf{1}_E(\cdot)$ can be efficiently computed.*

The proof of Claim 2.2 relies heavily on ideas from [LZ15]. We prove this in Section 4. As a consequence, we have the following corollary.

Corollary 2.3. *Let $\ell : \{0,1\}^n \to \mathbb{R}$ where $T = \|\widehat{\ell}\|_{L_1}$, $S_0 = \max_{S : \widehat{\ell}(S) \neq 0} |S|$. and $g : \{0,1\}^n \to \mathbb{R}$ be defined as Claim 2.2. We assume that $\ell = \sum_S \ell_S \cdot \chi_S(x)$ where ℓ is specified by the list $\{\ell_S\}$. Then, $\langle g, \ell \rangle$ can be computed to accuracy ϵ with probability $1 - \kappa$ using $\gamma(\mu, S_0, \epsilon/T) \cdot \log(|\mathsf{supp}(\widehat{\ell})|/\kappa)$ samples from $T_\mu f$ in time $\gamma(\mu, |S_0|, \epsilon/T) \cdot n \cdot |\mathsf{supp}(\widehat{\ell})| \cdot \log(|\mathsf{supp}(\widehat{\ell})|/\kappa)$.*

Proof. Note that $\langle \ell, g \rangle = \langle \sum_S \ell_S \cdot \chi_S(x), g \rangle = \sum_S \ell_S \cdot \langle \chi_S, g \rangle = \sum_S \ell_S \cdot \widehat{g}(S)$. Claim 2.2 implies that using $\gamma(\mu, S_0, \epsilon/T) \cdot \log(|\mathsf{supp}(\widehat{\ell})|/\delta)$ samples from $T_\mu f$, $\widehat{g}(S)$ can be computed for any $S \in \mathsf{supp}(\widehat{\ell})$ to accuracy ϵ/T with confidence $1 - \kappa/T$. As a corollary, using $\gamma(\mu, S_0, \epsilon/T) \cdot \log(|\mathsf{supp}(\widehat{\ell})|/\kappa)$ samples from $T_\mu f$, we can compute $\widehat{g}(S)$ for all $S \in \mathsf{supp}(\widehat{\ell})$ to accuracy ϵ/T with confidence $1 - \kappa$. Further, the time complexity of this algorithm is $\gamma(\mu, S_0, \epsilon/T) \cdot n \cdot |\mathsf{supp}(\widehat{\ell})| \cdot \log(|\mathsf{supp}(\widehat{\ell})|/\kappa)$. As a result, $\langle \ell, g \rangle$ can be computed to accuracy ϵ in the claimed time and sample complexity.

□

Recall that our task is to compute $f(0)$ to accuracy ϵ. The way we use Corollary 2.3 is as follows: By choosing $E \subseteq \{0,1\}^n$ and $\ell : \{0,1\}^n \to \mathbb{R}$ carefully, one can ensure that $\langle \ell, g \rangle \approx (T_\mu \mathbf{1}_E)(0) \cdot f(0)$. Thus, if we can can (approximately) compute $\langle \ell, g \rangle$, we will obtain an approximation for $f(0)$. The function ℓ is chosen so that $S_0 = O_\mu(\log(k/\epsilon))$, $T = (k/\epsilon)^{O_\mu(1)}$ and $|\mathsf{supp}(\widehat{\ell})| = (k/\epsilon)^{O_\mu(1)}$. Observe that if we plug in the values of S_0, T and $|\mathsf{supp}(\widehat{\ell})|$ in Corollary 2.3, the sample and time complexity of computing $\langle \ell, g \rangle$ (to error ϵ) is $(k/\epsilon)^{O_\mu(1)}$ and time complexity is $\mathsf{poly}(n, (k/\epsilon)^{O_\mu(1)})$. The precise details of this calculation is given in Section 3.

For the moment, we elaborate on how the set E and the function $\ell(\cdot)$ are chosen. $E \subseteq \{0,1\}^n$ is chosen so that $T_\mu \mathbf{1}_E(\cdot)$ has the following properties: $T_\mu \mathbf{1}_E(0) \geq 1/2$ and $T_\mu \mathbf{1}_E(x)$ decays exponentially as x moves away from the origin. The following lemma makes this precise.

Lemma 2.4. *Let $\{x_1, \ldots, x_k\}$ where $x_1 = 0$. Define the set $\mathsf{Far} = \{x_i : d_H(x_1, x_i) \geq (1/\mu^2) \cdot \log k\}$ and define the set $E = \{y \in \{0,1\}^n : d_H(x_1, y) \leq d_H(x_i, y)$ for all $x_i \in \mathsf{Far}\}$.*

- $(T_\mu \mathbf{1}_E)(0) \geq 1/2$.
- *For $x_i \in \mathsf{Far}$, $(T_\mu \mathbf{1}_E)(x_i) \leq e^{-\frac{1}{2} \cdot \mu^2 \cdot d_H(x_1, x_i)}$.*

Clearly, the function $\mathbf{1}_E(\cdot)$ can be computed in time $\mathsf{poly}(n, k)$. Further, $(T_\mu \mathbf{1}_E)(0)$ can be computed to additive error ϵ with confidence $1 - \kappa$ in time $\mathsf{poly}(n, k, 1/\epsilon) \cdot \log(1/\kappa)$.

While the above lemma is essentially identical to Lemma 3.2 in [LZ15], it is phrased a little differently in that paper. For the sake of completeness, we reprove this lemma in Appendix A.

Let $A = \mathsf{supp}(f)$ where $A = \{x_1, \ldots, x_k\}$ with $x_1 = 0$. Let $\mathbf{1}_E : \{0,1\}^n \to \{0,1\}$ be the corresponding function from Lemma 2.4 and $g : \{0,1\}^n \to \mathbb{R}$ be defined as $g(x) = f(x) \cdot (T_\mu \mathbf{1}_E)(x)$. From Lemma 2.4, we get that $g(x)$ decays exponentially in $|x|$ for $|x| \geq \mu^{-2} \cdot \log k$ where $|x|$ denotes the Hamming weight of x. Let $\mathcal{B}_r(0)$ denote the Hamming ball of radius r around the origin. Then, the above implies that if we set $r = O_\mu(\log(k/\epsilon))$, then g essentially vanishes outside $\mathcal{B}_r(0)$. Next, consider a function $\ell : \{0,1\}^n \to \mathbb{R}$ where we set $T = \|\widehat{\ell}\|_{L_1}$, $S_0 = \max_{S \in \mathsf{supp}(\widehat{\ell})} |S|$ such that $\ell(0) = 1$ and $|\ell(x)| \leq \eta$ for $x \in \mathsf{supp}(f) \cap \mathcal{B}_r(0)$. Then, it follows that (as shown in Section 3)

$$\left| \langle \ell, g \rangle - f(0) \cdot (T_\mu \mathbf{1}_E)(0) \right| \leq \eta + T \cdot e^{-\frac{\mu^2 \cdot r}{2}}.$$

If we set $\eta = \epsilon/16$, then it just remains to bound the second term. Thus, we seek to construct $\ell : \{0,1\}^n \to \{0,1\}$ which is 1 at the origin, at most η (in absolute value) for $x \in \mathsf{supp}(f) \cap \mathcal{B}_r(0)$ and T, S_0 and $|\mathsf{supp}(\widehat{\ell})|$ are as small as possible. In particular, in the above error term, we have two competing parameters, namely T and r i.e. as r increases, the value of $T = \|\widehat{\ell}\|_{L_1}$ corresponding to the optimal ℓ, also increases. Thus, it is not immediately obvious if there exists $\ell(\cdot)$ such that the second error term $T \cdot e^{-\frac{\mu^2 \cdot r}{2}}$ can be made vanishingly small. However, for a careful choice of ℓ (as we discuss shortly), the second term can also be made $\epsilon/16$. Thus, $|\langle \ell, g \rangle - f(0) \cdot (T_\mu \mathbf{1}_E)(0)| \leq \epsilon/8$. Thus, if we approximate $\langle \ell, g \rangle$ (as done in Corollary 2.3) and $(T_\mu \mathbf{1}_E)(0)$ (as done in Lemma 2.4), we obtain an ϵ-approximation to $f(0)$.

We now motivate our construction of the function $\ell(\cdot)$. For this, let $C = \mathsf{supp}(f) \cap \mathcal{B}_r(0)$ and let $C^{\downarrow}$ be the downset generated by C. Note that $x \in \{0,1\}^n$ can be identified as the characteristic vector of a subset of $[n]$ and hence, for the following discussion, we alternately

identify $\{0,1\}^n$ with $\mathcal{P}([n])$. We will first start with a suboptimal choice of ℓ which will motivate our final construction.

Corresponding to every $z \in C^{\downarrow}$, consider the monomial $\mathsf{AND}_z : \{0,1\}^n \to \{0,1\}$ defined as

$$\mathsf{AND}_z(x_1, \ldots, x_n) = \prod_{i:z_i=1} x_i.$$

We will define ℓ to be a linear combination of AND_z for $z \in C^{\downarrow}$ subject to the constraints

$$\ell(x) = \begin{cases} 0 & \text{if } x \in C^{\downarrow} \setminus \{0\}, \\ 1 & \text{if } x = 0. \end{cases}$$

Since ℓ is a linear combination of $\{\mathsf{AND}\}_{z \in C^{\downarrow}}$, let us assume that $\ell = \sum_{z \in C^{\downarrow}} \alpha_z \cdot \mathsf{AND}_z$. In terms of the $\boldsymbol{\zeta}$ transform for the poset $C^{\downarrow}$, we can express ℓ as $\ell = \boldsymbol{\zeta}^T(\sum_{z \in C^{\downarrow}} \alpha_z \cdot \mathbf{1}_z)$, where $\mathbf{1}_z$ is the indicator function of z. Thus, $\alpha_z = (\boldsymbol{\mu}^T \ell)(z)$ where $\boldsymbol{\mu}$ is the Möbius transform for $C^{\downarrow}$. Applying Proposition 1.1 and the value of ℓ on $C^{\downarrow}$, we obtain that $\alpha_z = (-1)^{|z|}$. Thus, $\ell(x) = \sum_{z \in C^{\downarrow}} (-1)^{|z|} \mathsf{AND}_z(x)$. It is not difficult to see that for this choice of ℓ, $\|\widehat{\ell}\|_{L_1} \leq |C^{\downarrow}| \leq k \cdot 2^r$ and $\max_{S:\widehat{\ell}(S) \neq 0} |S| \leq r$. This choice of ℓ itself yields non-trivial results. In particular, in an earlier version of this paper, the authors proved Theorem 1.1 with $\mu \gtrsim 0.555$, thereby giving the first polynomial time algorithm for noisy population recovery for any $\mu < 1$. However, to prove Theorem 1.1 for any $\mu > 0$, we require a more refined choice of ℓ. Henceforth, let us refer to the previous choice of ℓ as ℓ_0. In particular, the bottleneck in our argument comes from the fact that we bound $\|\widehat{\ell_0}\|_{L_1}$ by $k \cdot 2^r$. Instead, if we were able to bound $\|\widehat{\ell_0}\|_{L_1} \leq (1+\delta)^r$ for some $\delta < 1$, this would immediately imply improve the lower bound required on μ. If $\delta > 0$ could be made arbitrarily small, then we obtain Theorem 1.1 for all $\mu > 0$.

Towards a better choice of ℓ, we notice that while $\ell_0(x) = 0$ for $x \in C^{\downarrow} \setminus \{0\}$, it suffices to have $|\ell(x)| \leq \eta = \epsilon/16$ for $x \in C^{\downarrow} \setminus \{0\}$. Unfortunately, it is not clear how this relaxed requirement on ℓ can be exploited by the above analysis. To circumvent this, we consider a new family of functions $\{\mathsf{AND}_{\delta,z}\}_{z \in C^{\downarrow}}$ defined as follows.

$$\text{For } x \in C^{\downarrow},\ \mathsf{AND}_{\delta,z}(x) = \mathbf{1}_{x \succeq z} \cdot (1-\delta)^{|x|-|z|},$$

$$\text{and } \widehat{\mathsf{AND}_{\delta,z}} \text{ is supported on } C^{\downarrow}.$$

It is not difficult to see that the above conditions uniquely define $\mathsf{AND}_{\delta,z}$. For points in $C^{\downarrow}$, one can view $\mathsf{AND}_{\delta,z}(\cdot)$ as a *noise attenuated* version of the function $\mathsf{AND}_z(\cdot)$ (obtained by setting $\delta = 0$). We now set $\ell(x) = \sum_{z \in C^{\downarrow}} \alpha_z \cdot \mathsf{AND}_{\delta,z}(x)$. Obtaining the coefficients $\{\alpha_z\}_{z \in C^{\downarrow}}$ can be viewed as a sort of *noise attenuated Möbius inversion*. The flexibility afforded by the parameter δ allows us to exploit the relaxed constraints on ℓ and bound α_z by $\delta^{|z|} \cdot (1/\eta)^{O(\delta^{-1} \cdot \log \delta^{-1})}$. To prove this, we combine basic properties of the Möbius transform on $C^{\downarrow}$ with the robust local inverse from Theorem 1.2. Intuitively, since the function $\mathsf{AND}_{\delta,z}(\cdot)$ combines properties of AND_z with noise attenuation, it is not surprising that the properties of Möbius transform and the robust local inverse are useful in bounding $\{\alpha_z\}_{z \in C^{\downarrow}}$. Further, we show that

$$\|\widehat{\ell}\|_{L_1} \quad \leq \quad k^2 \cdot (1+2\delta)^r \cdot (1/\eta)^{O(\delta^{-1} \cdot \log \delta^{-1})}.$$

This proof of this inequality again uses the structure of $C^{\downarrow}$ as well as bounds on $\|\widehat{\mathsf{AND}_{\delta,z}}\|_{L_1}$ (this proof is given in Section 5). The above bound is incomparable to the bound of $k \cdot 2^r$ we obtained

for the first choice of $\ell = \sum_{z \in C^{\downarrow}} (-1)^{|z|} \cdot \mathsf{AND}_z$. In particular, we pay a dependence on η to bound $\|\widehat{\ell}\|_{L_1}$ whereas the bound on $\|\widehat{\ell_0}\|_{L_1}$ had no dependence on η. However, the place where we make a significant gain is that base of the exponential factor in r can be made arbitrarily close to 1 by choosing a suitably small $\delta > 0$. We summarize the properties of ℓ in the next theorem.

Theorem 2.1. *Let $C \subseteq \{0,1\}^n$ be as defined above where $|C| \leq k$ and $r = \max_{x \in C} |x|$. Given any $\delta, \eta > 0$, there exists $\ell : \{0,1\}^n \to \mathbb{R}$ which is a linear combination of $\{\mathsf{AND}_{\delta,z}\}_{z \in C^{\downarrow}}$ such that*

- *$\ell(0) = 1$ and $|\ell(x)| \leq \eta$ for $x \in C^{\downarrow} \setminus 0$.*
- *$\|\widehat{\ell}\|_1 \leq k^2 \cdot (1+2\delta)^r \cdot (2/\eta)^{\delta^{-1} \cdot \log(2\delta^{-1})}$.*
- *$\widehat{\ell}$ is supported on $C^{\downarrow}$ and hence $\max_{S : \widehat{\ell}(S) \neq 0} |S| = r$.*

Further, let $\ell(x) = \sum_{S \in C^{\downarrow}} \ell_S \cdot \chi_S(x)$. Then, for every $S \in C^{\downarrow}$, ℓ_S can be computed in time $\mathsf{poly}(|C^{\downarrow}|, n)$.

We compare the above theorem with an analogous result in Lovett and Zhang [LZ15] who show the existence of $\ell_{\mathsf{LZ}} : \{0,1\}^n \to \mathbb{R}$ which satisfies

- $\ell_{\mathsf{LZ}}(0) = 1$ and $\ell_{\mathsf{LZ}}(x) = 0$ for $x \in C \setminus 0$.
- $\|\widehat{\ell_{\mathsf{LZ}}}\|_{L_1} \leq k \cdot k^{\log r}$ and $\max_{S : \widehat{\ell_{\mathsf{LZ}}}(S) \neq 0} |S| \leq \log k$.

(This result is implied by Propersition 3.6 in their paper.)
We now compare ℓ_{LZ} with the function ℓ from Theorem 2.1

- $\ell_{\mathsf{LZ}}(x) = 0$ for $x \in C \setminus 0$ whereas we achieve the incomparable guarantee of $|\ell(x)| \leq \eta$ for $x \in C^{\downarrow} \setminus 0$.
- $\widehat{\ell_{\mathsf{LZ}}}$ is supported on a subset of $\mathcal{B}_{\log k}(0)$ whereas $\widehat{\ell}$ is supported on a subset of $\mathcal{B}_r(0)$. Thus, in terms of compactness of $\widehat{\ell}$, Lovett and Zhang achieve a superior guarantee.
- $\|\ell_{\mathsf{LZ}}\|_{L_1} \leq k \cdot k^{\log r}$ whereas for any $\delta, \eta > 0$, we achieve $\|\widehat{\ell}\|_1 \leq k^2 \cdot (1+2\delta)^r \cdot (2/\eta)^{\delta^{-1} \cdot \log \delta^{-1}}$. Our bound has worse asymptotic dependence on r (and a dependence on η). However, when the value of η and r are eventually plugged in (to $\epsilon/16$ and $r = O_\mu(\log(k/\epsilon))$ resp.), the bound on $\|\widehat{\ell}\|_1$ remains $k^{O(1)}$ (for a fixed $\mu > 0$) whereas the bound on $\|\ell_{\mathsf{LZ}}\|_{L_1}$ becomes $k^{O(\log \log k)}$. This is the crucial place where we gain over Lovett and Zhang [LZ15].

This concludes the proof overview. We now give the proof of the main theorem.

3 Proof of Theorem 1.1

Recall that we are assuming that $\mathsf{supp}(f) = \{x_1, \ldots, x_k\}$ where $x_1 = 0$. Also, from our discussion in the preceding section, to prove Theorem 1.1, it suffices to show that $f(0)$ can be approximated to ϵ with $(k/\epsilon)^{\tilde{O}(1/\mu)}$ samples and in time $\mathsf{poly}((k/\epsilon)^{\tilde{O}(1/\mu)}, n)$. Let E be the set defined in Lemma 2.4 and let $g : \{0,1\}^n \to \mathbb{R}$ be defined as $g(x) = f(x) \cdot (T_\mu \cdot \mathbf{1}_E)(x)$.

Let $r \geq \mu^{-2} \cdot \log k$ (whose precise value will be fixed later). Let $C = \mathsf{supp}(f) \cap \mathcal{B}_r(0)$. Let $\delta, \eta > 0$ whose values will be fixed later and let $\ell : \{0,1\}^n \to \mathbb{R}$ be the function from Theorem 2.1 corresponding to the parameters C, r, δ and η. As we have mentioned before, $f(0) \cdot (T_\mu \cdot \mathbf{1}_E)(0) \approx \langle \ell, g \rangle$. Thus, our algorithm to approximate $f(0)$ will be to approximate $\langle \ell, g \rangle$ (call the approximation $\widetilde{\langle \ell, g \rangle}$) and $(T_\mu \cdot \mathbf{1}_E)(0)$ (call the approximation Υ) and return $\widetilde{\langle \ell, g \rangle}/\Upsilon$.

We first bound the difference between $\langle \ell, g \rangle$ and $(T_\mu \cdot \mathbf{1}_E)(0) \cdot f(0)$ in terms of r, k, δ and η.

Claim 3.1.

$$\left|\langle \ell, g \rangle - f(0) \cdot (T_\mu \cdot \mathbf{1}_E)(0)\right| \leq \eta + \|\widehat{\ell}\|_{L_1} \cdot e^{-\frac{\mu^2 \cdot r}{2}}.$$

Proof. Using $\ell(0) = 1$ and the definition of g,

$$\begin{aligned} \left|\langle \ell, g \rangle - f(0) \cdot (T_\mu \cdot \mathbf{1}_E)(0)\right| &= \left|\langle \ell, g \rangle - \ell(0) \cdot g(0)\right| \\ &\leq \sum_{x \in C \setminus 0} \left|\ell(x) \cdot g(x)\right| + \sum_{x \notin C} \left|\ell(x) \cdot g(x)\right|. \end{aligned} \tag{1}$$

Next, we bound the first sum.

$$\begin{aligned} \sum_{x \in C \setminus 0} \left|\ell(x) \cdot g(x)\right| &\leq \sup_{x \in C \setminus 0} |\ell(x)| \cdot \sum_{x \in C \setminus 0} \left|g(x)\right| \\ &\leq \eta \cdot \sum_{x \in C} |g(x)| = \eta \cdot \sum_{x \in C} |f(x) \cdot (T_\mu \mathbf{1}_E)(x)| \\ &\leq \eta \cdot \sum_{x \in C} |f(x)| \leq \eta. \end{aligned} \tag{2}$$

In the above, the second inequality follows from the property of ℓ from Theorem 2.1, the third inequality uses $\|T_\mu \mathbf{1}_E\|_\infty \leq 1$ and the last inequality uses $\|f\|_1 = 1$. Next, we bound the second sum.

$$\begin{aligned} \sum_{x \notin C} \left|\ell(x) \cdot g(x)\right| &\leq \sup_x |\ell(x)| \cdot \sum_{x \notin C} \left|g(x)\right| \\ &\leq \|\widehat{\ell}\|_{L_1} \cdot \sum_{x \notin C} |g(x)| = \|\widehat{\ell}\|_{L_1} \cdot \sum_{x \in C} |f(x) \cdot (T_\mu \mathbf{1}_E)(x)| \\ &\leq \|\widehat{\ell}\|_{L_1} \cdot \sum_{x \in C} |f(x)| \cdot e^{-\frac{\mu^2 \cdot r}{2}} \leq \|\widehat{\ell}\|_{L_1} \cdot e^{-\frac{\mu^2 \cdot r}{2}} \end{aligned} \tag{3}$$

The second inequality uses that for all x, $|\ell(x)| \leq \|\widehat{\ell}\|_{L_1}$, the third inequality uses Lemma 2.4 whereas the last inequality uses $\|f\|_1 = 1$. Plugging (2) and (3) in (1), we obtain the claim. □

Note that $\|\widehat{\ell}\|_1 \leq k^2 \cdot (1+2\delta)^r \cdot (2/\eta)^{\delta^{-1}\cdot\log(2\delta^{-1})}$. If we set,

- $\eta = \epsilon/4$,
- $\delta = \mu^2/16$,
- $r = (100/\mu^4)\cdot\log(1/\mu)\cdot\log(k/\epsilon)$,

then using Claim 3.1, we have $\left|\langle \ell, g\rangle - f(0)\cdot(T_\mu \cdot \mathbf{1}_E)(0)\right| \leq \epsilon/8$.

Let $\ell(x) = \sum_{S\in C^\downarrow} \ell_S \cdot \chi_S(x)$. Using Theorem 2.1, we can assume that we have the complete list $\{\ell_S\}_{S\in C^\downarrow}$ in time $\mathsf{poly}(n, |C^\downarrow|) = \mathsf{poly}(n, k\cdot 2^r) = \mathsf{poly}((k/\epsilon)^{O_\mu(1)}, n)$. Applying Corollary 2.3, using $(k/\epsilon)^{O_\mu(1)}\cdot\log(1/\kappa)$ and time $n\cdot(k/\epsilon)^{O_\mu(1)}\cdot\log(1/\kappa)$, with confidence $1-\kappa$, we can obtain $\widetilde{\langle \ell, g\rangle}$ such that $\left|\widetilde{\langle \ell, g\rangle} - \langle \ell, g\rangle\right| \leq \frac{\epsilon}{16}$. This implies that

$$\begin{aligned}\left|\widetilde{\langle \ell, g\rangle} - f(0)\cdot(T_\mu\cdot\mathbf{1}_E)(0)\right| &\leq \left|\widetilde{\langle \ell, g\rangle} - \langle \ell, g\rangle\right| + \left|\langle \ell, g\rangle - f(0)\cdot(T_\mu\cdot\mathbf{1}_E)(0)\right| \\ &\leq \epsilon/16 + \epsilon/16 = \epsilon/8.\end{aligned}$$

Using Lemma 2.4, we can compute an approximation Υ with confidence $1-\kappa$ such that $|(T_\mu\cdot\mathbf{1}_E)(0) - \Upsilon| \leq \epsilon/8$ and $\Upsilon \geq 1/2$ in time $\mathsf{poly}(n,k,1/\epsilon)\cdot\log(1/\kappa)$. Thus,

$$\left|\frac{\widetilde{\langle \ell, g\rangle}}{\Upsilon} - f(0)\cdot\frac{(T_\mu\cdot\mathbf{1}_E)(0)}{\Upsilon}\right| \leq \frac{\epsilon}{8\cdot\Upsilon} \leq \frac{\epsilon}{4},$$

where the last inequality uses $\Upsilon \geq 1/2$. Further, with probability $1-\kappa$, we have

$$\left|\frac{(T_\mu\cdot\mathbf{1}_E)(0)}{\Upsilon} - 1\right| \leq \epsilon/4.$$

Thus, with probability $1-2\kappa$,

$$\begin{aligned}\left|\frac{\widetilde{\langle \ell, g\rangle}}{\Upsilon} - f(0)\right| &\leq \left|\frac{\widetilde{\langle \ell, g\rangle}}{\Upsilon} - f(0)\cdot\frac{(T_\mu\cdot\mathbf{1}_E)(0)}{\Upsilon}\right| + \left|f(0) - f(0)\cdot\frac{(T_\mu\cdot\mathbf{1}_E)(0)}{\Upsilon}\right| \\ &\leq \frac{\epsilon}{4} + f(0)\cdot\left|1 - \frac{(T_\mu\cdot\mathbf{1}_E)(0)}{\Upsilon}\right| \leq \frac{\epsilon}{4} + f(0)\cdot\frac{\epsilon}{4} \leq \frac{\epsilon}{2}.\end{aligned}$$

This concludes the proof of Theorem 1.1.

4 Proof of Claim 2.2

We begin by restating Claim 2.2.

Claim. *Let $g : \{0,1\}^n \to \mathbb{R}$ be defined as $g(x) = f(x)\cdot(T_\mu\mathbf{1}_E)(x)$. For $S \subset [n]$, $\widehat{g}(S)$ can be computed to additive accuracy ϵ with probability $1-\kappa$ using $\gamma(\mu,|S|,\epsilon)\cdot\log(1/\kappa)$ samples from $T_\mu f$ in time $n\cdot\gamma(\mu,|S|,\epsilon)\cdot\log(1/\delta)$ where $\gamma(\mu,|S|,\epsilon) = \mathsf{poly}((1/\mu)^{|S|}, 1/\epsilon)$. Here, we assume that $\mathbf{1}_E(\cdot)$ can be efficiently computed.*

Since $g(x) = f(x) \cdot (T_\mu \mathbf{1}_E)(x)$, we get that

$$\widehat{g}(S) = \langle (X_S f), (T_\mu \mathbf{1}_E) \rangle = \langle (T_\mu X_S f), \mathbf{1}_E \rangle.$$

We now make two observations. The first is that for any $S \subseteq [n]$, $T_{\mu,S}$ is a self-adjoint operator. The second is that if $S, S' \subseteq [n]$ are disjoint sets, then the operators $X_{S'}$ and $T_{\mu,S}$ commute. Decomposing $T_\mu = T_{\mu,S} T_{\mu,\overline{S}}$, we have

$$T_\mu X_S f = T_{\mu,S} T_{\mu,\overline{S}} X_S f = T_{\mu,S} X_S T_{\mu,\overline{S}} f = T_{\mu,S} X_S T_{\mu,S}^{-1} T_\mu f.$$

Thus, we get

$$\widehat{g}(S) = \langle T_{\mu,S} X_S T_{\mu,S}^{-1} T_\mu f, \mathbf{1}_E \rangle = \mathbf{E}_{z \sim T_\mu f} \langle T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z, \mathbf{1}_E \rangle$$

An easy but crucial fact is the following.

Proposition 4.1. *$\langle T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z, \mathbf{1}_E \rangle$ can be computed in time $\mathsf{poly}(n, 2^{|S|})$.*

Proof. To see this, define $A_{z,S} = \{y : y_{\overline{S}} = z_{\overline{S}}\}$. Observe that

$$\mathsf{supp}(T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z) \subseteq A_{z,S} \text{ and } |A_{z,S}| = 2^{|S|}.$$

Further, $T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z$ can be computed on any point in $A_{z,S}$ in time $2^{O(|S|)}$. Using the fact that $\mathbf{1}_E(\cdot)$ can be efficiently evaluated, we conclude that $\langle T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z, \mathbf{1}_E \rangle$ can be evaluated in time $\mathsf{poly}(n, 2^{|S|})$. □

Based on the above relation, our procedure to estimate $\widehat{g}(S)$ will be a simple random sampling procedure. Let M be a sufficiently large number (which will be fixed soon).

- Sample $z_1, \ldots, z_M \sim T_\mu f$.
- Return $\widetilde{g}_S = M^{-1} \cdot \big(\sum_{i=1}^{M} \langle T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z, \mathbf{1}_E \rangle \big)$.

To establish an upper bound on M, we recall the following facts from [LZ15] (Claim 3.5 in [LZ15]).

Claim 4.2. *$\|T_{\mu,i}\|_{1 \to 1} = 1$ and $\|T_{\mu,i}^{-1}\|_{1 \to 1} = 1/\mu$.*

The above immediately implies

$$\|T_{\mu,S}\|_{1 \to 1} \leq 1, \qquad \|T_{\mu,S}^{-1}\|_{1 \to 1} \leq (1/\mu)^{|S|}. \tag{4}$$

Using $\|X_S\|_{1 \to 1} \leq 1$, this implies that $\|T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z\|_1 \leq (1/\mu)^{|S|}$.

$$\langle T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z, \mathbf{1}_E \rangle \leq \|T_{\mu,S} X_S T_{\mu,S}^{-1} \mathbf{1}_z\|_1 \leq (1/\mu)^{|S|}.$$

An application of Chernoff bound yields that if $M = \mathsf{poly}(1/\epsilon, 1/|\mu|^{|S|}) \cdot \log(1/\kappa)$, then with probability $1 - \kappa$, $|\widetilde{g}(S) - \widehat{g}(S)| \leq \epsilon$.

5 Proof of Theorem 2.1

Towards the proof of Theorem 2.1, we first recall the following basic facts about $\mathsf{AND}_z(\cdot)$.

Proposition 5.1. *For any $z \in \{0,1\}$, the function $\mathsf{AND}_z : \{0,1\}^n \to \{0,1\}$ satisfies the following:*

- $\widehat{\mathsf{AND}_z}$ *is supported on the subsets of* $\{i : z_i = 1\}$,
- *and* $\|\mathsf{AND}_z\|_{L_1} = 1$.

Recall that $C = \{x_1, \ldots, x_k\} \subseteq \{0,1\}^n$ where $|x_i| \leq r$ (we are assuming that the size of the set C is k as opposed to at most k). The next proposition proves important structural properties of the function $\mathsf{AND}_{\delta,z}(\cdot)$.

Proposition 5.2. *Let $0 \leq \delta \leq 1$. Then, for any point $z \in C^{\downarrow}$, there exists $\mathsf{AND}_{\delta,z} : \{0,1\}^n \to \mathbb{R}$, with the following properties:*

- *For $y \in C^{\downarrow}$, $\mathsf{AND}_{\delta,z}(y) = \mathbf{1}_{y \succeq z} \cdot (1-\delta)^{|y|-|z|}$,*
- $\widehat{\mathsf{AND}_{\delta,z}}(S) \neq 0$ *only if $S \in C^{\downarrow}$.*
- $\|\widehat{\mathsf{AND}_{\delta,z}}\|_{L_1} \leq k \cdot (1+\delta)^{r-|z|}$.

Proof. Let $\mathsf{Sym}_j : \mathbb{R}^n \to \mathbb{R}$ as the elementary symmetric polynomial of degree j. We first construct the function $\mathsf{AND}_{\delta,0}$ i.e. the function $\mathsf{AND}_{\delta,z}$ where z is the origin. Towards constructing $\mathsf{AND}_{\delta,0}$, we define the function $h_{\delta,0} : \{0,1\}^n \to \mathbb{R}$ as

$$h_{\delta,0}(y) = \sum_{j=0}^{r} (-\delta)^j \cdot \mathsf{Sym}_j(y) = \sum_{j=0}^{r} \sum_{S \in \binom{[n]}{j}} (-\delta)^j \cdot \mathsf{AND}_S(y).$$

Thus, for any $y \in \{0,1\}^n$, $|y| \leq r$,

$$h_{\delta,0}(y) = \sum_{j=0}^{r} (-\delta)^j \binom{|y|}{j} = (1-\delta)^{|y|}.$$

Next, observe that if $S \notin C^{\downarrow}$, $\mathsf{AND}_S(y) = 0$. We define $\mathsf{AND}_{\delta,0} : \{0,1\}^n \to \mathbb{R}$ as

$$\mathsf{AND}_{\delta,0}(y) = \sum_{j=0}^{r} \sum_{S \in C^{\downarrow} : |S| = j} (-\delta)^j \cdot \mathsf{AND}_S(y).$$

In comparison to $h_{\delta,0}(y)$, the only terms dropped in $\mathsf{AND}_{\delta,0}(y)$ are $\mathsf{AND}_S(y)$ for $S \notin C^{\downarrow}$. Thus, for $y \in C^{\downarrow}$,

$$\mathsf{AND}_{\delta,0}(y) = \sum_{j=0}^{r} (-\delta)^j \binom{|y|}{j} = (1-\delta)^{|y|}.$$

Thus, this satisfies the first requirement. For the second requirement, we observe that $\widehat{\mathsf{AND}_S}$ is supported on $S^{\downarrow}$. Since $C^{\downarrow}$ is closed under downward closure, we get that $\widehat{\mathsf{AND}_{\delta,0}}$ is supported on $C^{\downarrow}$. For the final item, note that

$$\|\widehat{\mathsf{AND}_{\delta,0}}\|_{L_1} \leq \sum_{j=0}^{r} \sum_{S \in C^{\downarrow}:|S|=j} \delta^j \cdot \|\widehat{\mathsf{AND}_S}\|_{L_1} \leq \sum_{j=0}^{r} \sum_{S \in C^{\downarrow}:|S|=j} \delta^j.$$

The last inequality uses Proposition 5.1. Note that $|C^{\downarrow} \cap \{S : |S| = j\}| \leq k \cdot \binom{r}{j}$. Thus,

$$\|\widehat{\mathsf{AND}_{\delta,0}}\|_{L_1} \leq \sum_{j=0}^{r} \sum_{S \in C^{\downarrow}:|S|=j} \delta^j \leq \sum_{j=0}^{r} \binom{r}{j} \cdot k \cdot \delta^j = k(1+\delta)^r.$$

This finishes the construction of $\mathsf{AND}_{\delta,0}$. For $z \in C^{\downarrow} \setminus \{0\}$, let $\mathcal{I}_z = \{i : z_i = 1\}$. Define $\mathsf{AND}_{\delta,0,\mathcal{I}_z} : \mathbb{R}^{n \setminus \mathcal{I}_z} \to \mathbb{R}$ as the function $\mathsf{AND}_{\delta,0}$ when the ambient dimensions are restricted to $[n] \setminus \mathcal{I}_z$. Note that correspondingly, we also project $C^{\downarrow}$ to the coordinates $[n] \setminus \mathcal{I}_z$.

$$\mathsf{AND}_{\delta,z}(y) = \mathsf{AND}_z(y) \cdot \mathsf{AND}_{\delta,0,\mathcal{I}_z}(y).$$

First, by definition of $\mathsf{AND}_{\delta,0,\mathcal{I}_z}(y)$, it follows that for every $y \in C^{\downarrow}$,

$$\mathsf{AND}_{\delta,0,\mathcal{I}_z}(y) = (1-\delta)^{|y_{[n]\setminus\mathcal{I}_z}|} = (1-\delta)^{|y|-|z|}.$$

This implies that $\mathsf{AND}_{\delta,z}(y) = \mathbf{1}_{y \succeq z} \cdot (1-\delta)^{|y|-|z|}$.

Next, by Proposition 5.1, $\widehat{\mathsf{AND}_z}$ is supported on the sets $\mathcal{I}_z$ and by the first part of our proof, $\widehat{\mathsf{AND}_{\delta,0,\mathcal{I}_z}}$ is supported on the projection of $C^{\downarrow}$ to the coordinates in $[n] \setminus \mathcal{I}_z$. This together implies that $\widehat{\mathsf{AND}_{\delta,z}}$ is supported on $C^{\downarrow}$.

Finally, by Proposition 5.1, $\|\widehat{\mathsf{AND}_z}\|_{L_1} = 1$ and by the first part of our proof, $\|\widehat{\mathsf{AND}_{\delta,0,\mathcal{I}_z}}\|_{L_1} \leq k \cdot (1+\delta)^{r-|z|}$. Combining these two, we get $\|\widehat{\mathsf{AND}_{\delta,z}}\|_{L_1} \leq k \cdot (1+\delta)^{r-|z|}$. This finishes the proof. □

Proof of Theorem 2.1. Recall the matrix $A_{\mu,r} \in \mathbb{R}^{(r+1)\times(r+1)}$ is defined as $A_{\mu,r}(i,j) = \binom{i}{j}\mu^j(1-\mu)^{i-j}$ where the rows and columns are indexed by $0 \leq i,j \leq r$. Using Theorem 1.2, there exists $v \in \mathbb{R}^{r+1}$ such that $\|A_{\mu,r} \cdot v - e_0\| \leq \eta$ where $\|v\|_\infty \leq (2/\eta)^{(1/\delta)\log(2/\delta)}$. Further, v can be computed in time $\mathsf{poly}(r)$.
We define $\ell(y) = \sum_{z \in C^{\downarrow}} v_{|z|} \cdot \delta^{|z|} \cdot \mathsf{AND}_{\delta,z}(y)$. First, it easily follows that $\widehat{\ell}$ is supported on $C^{\downarrow}$. For $y \in C^{\downarrow}$ define the set $\mathsf{Down}_{y,t} = \{z \in C^{\downarrow} : z \preceq y \text{ and } |z| = t\}$. Since $C^{\downarrow}$ is closed under downward closure, $|\mathsf{Down}_{y,t}| = \binom{|y|}{t}$. Note that $\mathsf{AND}_{\delta,z}(y) = 0$ for $z \notin \cup_{0 \leq t \leq |y|} \mathsf{Down}_{y,t}$. Thus, for any $y \in C^{\downarrow}$,

$$\ell(y) = \sum_{z \in C^{\downarrow}} v_{|z|} \cdot \delta^{|z|} \cdot \mathsf{AND}_{\delta,z}(y) = \sum_{0 \leq t \leq |y|} v_t \cdot \delta^t \cdot (1-\delta)^{|y|-t} \cdot \binom{|y|}{t} = (A_{\mu,r} \cdot v)_{|y|}.$$

Using Theorem 1.2, $\ell(0) = 1$ and for $x \in C^{\downarrow} \setminus \{0\}$, $|\ell(x)| \leq \eta$. To prove Theorem 2.1, all that

remains is to bound $\|\widehat{\ell}\|_{L_1}$.

$$\begin{aligned}
\|\widehat{\ell}\|_{L_1} &\leq \sum_{z \in C^{\downarrow}} |v_{|z|}| \cdot \delta^{|z|} \cdot \|\widehat{\mathsf{AND}_{\delta,z}}\|_{L_1} \text{ (follows from definition of } \ell) \\
&\leq \sum_{z \in C^{\downarrow}} |v_{|z|}| \cdot \delta^{|z|} \cdot k \cdot (1+\delta)^{r-|z|} \text{ (using Proposition 5.2)} \\
&\leq \|v\|_{\infty} \cdot \sum_{z \in C^{\downarrow}} \delta^{|z|} \cdot k \cdot (1+\delta)^{r-|z|} \\
&= \|v\|_{\infty} \cdot \sum_{0 \leq j \leq r} \delta^{j} \cdot k \cdot (1+\delta)^{r-j} \cdot |\{z \in C^{\downarrow} : |z| = j\}|. \qquad (5)
\end{aligned}$$

Since $|C| \leq k$ and $C^{\downarrow} \subseteq B_r(0)$, it easily follows that

$$|\{z \in C^{\downarrow} : |z| = j\}| \leq k \cdot \binom{r}{j}.$$

Plugging this in (5), we get

$$\begin{aligned}
\|\widehat{\ell}\|_{L_1} &\leq \|v\|_{\infty} \sum_{0 \leq j \leq r} \delta^{j} \cdot k \cdot (1+\delta)^{r-j} \cdot k \cdot \binom{r}{j} \\
&= k^2 \cdot \|v\|_{\infty} \cdot (1+2\delta)^{r}.
\end{aligned}$$

Using $\|v\|_{\infty} \leq (2/\eta)^{(1/\delta)\cdot\log(2/\delta)}$, we get the final bound on $\|\widehat{\ell}\|_{L_1}$. □

Acknowledgments

A.D. is grateful to Rocco Servedio for many illuminating conversations about this problem.

References

[BIMP13] Lucia Batman, Russell Impagliazzo, Cody Murray, and Ramamohan Paturi. Finding heavy hitters from lossy or noisy data. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 347–362. Springer, 2013.

[DRWY12] Zeev Dvir, Anup Rao, Avi Wigderson, and Amir Yehudayoff. Restriction access. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 19–33. ACM, 2012.

[LZ15] Shachar Lovett and Jiapeng 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*, pages 137–142. ACM, 2015.

[MS13] Ankur Moitra and Michael Saks. A polynomial time algorithm for lossy population recovery. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 110–116. IEEE, 2013.

[Sta97] Richard Stanley. *Enumerative Combinatorics*. Cambridge University Press, 1997.

[WY12] Avi Wigderson and Amir Yehudayoff. Population recovery and partial identification. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 390–399. IEEE, 2012.

A Proof of Lemma 2.4

We begin by restating Lemma 2.4.

Lemma. *Let* $\{x_1, \ldots, x_k\}$ *where* $x_1 = 0$. *Define the set* $\mathsf{Far} = \{x_i : d_H(x_1, x_i) \geq (1/\mu^2) \cdot \log k\}$ *and define the set* $E = \{y \in \{0,1\}^n : d_H(x_1, y) \leq d_H(x_i, y)$ *for all* $x_i \in \mathsf{Far}\}$.

- $(T_\mu \mathbf{1}_E)(0) \geq 1/2$.
- *For* $x_i \in \mathsf{Far}$, $(T_\mu \mathbf{1}_E)(x_i) \leq e^{-\frac{1}{2}\cdot\mu^2 \cdot d_H(x_1,x_i)}$.

Clearly, the function $\mathbf{1}_E(\cdot)$ *can be computed in time* $\mathsf{poly}(n,k)$. *Further,* $(T_\mu \mathbf{1}_E)(0)$ *can be computed to additive error* ϵ *in time* $\mathsf{poly}(n, k, 1/\epsilon)$.

Proof. We first lower bound $(T_\mu \mathbf{1}_E)(x_1)$. Let $s = \log(k)/\mu^2$. By definition,

$$\begin{aligned}(T_\mu \mathbf{1}_E)(x_1) = \Pr_{e\sim D_\mu}[x_1 + e \in E] &= 1 - \Pr_{e\in D_\mu}[x_1 + e \notin E] \\ &\geq 1 - \left(\sum_{i : d_H(x_1,x_i) \geq s} \Pr_{y \sim x_1 + e}[d_H(x_1, y) \geq d_H(x_i, y)] \right)\end{aligned}$$

The last inequality follows by the definition of E and union bound. To lower bound the right hand side, let us define $S_i = \{j \in [n] :$ x_i and x_i differ in the j^{th} coordinate$\}$. If $d_H(x_i, x_1) \geq s$, then $|S_i| \geq s$. For such a point $x_i \in S$,

$$\Pr_{y\sim x_1 + D_\mu}[d_H(x_1, y) \geq d_H(x_i, y)] = \Pr_{e\sim D_\mu}\left[\sum_{j\in S_i} e_j \geq |S_i|/2\right]$$

To bound the above sum, we recall the Chernoff bound.

Proposition. *Let* $X_1, \ldots, X_n$ *be* n *independent* $\{0,1\}$ *random variables such that* $1 \leq i \leq n$, $\mathbf{E}[X_i] = p$. *If* $q > p$, *then,*

$$\Pr\left[X_1 + \ldots + X_n \geq n \cdot q\right] \leq \exp\left(-\frac{n}{2} \cdot \left(\frac{q}{p} - 1\right)^2 \right).$$

Applying the above proposition, we get that

$$\Pr_{y\sim x_1 + D_\mu}[d_H(x_1, y) \geq d_H(x_i, y)] \leq \exp\left(\frac{-|S_i|}{2} \cdot \mu^2\right) \leq \frac{1}{2k}.$$

This implies that

$$(T_\mu \mathbf{1}_E)(x_1) \geq 1 - \left(\sum_{i : d_H(x_1,x_i) \geq s} \Pr_{y \sim x_1 + e}[d_H(x_1, y) \geq d_H(x_i, y)] \right) \geq \frac{1}{2}.$$

We now upper bound $(T_\mu \mathbf{1}_E)(x_i)$ for $x_i \in \mathsf{Far}$. Note that $(T_\mu \mathbf{1}_E)(x_i) = \Pr_{e\sim D_\mu}[x_i + e \in E]$. Note that if $x_i + e \in E$, then $d_H(x_i + e, x_1) \leq d_H(x_i + e, x_i)$. This implies that $\sum_{j\in S_i} e_j \geq |S_i|/2$. Applying the Chernoff bound, we have

$$(T_\mu \mathbf{1}_E)(x_i) = \Pr_{e\sim D_\mu}[x_i + e \in E] \leq \Pr_{e\sim D_\mu}\left[\sum_{j\in S_i} e_j \geq \frac{|S_i|}{2}\right] \leq e^{-\frac{1}{2}\cdot\mu^2\cdot d_H(x_1,x_i)}.$$

The fact that $\mathbf{1}_E(\cdot)$ can be computed in time $\mathsf{poly}(n, k)$ follows from the definition of E. Further, since $\mathbf{1}_E(\cdot)$ is computable in time $\mathsf{poly}(n, k)$ and D_μ is samplable in time $\mathsf{poly}(n)$, we immediately get that

$$(T_\mu \mathbf{1}_E)(x_1) = \Pr_{e\sim D_\mu}[x_1 + e \in E],$$

can be approximated to ϵ in time $\mathsf{poly}(n, k, 1/\epsilon) \cdot \log(1/\kappa)$ with confidence $1 - \kappa$. □

B Robust local inverse from [MS13]

Recall that the matrix $A_{\mu,n} \in \mathbb{R}^{(n+1)\times(n+1)}$ is defined to be

$$A_{\mu,n}(i, j) = \binom{i}{j} \cdot \mu^j \cdot (1 - \mu)^{i-j},$$

where $\binom{i}{j} = 0$ if $j > i$. Following Moitra and Saks [MS13], we now define an ϵ-local inverse.

Definition 2. *Let $w \in \mathbb{R}^{n+1}$ such that $\|A_{\mu,n} \cdot w - e_0\|_\infty \leq \epsilon$. Such a vector w is said to be an ϵ-local inverse of $A_{\mu,n}$.*

Further, $\|w\|_\infty$ is defined to be the sensitivity of such a vector. Definition 2.1 from [MS13] defines $\sigma_n(\mu, \epsilon)$ to be

$$\sigma_n(\mu, \epsilon) = \min_{\|A_{\mu,n}\cdot w - e_0\|_\infty \leq \epsilon} \|w\|_\infty.$$

The next observation states that the w achieving the optimum in the above definition can be found using linear programming.

Observation B.1. *Using linear programming, it is possible to find $w \in \mathbb{R}^{n+1}$ in time $\mathsf{poly}(n)$ such that $\|A_{\mu,n} \cdot w - e_0\|_\infty \leq \epsilon$, such that $\|w\|_\infty = \sigma_n(\mu, \epsilon)$.*

We now restate Theorem 2.2 from [MS13] which gives an upper bound on $\sigma_n(\mu, \epsilon)$.

Theorem. *For all positive integers n and $\mu, \epsilon > 0$, $\sigma_n(\mu, \epsilon) = (1/\epsilon)^{f(\mu)}$ where $f(\mu) = (1/\mu) \cdot \log(2/\mu)$.*

Now choose $\epsilon_0 = \frac{\epsilon}{1+\epsilon}$ in this theorem. Let α_0 be the zeroth coordinate of $A_{\mu,n} \cdot w$. Note that $1 + \frac{\epsilon}{1+\epsilon} \geq \alpha_0 \geq 1 - \frac{\epsilon}{1+\epsilon}$. Define $v = w/\alpha_0$. Then the zeroth coordinate of $A_{\mu,n} \cdot w$ is 1. For the other coordinate $i \neq 0$, we have:

$$|(A_{\mu,n} \cdot v)_i| = |(A_{\mu,n} \cdot w)_i|/\alpha_0 \leq \frac{\epsilon}{1+\epsilon} \cdot \left(1 - \frac{\epsilon}{1+\epsilon}\right)^{-1} \leq \epsilon$$

Also we have: $\|v\|_\infty = (1/\alpha_0) \cdot \|w\|_\infty \leq (1 - \frac{\epsilon}{1+\epsilon})^{-1}((1 + \epsilon)/\epsilon)^{(2/\mu)\cdot\log(1/\mu)} \leq (2/\epsilon)^{(2/\mu)\cdot\log(1/\mu)}$. This proves Theorem 1.2.

ECCC ISSN 1433-8092
http://eccc.hpi-web.de