

Bounds for list-decoding and list-recovery of random linear codes*

Venkatesan Guruswami², Ray Li¹, Jonathan Mosheiff², Nicolas Resch², Shashwat Silas¹,
and Mary Wootters¹

¹*Stanford University*

²*Carnegie Mellon University*

April, 2020

Abstract

A family of error-correcting codes is list-decodable from error fraction p if, for every code in the family, the number of codewords in any Hamming ball of fractional radius p is less than some integer L that is independent of the code length. It is said to be list-recoverable for input list size ℓ if for every sufficiently large subset of codewords (of size L or more), there is a coordinate where the codewords take more than ℓ values. The parameter L is said to be the “list size” in either case. The capacity, i.e., the largest possible rate for these notions as the list size $L \rightarrow \infty$, is known to be $1 - h_q(p)$ for list-decoding, and $1 - \log_q \ell$ for list-recovery, where q is the alphabet size of the code family.

In this work, we study the list size of *random linear codes* for both list-decoding and list-recovery as the rate approaches capacity. We show the following claims hold with high probability over the choice of the code (below q is the alphabet size, and $\varepsilon > 0$ is the gap to capacity).

- A random linear code of rate $1 - \log_q(\ell) - \varepsilon$ requires list size $L \geq \ell^{\Omega(1/\varepsilon)}$ for list-recovery from input list size ℓ . This is surprisingly in contrast to completely random codes, where $L = O(\ell/\varepsilon)$ suffices w.h.p.
- A random linear code of rate $1 - h_q(p) - \varepsilon$ requires list size $L \geq \lceil h_q(p)/\varepsilon + 0.99 \rceil$ for list-decoding from error fraction p , when ε is sufficiently small.
- A random *binary* linear code of rate $1 - h_2(p) - \varepsilon$ is list-decodable from *average* error fraction p with list size with $L \leq \lceil h_2(p)/\varepsilon \rceil + 2$. (The average error version measures the average Hamming distance of the codewords from the center of the Hamming ball, instead of the maximum distance as in list-decoding.)

The second and third results together precisely pin down the list sizes for binary random linear codes for both list-decoding and average-radius list-decoding to three possible values.

Our lower bounds follow by exhibiting an explicit subset of codewords so that this subset—or some symbol-wise permutation of it—lies in a random linear code with high probability. This uses a recent characterization of (Mosheiff, Resch, Ron-Zewi, Silas, Wootters, 2019) of configurations of codewords that are contained in random linear codes. Our upper bound follows from a refinement of the techniques of (Guruswami, Håstad, Sudan, Zuckerman, 2002) and strengthens a previous result of (Li, Wootters, 2018), which applied to list-decoding rather than average-radius list-decoding.

*RL, SS and MW are partially funded by NSF-CAREER grant CCF-1844628, NSF-BSF grant CCF-1814629, and a Sloan Research Fellowship. RL is partially supported by NSF GRFP grant DGE-1656518. SS is partially supported by a Google Graduate Fellowship. VG, JM, and NR are partially funded by NSF grants CCF-1563742 and CCF-1814603. NR is partially supported by NSF grants CCF-1527110, CCF-1618280, CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808 and a Sloan Research Fellowship.

1 Introduction

In coding theory, one is interested in the combinatorial properties of sets $\mathcal{C} \subseteq \mathbb{F}_q^n$.¹ Such a set \mathcal{C} is called a *code of length n* over the *alphabet* \mathbb{F}_q , and the elements $c \in \mathcal{C}$ are called *codewords*.

List-decoding, introduced by Elias and Wozencraft in the 1950’s [Eli57, Woz58], is such a combinatorial property. For $p \in [0, 1]$ and integer $L \geq 1$, we say that a code $\mathcal{C} \subseteq \mathbb{F}_q^n$ is (p, L) -*list-decodable* if, for all $z \in \mathbb{F}_q^n$,

$$|\{c \in \mathcal{C} : \delta(c, z) \leq p\}| < L,$$

where $\delta(x, y) = \frac{1}{n} |\{i : x_i \neq y_i\}|$ denotes relative Hamming distance. That is, \mathcal{C} is list-decodable if not too many codewords of \mathcal{C} live in any small enough Hamming ball. In this paper, we are interested in the trade-offs between p , L , and the rate of the code \mathcal{C} . The *rate* R of \mathcal{C} is defined as $R = \frac{\log_q |\mathcal{C}|}{n}$. The rate lies in the interval $[0, 1]$, and larger is better.

Variations of list-decoding. In this work, we consider standard list-decoding along with two variations.

The first variation is a strengthening of list-decoding known as *average-radius list-decoding*. A code \mathcal{C} is (p, L) -*average-radius list-decodable* if for any set $\Lambda \subseteq \mathbb{F}_q^n$ of size L ,

$$\frac{1}{L} \sum_{v \in \Lambda} \delta(c, v) \geq p.$$

It is not hard to see that (p, L) -average-radius list-decodability implies (p, L) -list-decodability, and this stronger formulation has led to stronger lower bounds than are achievable otherwise [GN14]. In addition to stronger lower bounds, average-radius list-decoding—essentially replacing a maximum with an average in the definition of list-decoding—is a natural concept, and it has helped establish connections between list-decoding and compressed sensing [CGV13].

The second variation, known as *list-recovery*, is a version where the “noise” is replaced by uncertainty about each symbol of the received word z . Formally, we say that a code \mathcal{C} is (ℓ, L) -*list-recoverable* if for any sets $S_1, \dots, S_n \subseteq \mathbb{F}_q$ with $|S_i| \leq \ell$ for all i ,

$$|\{c \in \mathcal{C} : c_i \in S_i \forall i\}| < L.$$

List-recovery was originally used as a stepping-stone to list-decoding and unique-decoding (e.g., [GI01, GI02, GI03, GI04]) but it has since become a useful primitive in its own right, with applications beyond coding theory [INR10, NPR11, GNP⁺13, HIOS15, DMOZ19].

Pinning down the output list size. We are motivated by the problem of pinning down the output list size L for (average-radius) list-decoding and for list-recovery. For all three of these problems, given q and p (respectively, q and ℓ), there exists an *optimal rate*, denoted R^* . Namely, R^* is the largest rate so that, for any $\varepsilon > 0$, there are q -ary codes of rate $R^* - \varepsilon$ and arbitrarily large length, which are (p, L) -(average-weight)-list-decodable (resp. (ℓ, L) -list-recoverable), for some $L(q, p, \varepsilon)$ (resp. $L(q, \ell, \varepsilon)$). Importantly, L must not depend on the length of the code. The *list-decoding capacity theorem* gives the dependence of R^* on q and p (resp. q and ℓ): $R^* = 1 - h_q(p)$ for both

¹Here and throughout the paper, \mathbb{F}_q denotes the finite field with q elements. In this we work only consider linear codes, so we always assume that the alphabet is a finite field.

standard and average-radius list-decoding, [Eli91, ZP81] and $R^* = 1 - \log_q(\ell)$ for list-recovery (e.g., [RW18]).

We are interested in the trade-off between the list size L , the parameters p, q, ℓ of the problem, and this gap ε ; we refer to ε as the *gap to capacity*. Pinning down the list size L is an important problem. For example, for many of the algorithmic applications within coding theory, the list size represents a bottleneck on the running time of an algorithm that must check each item in the list before pruning it down [GI04, DL12, GX12, GX13, GK16]. For applications in pseudorandomness, for example to expanders or extractors, the list size corresponds to the expansion or to the amount of entropy in the input, respectively, and it is of interest to precisely pin down these quantities.

We make progress on pinning down the output list sizes for the case of *random linear codes*. A random linear code is a uniformly random subspace of \mathbb{F}_q^n of certain dimension. The list-decodability of random linear codes has been well studied for many reasons [ZP81, GHK11, CGV13, Woo13, RW14a, RW18, LW18]. First, it is a natural mathematical question that studies the interplay between two fundamental notions in \mathbb{F}_q^n : subspaces and Hamming balls. Second, there are constructions of codes which use random linear codes (and their list-decodability) as a building block [GI04, GR08, HW18, HRW17], and improvements in the parameters of random linear codes will lead to improvements in these constructions as well. Third, random linear codes can be seen as one way to partially derandomize completely random codes; this is especially motivating in the binary (or fixed alphabet) case, where we do not know of any explicit constructions of optimally list-decodable codes, linear or otherwise.

1.1 Contributions

Our main results are improved bounds on the list size of random linear codes. We defer the formal theorem statements until after we have set up notation, but we informally summarize our results here. Below, we consider codes of *rate* $R^* - \varepsilon$, where as above we use R^* to denote best achievable rate for each particular problem.

- (1) **Lower bound on the list size for list-recovery of random linear codes.** We show that if a random linear code of rate $R^* - \varepsilon$ is list-recoverable with high probability with input list sizes ℓ and output list size L , then we must have $L = \ell^{\Omega(1/\varepsilon)}$. This is in contrast to completely random codes, for which the output list size is $L = O(\ell/\varepsilon)$ with high probability.

This gap between random linear codes and completely random codes demonstrates that in some sense zero-error list-recovery behaves more like *erasure-list-decoding* [Gur03] than it does like list-decoding with errors. Such a gap is present between general and linear codes in erasure list-decoding, but as we see below, there is no such gap for list-decoding from errors.

Our result extends to the setting of list-recovery with erasures as well. The formal theorem statement and proof can be found in Section 3.

- (2) **Better lower bounds on the list size for list-decoding random linear codes.** We show that if a q -ary random linear code of rate $R^* - \varepsilon$ is list-decodable with high probability up to radius p with an output list size of L , then we must have $L \geq \left\lfloor \frac{h_q(p)}{\varepsilon} + 0.99 \right\rfloor$. By [LW18], this result is tight for list-decoding of binary random linear codes up to a small additive factor. As an immediate corollary, $L \geq \left\lfloor \frac{h_q(p)}{\varepsilon} + 0.99 \right\rfloor$ for *average-radius* list-decoding of random linear codes as well, and, as we will see below, this is also tight for binary random linear codes, up

to a small additive factor. We conjecture that the leading constant $h_q(p)$ is also correct for $q > 2$.

Previous work [GN14] has established that $L = \Omega(1/\varepsilon)$, but to the best of our knowledge this is the first work that pins down the leading constant. In particular, [GN14] shows that, in the situation above, we have $L \geq c_{p,q}/\varepsilon$, where $c_{p,q}$ is a constant that goes to zero as p goes to $1 - 1/q$. In contrast, we show below that the leading constant is at least $h_q(p)$, which goes to 1 as p goes to $1 - 1/q$.

The formal theorem statement and proof can be found in Section 4.

- (3) **Completely pinning down the list size for average-radius list-decoding of binary random linear codes.** We prove a new upper bound on the *average-radius* list-decodability of binary random linear codes, which matches our lower bound, even up to the leading constant. More precisely, we show that with high probability, a random binary linear code of rate $R^* - \varepsilon$ is *average-radius* list-decodable up to radius p with $L \leq \lfloor h_2(p)/\varepsilon \rfloor + 2$.

Such a bound was known for standard list-decoding [LW18], but our upper bound holds even for the stronger notion of average-radius list-decoding, and improves the additive constant by 1.² In particular, this shows that for both list-decoding and average-radius list-decoding of binary random linear codes, the best possible L is concentrated on at most three values: $\lfloor h_q(p)/\varepsilon \rfloor + 2$, $\lfloor h_q(p)/\varepsilon \rfloor + 1$ and $\lfloor h(p)/\varepsilon + 0.99 \rfloor$. This tight concentration demonstrates the sharpness of our upper and lower bound techniques.

The formal theorem statement and proof can be found in Section 5.

1.2 Overview of techniques

In this section, we give a brief overview of our techniques.

Lower bounds. To illustrate the techniques for our lower bounds, we warm up with a back-of-the-envelope calculation which suggests why the “right” answer for our result (1) above is $\ell^{\Omega(1/\varepsilon)}$.

Consider a random linear code $\mathcal{C} \subset \mathbb{F}_q^n$, of rate $R = 1 - \log_q(\ell) - \varepsilon$, where $\varepsilon \in (0, \frac{1}{2})$. That is, \mathcal{C} is the kernel³ of a uniformly random matrix sampled from $\mathbb{F}_q^{(1-R)n \times n}$. Suppose that ℓ is a prime power, and $q = \ell^t$ for some $t \geq 2$. Thus, \mathbb{F}_ℓ is a sub-field of \mathbb{F}_q . Let D be an integer slightly smaller than $\frac{1}{2\varepsilon}$ and let $L = \ell^D \geq \ell^{\Omega(1/\varepsilon)}$. We claim that \mathcal{C} is unlikely to be (ℓ, L) -list-recoverable.

Given a matrix $M \in \mathbb{F}_q^{n \times D}$, we write $M \subseteq \mathcal{C}$ (“ \mathcal{C} contains M ”) to mean that each of the columns of M is a codeword in \mathcal{C} .

Let \mathcal{M} denote the set of all full-rank matrices $M \in \mathbb{F}_q^{n \times D}$ that have the following property: for every row M_i of M there exists some $x_i \in \mathbb{F}_q^*$ such that all entries of M_i belong to the set $x_i \cdot \mathbb{F}_\ell$.

We will show that \mathcal{M} is *bad* and *abundant*. By *bad* we mean that a linear code containing a matrix from \mathcal{M} cannot be (ℓ, L) -list-recoverable. We say that \mathcal{M} is *abundant* (for the rate R) if a random linear code of rate R is likely to contain at least one matrix from \mathcal{M} . Clearly, the combination of these properties means that \mathcal{C} is unlikely to be (ℓ, L) -list-recoverable.

²Under our definition of list-decoding, [LW18] show (p, L) list-decodability with $L = \lfloor h(p)/\varepsilon \rfloor + 3$.

³This is one of several natural models for a random linear code. Another possible model is taking a uniformly random subspace of dimension Rn . It is not hard to see that the total variation distance between these distributions is exponentially small. In particular, our model yields a code of dimension exactly Rn with probability $1 - \exp(-\Omega(n))$.

We first prove that \mathcal{M} is bad. Assume that \mathcal{C} contains some matrix $M \in \mathcal{M}$. By linearity of the code, \mathcal{C} also contains every vector of the form Mu , for $u \in \mathbb{F}_q^D$. In particular, consider the set of vectors $B := \{Mu \mid u \in \mathbb{F}_q^D\} \subseteq \mathcal{C}$. Observe that \mathcal{C} cannot be (ℓ, L) -list-recoverable, since B is a “bad list” for list-recoverability with these parameters: First, since M has full-rank, B is of cardinality $\ell^D = L$. Now, given $i \in [n]$, we need to show that there exists a subset $S_i \subseteq \mathbb{F}_q$ with $|S_i| = \ell$, such that $v_i \in S_i$ for all $v \in B$. We take S_i to be the set $x_i \cdot \mathbb{F}_\ell$, which contains all entries of the row M_i . For $j \in [D]$, write $M_{i,j} = x_i \cdot w_j$ ($w_j \in \mathbb{F}_q^\ell$), and let $v = Mu$ for some $u \in \mathbb{F}_q^\ell$. Then

$$v_i = \sum_{j=1}^D M_{i,j}u_j = x_i \cdot \left(\sum_{j=1}^D w_j \cdot u_j \right) \in x_i \cdot \mathbb{F}_\ell,$$

and we conclude that \mathcal{M} is bad.

Showing that \mathcal{M} is abundant is harder, and at this stage we only provide some intuition for this fact. Let us compute the expected number of matrices $M \in \mathcal{M}$ that are contained in \mathcal{C} . First, we estimate the cardinality of \mathcal{M} . One may generate a matrix in \mathcal{M} by choosing each of its rows in an essentially independent fashion.⁴ Choosing a row amounts to choosing one of $\frac{q-1}{\ell-1}$ sets of the form $x \cdot \mathbb{F}_\ell$ ($x \in \mathbb{F}_q^*$) and then taking each entry to be an element of that set. Accounting for multiple counting of the all-zero row, the number of possible rows is thus $\frac{q-1}{\ell-1} \cdot (\ell^D - 1) + 1$, which we approximate as $q \cdot \ell^{D-1}$. Thus, $|\mathcal{M}| \approx (q \cdot \ell^{D-1})^n$. Next, it is not hard to see that a random linear code contains a given matrix of rank r with probability $q^{-(1-R) \cdot r \cdot n}$. Consequently, for $M \in \mathcal{M}$, we have $\Pr[M \subseteq \mathcal{C}] = q^{-(1-R)Dn}$. Therefore,

$$\begin{aligned} \mathbb{E} \{ |\mathcal{M} \cap \mathcal{C}| \} &= |\mathcal{M}| q^{-(1-R)Dn} \approx (q \cdot \ell^{D-1})^n \cdot q^{-(1-R)Dn} \\ &= (\ell^{-1} \cdot q^{1-\varepsilon D})^n = \ell^{(-1+t(1-\varepsilon D))n}, \end{aligned}$$

where, the penultimate equality is due to substituting $1 - \log_q(\ell) - \varepsilon$ for R . Finally, since $t \geq 2$ and $D < \frac{1}{2\varepsilon}$, the right-hand side of the above is $\ell^{\Omega(n)}$. Thus, in expectation, \mathcal{C} contains many “bad” lists for list-recovery.

Of course, this back-of-the-envelope calculation does *not* yield the result advertised above. It might be the case that, even though the expected number of $M \in \mathcal{M}$ so that $M \subseteq \mathcal{C}$ is large, the probability that such an M exists is still small. In fact, as [MRRZ⁺19] shows, there are simple examples where this does happen. Thus, proving that \mathcal{M} is abundant requires more work.

A standard approach to show that \mathcal{M} is abundant would be via the second-moment method. Recently, [MRRZ⁺19] gave a general theorem which encompasses second-moment calculations in this context. In particular, they showed that there is essentially only one reason that a set \mathcal{M} might not be abundant: there exists some matrix $A \in \mathbb{F}_q^{D \times D'}$, such that the set $\{MA \mid M \in \mathcal{M}\}$ is small. If this occurs, we say that \mathcal{M} is *implicitly rare*.⁵ They used this result to study the list-decodability of random Low-Density Parity-Check codes, but we can use their result to do our second moment calculation. We show that our example of \mathcal{M} above⁶ is not implicitly rare, by showing that there is

⁴We say “essentially” since the resulting matrix might not have full rank, but this happens with negligibly small probability.

⁵The term “implicitly rare” is used by the first version of [MRRZ⁺19], available at <https://arxiv.org/abs/1909.06430v1>.

⁶More precisely, we study an example similar to this one; the example above was slightly tweaked to simplify the exposition for this back-of-the-envelope explanation.

no such linear map A . This establishes that the back-of-the-envelope calculation is in fact correct. Appealing to the machinery of [MRRZ⁺19], rather than applying the second moment method from scratch, allows us to get tighter constants with slightly less work, and gives a more principled approach to our lower bounds; indeed, our result (2) follows the same outline.

The intuition for our second result (2) is similar: we give an example of a class \mathcal{M} which is *bad* for list-decoding and *abundant*. We define \mathcal{M} as follows: Let $u \in \mathbb{F}_q^D$ be a random vector with independent Bernoulli $_q(p)$ entries, namely, each entry is 0 with probability $1 - p$, and chosen uniformly from \mathbb{F}_q^* with probability p . Let x be uniformly sampled from \mathbb{F}_q . Let τ denote the distribution (over \mathbb{F}_q^D) of the random vector $u + x \cdot 1_D$. Finally, define \mathcal{M} to be the set of all matrices $M \in \mathbb{F}_q^{n \times D}$, such that a uniformly sampled row of M has the distribution τ . As before, we show that \mathcal{M} is abundant by showing that \mathcal{M} is not implicitly rare and using the result of [MRRZ⁺19].

Upper bounds. Our argument for our upper bound result (3) closely follows that of [LW18], which itself builds on the argument of [GHSZ02]. The argument imagines building the random linear code one dimension at a time and uses a potential function to show that, so long as we do not add too many dimensions, no ball intersects the code too much. We now provide an informal overview of our approach, specifically comparing and contrasting it with the arguments of Guruswami, Håstad, Sudan and Zuckerman [GHSZ02]; and Li and Wootters [LW18].

Let $R = 1 - h(\rho) - \varepsilon$ and put $k := Rn$ (which we assume for exposition is an integer). Note that sampling a random linear code of rate R is the same as sampling $b_1, \dots, b_k \in \mathbb{F}_2^n$ independently and uniformly at random and outputting $\text{span}\{b_1, \dots, b_k\}$. Consider the “intermediate” codes $\mathcal{C}_i = \text{span}\{b_1, \dots, b_i\}$; [LW18] (following [GHSZ02]) define a potential function $S_{\mathcal{C}_i}$ and endeavor to show that $S_{\mathcal{C}_i}$ does not grow too quickly. The work [GHSZ02] demonstrated that this holds in expectation; the work [LW18] improved their argument to show that it holds with high probability. In both cases the potential function is such that it is easy to show that, so long as $S_{\mathcal{C}}$ is $O(1)$, the code \mathcal{C} is list-decodable.

The potential function in these works keeps track of the radius p list-size at each vector $x \in \mathbb{F}_2^n$, that is, the cardinalities $|\{c \in \mathcal{C}_i : \delta(x, c) \leq p\}|$ for $i = 1, \dots, k$, and shows that so long as i is not too large all these cardinalities remain at most L . For average-radius list-decoding, we instead keep track of a sort of “weighted” list size, where codewords that are very close x are weighted more heavily. We can reuse much of the analysis from [LW18] to demonstrate that on the k -th step the potential function is still bounded by a constant (in fact, it is at most 2). The real novelty in our argument is a demonstration that, assuming this potential function is small, the code is indeed (p, L) -average-radius list-decodable. This step is more involved than the argument in [GHSZ02, LW18] to establish (p, L) -list-decodability.

1.3 Related work

We now highlight some related work. In what follows, ε is always the “gap-to-capacity”, i.e., if the capacity for a particular problem is R^* , then the result concerns codes of rate $R^* - \varepsilon$.

Lower bounds for list sizes of arbitrary codes. It is known that a typical (i.e., uniformly random) list-decodable code of rate $R^* - \varepsilon$ has list size $L = \Theta(1/\varepsilon)$, and a natural question to ask is whether *every* code requires a list of size $L = \Omega(1/\varepsilon)$. Blinovsky ([Bli86, Bli05]) showed that lists

of size $\Omega_p(\log(1/\varepsilon))$ are necessary for list-decoding a code of rate $R^* - \varepsilon$. Later, Guruswami and Vadhan [GV05] considered the high-noise regime where $p = 1 - 1/q - \eta$ and showed that lists of size $\Omega_q(1/\eta^2)$ are necessary. Finally, Guruswami and Narayanan [GN14] showed that for *average-radius* list-decoding, the list size must be $\Omega_p(1/\sqrt{\varepsilon})$.

Existing lower bounds for random linear codes. For the special case of random linear codes, Guruswami and Narayanan [GN14] showed that lists of size $c_{p,q}/\varepsilon$ are necessary. The constant $c_{p,q}$ is not explicitly computed (and in fact relies on a constant from [GHK11] which we discuss below), but one can deduce from the proof that if p tends to $1 - 1/q$ then $c_{p,q}$ will tend to 0. Their lower bound follows from a second moment method argument, i.e., they consider a certain random variable X whose positivity is equivalent to the failure of a random linear code to be list-decodable, and then show that $\text{Var}(X) = o(\mathbb{E}[X])^2$. In this sense our approach is similar to theirs, because we rely on results from [MRRZ⁺19] which themselves are proved using a second moment method. However, we are able to get stronger results (in the sense that our leading constant does not decay as $p \rightarrow 1 - 1/q$, and moreover is optimal for binary codes). One of the reasons may be the notion of “implicit rareness” from [MRRZ⁺19], which provides a useful characterization of the lists contained in a random linear code.

The work [GN14] also established lower bounds on list-decoding random linear codes from *erasures*. While we do not discuss list-decoding from erasures in this work (except in the sense that erasure list-recovery is a generalization of list-decoding from erasures), this result is relevant to our work because [GN14] established an exponential lower bound of the form $L \geq \exp(\Omega(1/\varepsilon))$, in contrast to the list size $O(1/\varepsilon)$ that is attained by uniformly random codes. Thus, our results suggest that (zero-error) list-recovery behaves more like list-decoding from erasures than from errors, at least with respect to the list size of random (linear) codes.

Existing upper bounds for random linear codes. We now turn our attention to upper bounds on list sizes for random linear codes. A long line of works [ZP81, GHSZ02, GHK11, CGV13, Woo13, RW14a, RW18, LW18] has studied this problem, and we highlight the most relevant results now. Zyablov and Pinsker [ZP81] showed that random linear codes of rate $R^* - \varepsilon$ have lists of size at most $q^{1/\varepsilon}$.⁷ Guruswami, Håstad, Sudan and Zuckerman [GHSZ02] first showed the existence of capacity-achieving binary linear codes with lists of size $O(1/\varepsilon)$. Li and Wootters [LW18] revisited their techniques and showed that in fact random linear codes of rate $R^* - \varepsilon$ have lists of size $O(1/\varepsilon)$ with high probability; moreover they computed the constant coefficient in the big-Oh notation. However, neither of these results apply to either average-radius list-decoding or to list-recovery. As discussed above in Section 1.2, our new upper bound is the result of an improvement of the techniques of [LW18], which extends their result to average-radius list-decoding.

As for larger alphabets, Guruswami, Håstad and Kopparty [GHK11] showed that there exists a constant $C_{p,q}$ for which random linear codes are $(p, C_{p,q}/\varepsilon)$ -list-decodable with high probability. Unfortunately, if p tends to $1 - 1/q$ then this constant tends to infinity. To address this, an ongoing line of works [CGV13, Woo13, RW14b, RW18] has studied the list-decodability of random linear codes in the “high-noise regime” where p is close to $1 - 1/q$; these results also apply to average-radius list-decodability. These results imply that for binary random linear codes, when $p = 1 - 1/q - \Theta(\sqrt{\varepsilon})$, random linear codes with rate $R^* - \varepsilon$ are average-radius list-decodable with

⁷For list-recovery with input lists of size ℓ , the argument of [ZP81] shows that the list size is at most $q^{\ell/\varepsilon}$. Furthermore, their results for list-decoding also apply to average-radius list-decoding.

list sizes $O(1/\varepsilon)$. However, the constant hiding in the big-Oh is not correct (in particular, the authors do not see how to make it smaller than 2). Moreover, these results only hold in a particular parameter regime for p and ε , and degrade as the alphabet size grows.

As for list-recovery, a result by Rudra and Wootters [RW18] guarantees that random linear codes with rate $R^* - \varepsilon$ over sufficiently large alphabets \mathbb{F}_q have lists of sizes at most $(q\ell)^{O(\log(\ell)/\varepsilon)}$. To the best of our knowledge, no lower bounds were known.

Relevant results for other ensembles of codes. Lastly, we discuss some other results concerning other code ensembles. First of all, recent work of [MRRZ⁺19] shows that a random code from Gallager’s ensemble of LDPC codes [Gal62] achieves list-decoding capacity with high probability. More generally, they show that random LDPC codes have similar combinatorial properties to random linear codes, including list-decoding, average-radius list-decoding, and list-recovery. As part of their approach, they develop techniques to characterize the lists that appear in a random linear code with high probability, which we utilize for our work.

Finally, we note that there are no known explicit constructions of list-decodable codes of rate $R^* - \varepsilon$ which achieve a list size even of $O(1/\varepsilon)$. Over large alphabets, the best explicit constructions of capacity-achieving list-decodable or list-recoverable codes have list sizes at least $(1/\varepsilon)^{\Omega(1/\varepsilon)}$ (e.g., [KRSW18, KRRZ⁺19]). Further, if one insists on *binary* codes, or even codes over alphabets of size independent of ε , we do not know of *any* explicit constructions of list-decodable codes with rate approaching R^* .

Two-point concentration. We showed that the optimal list size L of a random linear code is concentrated on at most three values for both list-decoding and average-radius list-decoding: $\lceil h(p)/\varepsilon \rceil + 2$, $\lceil h(p)/\varepsilon \rceil + 1$, and, if the value is different, $\lceil h(p)/\varepsilon + 0.99 \rceil$.

In [LW18, Theorem 2.5], it was also shown that the optimal list size of a *completely* random code is concentrated on two or three values for list-decoding. This type of concentration is also well studied in graph theory, where it is known that in Erdős-Rényi graphs, a number of graph parameters are concentrated on two values. Examples include the clique number (size of the largest clique) [Mat72, BE76], the chromatic number [Luc91, AK97, AN05], and the diameter [RW10].

1.4 Discussion and open problems

In this work, we have made progress on pinning down the output list sizes for (average-radius) list-decoding and list-recovery for random linear codes. Before we continue with the technical portion of the paper, we highlight some open questions and future directions.

- We showed that random linear codes of rate $1 - h_q(p) - \varepsilon$ are not (p, L) -list-decodable for $L \sim \frac{h_q(p)}{\varepsilon}$. We conjecture this lower bound is tight, i.e. that random linear codes of rate $1 - h_q(p) - \varepsilon$ are (p, L) -(average-radius) list-decodable for $L = \frac{h_q(p)}{\varepsilon}(1 + o(1))$, where the $o(1) \rightarrow 0$ as $\varepsilon \rightarrow 0$. Our Theorem 5.1 (and earlier in [LW18] for list-decoding) shows it is true for $q = 2$, and we conjecture this is true for larger q .
- Our results show that list-decoding and average-radius list-decoding have essentially the same output list sizes over binary alphabets, for random linear codes. It would be interesting to extend this to larger alphabets, or even to more general families of codes. This is especially interesting given that there is an exponential gap in the best known lower bounds (on the

list-size for arbitrary codes) between list-decoding and average-radius list-decoding for general codes.

- We have used different techniques for our upper and lower bounds. However, we think it is an interesting direction to use the characterization of [MRRZ⁺19]—which we used to prove our lower bounds—to prove upper bounds as well. This would entail showing that every sufficiently bad list is implicitly rare.
- Finally, we note that our lower bounds for list-recovery rely on the field \mathbb{F}_q being an extension field (that is, $q = p^t$ for some $t > 1$). It is an interesting question whether or not an exponential lower bound also holds over prime fields. We note that other lower bounds on list-decoding and list-recovery for Reed-Solomon codes also apply only to extension fields [GR05, BSKR09]; perhaps all of these bounds taken together are evidence that better list-decodability may be possible in general over prime fields.

1.5 Organization

In Section 2 we set up notation and formally state the results of [MRRZ⁺19] that we build on for our lower bounds. In Section 3 we state and prove our lower bound on list-recovery of random linear codes. In Section 4 we state and prove our lower bound on the list-decodability of random linear codes. In Section 5 we prove our upper bound on the list-decodability of random linear codes.

2 Preliminaries

In this section, we set notation and introduce the notions and results from [MRRZ⁺19] that we need for our lower bounds.

Notation. Unless otherwise specified, all logarithms are base 2. We use the notation $\exp(x)$ to mean e^x . For an integer a , we define $[a] := \{1, \dots, a\}$. For a vector $x \in \mathbb{F}_q^A$ and $I \subset [A]$, we use $x_I \in \mathbb{F}_q^{|I|}$ to denote the vector $(x_i)_{i \in I}$ with coordinates from I in increasing order. We use $\mathbf{1}_D$ to denote the all ones vector of length D .

We use several notions from information theory. Define the q -ary entropy $h_q : [0, 1] \rightarrow [0, 1]$ by

$$h_q(p) \stackrel{\text{def}}{=} x \log_q(q - 1) - x \log_q x - (1 - x) \log_q(1 - x)$$

We assume $q = 2$ if q is omitted from the subscript.

For a random variable X with domain \mathcal{X} , we use $H(X)$ to denote the *entropy* of X :

$$H(X) = - \sum_{x \in \mathcal{X}} \Pr_X(x) \log(\Pr_X(x)).$$

For a probability distribution τ , we may also use $H(\tau)$ to denote the entropy of a random variable with distribution τ .

Let X be a random variable supported on \mathcal{X} and Y be a random variable supported on \mathcal{Y} . We define the *conditional entropy* of Y given X as

$$H(Y|X) = - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log \frac{p(x, y)}{p(x)}.$$

It is easy to check that $H(X) - H(X|Y) = H(Y) - H(Y|X)$ and we call this the *mutual information* $I(X; Y)$:

$$I(X; Y) = H(X) - H(X|Y) = H(Y) - H(Y|X)$$

For random variables X, Y, Z , we define the *conditional mutual information* $I(X; Y|Z)$ by

$$I(X; Y|Z) = H(X|Z) - H(X|Y, Z) = H(Y|Z) - H(Y|X, Z)$$

Conditional entropy, mutual information, and conditional mutual information satisfy the *data processing inequality*: for any function f supported on the domain of Y , we have

$$H(X|f(Y)) \geq H(X|Y) \quad \text{and} \quad I(X; Y) \geq I(X; f(Y)) \quad \text{and} \quad I(X; Y|Z) \geq I(X; f(Y)|Z).$$

We also use *Fano's inequality*, which states that if X is a random variable supported on \mathcal{X} and Y is a random variable supported on \mathcal{Y} , and if $f : \mathcal{Y} \rightarrow \mathcal{X}$ is a function and $p_{err} = \Pr_{X,Y}[f(Y) \neq X]$

$$H(X|Y) \leq h(p_{err}) + p_{err} \cdot \log(|\mathcal{X}| - 1)$$

We define

$$H_q(X) \stackrel{\text{def}}{=} \frac{H(X)}{\log q}, \quad I_q(X; Y) = \frac{I(X; Y)}{\log q}.$$

and similarly for conditional entropy and conditional mutual information.

For a distribution τ on \mathbb{F}_q^L and a matrix $A \in \mathbb{F}_q^{L' \times L}$, we define the distribution $A\tau$ on $\mathbb{F}_q^{L'}$ in the natural way by

$$\Pr_{A\tau}(x) = \sum_{\{y \in \mathbb{F}_q^L : Ay = x\}} \Pr_{\tau}(y),$$

namely, $A\tau$ is the distribution of the random vector Ay , where $y \sim \tau$.

We have defined list-decoding, average-radius list-decoding, and list-recovery in the introduction. We will in fact consider a more general version of list-recovery, which also tolerates erasures:

Definition 2.1 (List-recovery from erasures). *A code $C \subset \mathbb{F}_q^n$ is (α, ℓ, L) -list-recoverable from erasures if the following holds. Let $S_1, \dots, S_n \subset \mathbb{F}_q$ be lists so that $|S_i| \leq \ell$ for at least αn values of i . Then*

$$|\{c \in C : \forall i \in [n], c_i \in S_i\}| < L.$$

We take $\alpha = 1$ if it is omitted.

Tools from [MRRZ⁺19]. As discussed in Section 1.2, for our lower bounds we use tools from the recent work [MRRZ⁺19]. We work with matrices $M \in \mathbb{F}_q^{n \times L}$ ($L \in \mathbb{N}$), where we view the columns of M as potential codewords in \mathcal{C} . We use the notation “ $M \subseteq \mathcal{C}$ ” to mean that the columns of M are all contained in \mathcal{C} .

We group together sets of such matrices M according to their row distribution.

Definition 2.2 ($\tau_M, \dim(\tau), \mathcal{M}_{n,\tau}$). *Given a matrix $M \in \mathbb{F}_q^{n \times L}$, the empirical row distribution defined by the rows of M over \mathbb{F}_q^L is called the type τ_M of M . That is, τ_M is the distribution so that for $v \in \mathbb{F}_q^L$,*

$$\Pr_{\tau_M}(v) = \frac{|\{i : \text{the } i\text{'th row of } M \text{ is equal to } v\}|}{n}.$$

For a distribution τ on \mathbb{F}_q^L , we use $\dim(\tau)$ to refer to $\dim(\text{span}(\text{supp}(\tau)))$. We use $\mathcal{M}_{n,\tau}$ to refer to the set of all matrices in $\mathbb{F}_q^{n \times L}$ which have empirical column distribution τ .

Remark 2.3. We remark that for some distributions τ over \mathbb{F}_q^L , the set $\mathcal{M}_{n,\tau}$ may be empty due to $n \cdot \Pr_\tau(v)$ not being an integer. For such τ we can define $\mathcal{M}_{n,\tau}$ to consist of matrices M with either $\lfloor n \cdot \Pr_\tau(v) \rfloor$ or $\lceil n \cdot \Pr_\tau(v) \rceil$ copies of v . This has a negligible effect on the analysis as we always take n to be sufficiently large compared to other parameters, so for clarity of exposition we ignore this technicality.

Given $M \in \mathcal{M}_{n,\tau}$, note that $\mathcal{M}_{n,\tau}$ consists exactly of those matrices obtained by permuting the rows of M . In particular, since that *random linear code* model is invariant to such permutations, all of the matrices in $\mathcal{M}_{n,\tau}$ have the same probability of being contained in \mathcal{C} .

As discussed in Section 1.2, we prove a lower bound by exhibiting a distributions τ over \mathbb{F}_q^L such that the corresponding set $\mathcal{M}_{n,\tau}$ is both *bad* and *abundant*. When $\mathcal{M}_{n,\tau}$ satisfies these properties, we say that τ itself is, respectively, *bad* and *abundant*.

The work [MRRZ⁺19] characterizes which distributions τ satisfy the *abundance* property, namely, which classes $\mathcal{M}_{n,\tau}$ are likely to have at least one of their elements appear (as a matrix) in a random linear code of a given rate. To motivate the definition below, suppose that the distribution τ has low entropy: $H_q(\tau) < \gamma \cdot \dim(\tau)$ for some $\gamma \in (0, 1)$. This implies that the class $\mathcal{M}_{n,\tau}$ is not too big: more precisely, it is not hard to see that $|\mathcal{M}_{n,\tau}| \leq q^{H_q(\tau) \cdot n} \leq q^{\gamma \dim(\tau) n}$. Using a calculation like we did in Section 1.2, we see that, since $\mathcal{M}_{n,\tau}$ is not very large, it is unlikely for a random linear code of rate less than $1 - \gamma$ to contain a matrix from $\mathcal{M}_{n,\tau}$. However, this is not the only reason that $\mathcal{M}_{n,\tau}$ might be unlikely to appear in a random linear code. As is shown in [MRRZ⁺19], it could also be because a random output of τ , subject to some linear transformation (perhaps to a space of smaller dimension), has low entropy. We call such distributions *implicitly rare*:

Definition 2.4 (γ -implicitly rare). We say that a distribution τ over \mathbb{F}_q^L is γ -implicitly rare if there exists a full-rank linear transformation $A : \mathbb{F}_q^L \rightarrow \mathbb{F}_q^{L'}$ where $L' \leq L$ such that

$$H_q(A\tau) < \gamma \cdot \dim(A\tau)$$

Observe that, by taking A to be the identity map, we recover the case where τ itself has low entropy. For an illustrative example of an implicitly rare distribution, we refer the reader to [MRRZ⁺19, Example 2.5].

Essentially, [MRRZ⁺19] shows that a row distribution τ is likely to appear in a random linear code (namely, τ satisfies the abundance property) if and only if it is *not* implicitly rare. The following Theorem follows from Lemma 2.7 in [MRRZ⁺19].⁸

Theorem 2.5 (Follows from Lemma 2.7 in [MRRZ⁺19]). Let $R \in (0, 1)$ and fix $\eta > 0$. Let τ be a $(1 - R - \eta)$ -implicitly rare distribution over \mathbb{F}_q^L ($L \in \mathbb{N}$), and let \mathcal{C} be a random linear code of rate R . Then

$$\Pr[\exists M \in \mathcal{M}_{n,\tau} : M \subseteq \mathcal{C}] \leq q^{-\eta n}$$

Conversely, suppose that τ is not even $(1 - R + \eta)$ -implicitly rare. Then

$$\Pr[\exists M \in \mathcal{M}_{n,\tau} : M \subseteq \mathcal{C}] \geq 1 - n^{O_{L,q}(1)} \cdot q^{-\eta n} .$$

⁸This is also given as Theorem 2.2 in the first version of [MRRZ⁺19], available at <https://arxiv.org/abs/1909.06430v1>.

3 Lower bounds for list-recovery with erasures

Our main result in this section is the following.

Theorem 3.1. *Fix $0 \leq \rho < 1$. Fix a prime power $\ell \geq 2$ and an integer $t \geq 2$, and let $q = \ell^t$. Fix $0 < \varepsilon \leq \frac{1-\rho}{20t}$ and let $L = \ell^{\lceil \frac{1-\rho}{20\varepsilon} \rceil}$. For $n \in \mathbb{N}$, let $\mathcal{C} \subseteq \mathbb{F}_q^n$ denote a random linear code of rate $R := 1 - (\rho + (1-\rho)\log_q(\ell)) - \varepsilon$. Then the probability of \mathcal{C} being $(1-\rho, \ell, L)$ -erasure list-recoverable is at most $q^{-\Omega(n)}$.*

3.1 Proof of Theorem 3.1

We will prove Theorem 3.1 below, after we build up the necessary building blocks. As discussed in Sections 1.2 and 2, to prove Theorem 3.1 we seek a distribution τ that is both *bad* and *abundant*. That is, \mathcal{C} should likely contain some matrix from $\mathcal{M}_{n,\tau}$, and the corresponding codewords should yield a counterexample to the list-recoverability of \mathcal{C}_n . We will describe our choice of τ in Definition 3.2; we will show that it is bad in Proposition 3.3; and finally we will show that it is implicitly rare (and hence abundant by Theorem 2.5) in Lemma 3.4.

Our construction of the distribution τ follows similar lines to that in Section 1.2.

Definition 3.2 (The bad distribution τ for list-recovery lower bounds). *Fix ρ, ℓ, t, q as in Theorem 3.1. Let $D \geq t$ be a positive integer. Let \mathbb{F}_ℓ be a subfield of \mathbb{F}_q , where $q = \ell^t$ and $t \geq 2$. Let $\alpha_1, \dots, \alpha_{(q-1)/(\ell-1)}$ be a set so that $\alpha_i \mathbb{F}_\ell^*$ are distinct cosets of \mathbb{F}_ℓ^* partitioning \mathbb{F}_q^* . Let $L = \ell^D$. Let $G \in \mathbb{F}_\ell^{L \times D}$ be the matrix whose rows are all of the distinct elements of \mathbb{F}_ℓ^D .*

Let σ be the distribution that with probability $1-\rho$ returns $\alpha_i u$ for (i, u) uniform in $\left\{1, \dots, \frac{q-1}{\ell-1}\right\} \times \mathbb{F}_\ell^D$; and with probability ρ returns a uniformly random element of \mathbb{F}_q^D .

Let τ be the distribution given by Gv for $v \sim \sigma$.

To motivate this construction, consider first the $\rho = 0$ case. Now consider a matrix $M \in \mathbb{F}_q^{n \times L}$ that has row distribution given by τ . If we ignore the coefficients α_i , the columns of M span a D -dimensional subspace of \mathbb{F}_ℓ^n . In particular, they are *bad*, in the sense that each coordinate of these codewords are contained in a list of size ℓ (namely, \mathbb{F}_ℓ). Moreover, as soon as any D linearly independent columns of M are contained in \mathcal{C} , all of the columns of M are contained in \mathcal{C} ; this suggests that it's relatively likely (compared to, say, a random matrix in $\mathbb{F}_q^{L \times n}$) that $M \subseteq \mathcal{C}$. These properties don't change when we multiply by the coefficients α_i : each coordinate is now contained in some list $\alpha_i \mathbb{F}_\ell$ rather than \mathbb{F}_ℓ (notice that the fact that the α_i are coset representatives means that all of these possible lists are distinct, other than zero), and it's still just as likely that $M \subseteq \mathcal{C}$. However, by throwing these multiples α_i into the mix, we have increased the size of $\mathcal{M}_{n,\tau}$, making τ more *abundant*.

To generalize to the $\rho > 0$ case, the construction essentially “frees” a ρ fraction of the coordinates relative to the $\rho = 0$ case. This further increases the size of $\mathcal{M}_{n,\tau}$ (making τ even more abundant), while still maintaining the badness property for list-recovery with a ρ fraction of erasures.

Proposition 3.3 (τ is bad). *Let τ be as in Definition 3.2. Let $\mathcal{C} \subseteq \mathbb{F}_q^n$, and let $M \in \mathcal{M}_{n,\tau}$. If $M \subseteq \mathcal{C}$, then \mathcal{C} is not $(1-\rho, \ell, L)$ -list-recoverable.*

Proof. Suppose that $M \subseteq \mathcal{C}$. Let $w_1, w_2, \dots, w_n \in \mathbb{F}_q^L$ be the rows of M . It suffices to show that there are input lists S_1, \dots, S_n so that $w_j \in S_j^L$ for all $j \in [n]$, and so that for at least $(1-\rho)n$

values of $j \in [n]$, we have $|S_j| \leq \ell$. Recall that each row w_j of M is of the form Gv_j where a $(1 - \rho)$ fraction of the v_j are of the form $\alpha_{i_j} \cdot u_j$ for $(i_j, u_j) \in [(q - 1)/(\ell - 1)] \times \mathbb{F}_\ell^D$, and a ρ fraction of the v_j are arbitrary vectors in \mathbb{F}_q^D .⁹

In the first case, set $S_j = \alpha_{i_j} \cdot \mathbb{F}_\ell$. Because the elements of G are all in \mathbb{F}_ℓ , all the coordinates of $w_j = Gv_j = \alpha_{i_j} Gu_j$ lie in S_j . Moreover by definition $|S_j| \leq \ell$. In the second case, set $S_j = \mathbb{F}_q$. By definition all the coordinates of $w_j \in \mathbb{F}_q^L$ lie in $S_j = \mathbb{F}_q$.

This completes the proof. \square

Next, we will show that τ is not implicitly rare, which will imply that τ is abundant.

Lemma 3.4 (τ is abundant). *Let τ be as in Definition 3.2. Then τ is not $\left(\rho + (1 - \rho) \log_q(\ell) + \frac{1 - \rho}{10D}\right)$ -implicitly rare.*

The proof of Lemma 3.4 is in Section 3.2 below. Before we prove Lemma 3.4, we use it to prove Theorem 3.1.

Proof of Theorem 3.1, assuming Lemma 3.4. Let $0 < \varepsilon \leq \frac{1 - \rho}{20t}$. Let τ be as in Definition 3.2, choosing $D = \left\lceil \frac{1 - \rho}{20\varepsilon} \right\rceil$. By our choice of ε , we indeed have $D \geq t$. Lemma 3.4 shows that τ is not $\left(\rho + (1 - \rho) \log_q(\ell) + \frac{1 - \rho}{10D}\right)$ -implicitly rare. By choice of D , we have $\frac{1 - \rho}{10D} > \varepsilon$. From Theorem 2.5 with $\eta = \frac{1 - \rho}{10D} - \varepsilon$, we see that for any sufficiently large n , a random code of rate

$$\left(1 - \left(\rho + (1 - \rho) \log_q(\ell) + \frac{1 - \rho}{10D}\right)\right) + \eta = 1 - \left(\rho + (1 - \rho) \log_q(\ell)\right) - \varepsilon$$

contains ℓ^D codewords given by a matrix $M \in \mathcal{M}_{n, \tau}$ with probability at least $1 - q^{\Omega(\varepsilon n)}$. By Proposition 3.3, if this occurs, then \mathcal{C} is not $(1 - \rho, \ell, L)$ -list-recoverable. \square

3.2 Proof of Lemma 3.4

In this section we prove Lemma 3.4, which will complete the proof of Theorem 3.1. We first prove the following technical lemma, which roughly states that a distribution with few ‘‘collisions’’ has entropy close to the uniform distribution.

Lemma 3.5. *Let Z be a finite set, and for $z \in Z$, let n_z be a nonnegative integer. Suppose that $N_1 = \sum_{z \in Z} n_z$ and that $N_2 = \sum_{z \in Z} \binom{n_z}{2}$. Then the distribution that samples an element $z \in Z$ with probability proportional to n_z has entropy at least*

$$\log(N_1) - \log\left(1 + \frac{2N_2}{N_1}\right).$$

Proof. The entropy is

$$\begin{aligned} \sum_{z \in Z} \frac{n_z}{N_1} \cdot \log\left(\frac{N_1}{n_z}\right) &= \log(N_1) - \sum_{z \in Z} \frac{n_z}{N_1} \cdot \log(n_z) \\ &\geq \log(N_1) - \log\left(\sum_{z \in Z} \frac{n_z}{N_1} \cdot n_z\right) = \log(N_1) - \log\left(\frac{N_1 + 2N_2}{N_1}\right), \end{aligned}$$

⁹As per Remark 2.3, we may ignore the rounding issue that ρn may not be an integer. This is without loss of generality, as we may replace τ with a very similar distribution so that a $\lceil \rho n \rceil$ fraction of the v_j are arbitrary in \mathbb{F}_q^D , and adjust all parameters by a term that is $o(1)$ as $n \rightarrow \infty$.

as desired. In the inequality, we used Jensen's inequality and that $\log(x)$ is concave. \square

Next, we prove Lemma 3.4. We prove it first for $\rho = 0$, and then use the $\rho = 0$ case to prove the general statement.

Proof of Lemma 3.4 for $\rho = 0$. Fix a matrix $A \in \mathbb{F}_q^{L' \times L}$, and let τ be as in Definition 3.2. Recall that the distribution $A\tau$ is given by Av for $v \sim \tau$.¹⁰ Our goal is to show that, for all A , the distribution $A\tau$ supported on $\mathbb{F}_q^{L'}$ has large entropy.

Let D' be the rank of $AG \in \mathbb{F}_q^{L' \times D}$.

First we show that $A\tau$ has dimension D' . By definition, $\alpha e_i \in \text{supp}(\sigma)$ for all $\alpha \in \mathbb{F}_q$ and all $i \in [D]$, so $\text{supp}(\tau)$ contains $\{G \cdot \alpha \cdot e_i : \alpha \in \mathbb{F}_q, i \in [D]\}$, and thus $\text{span}_{\mathbb{F}_q}(\text{supp}(\tau)) = G \cdot \mathbb{F}_q^D$. Hence,

$$\begin{aligned} \dim(\text{span}_{\mathbb{F}_q}(\text{supp}(A\tau))) &= \dim(\text{span}_{\mathbb{F}_q}(A \cdot \text{supp}(\tau))) \\ &= \dim(A \cdot \text{span}_{\mathbb{F}_q}(\text{supp}(\tau))) \\ &= \dim(AG \cdot \mathbb{F}_q^D) \\ &= D' \end{aligned}$$

as desired.

Next we show that the entropy $H_q(A\tau)$ is at least $D'(\log_q \ell + \frac{1}{10D})$. It suffices to prove that

$$H(A\tau) \geq \log \left(\ell^{D'} \cdot q^{0.1D'/D} \right).$$

Since $AG \in \mathbb{F}_q^{L' \times D}$ has rank D' , there exist D' linearly independent rows whose span contains all the rows of AG . Let $W \in \mathbb{F}_q^{D' \times D}$ be the submatrix of AG obtained by keeping these rows. Note that for this W , for all $v, v' \in \mathbb{F}_q^D$ we have $Wv = Wv'$ if and only if $AGv = AGv'$.

Since W has rank D' and $D' \leq D$, there are D' linearly independent columns of W . Suppose without loss of generality that they are the first D' columns of W . Thus, we may write $W = [W^{(1)} | W^{(2)}]$ where $W^{(1)} \in \mathbb{F}_q^{D' \times D'}$ is invertible and $W^{(2)} \in \mathbb{F}_q^{D' \times (D-D')}$. For any $v \in \mathbb{F}_q^D$, we may write $v = \begin{bmatrix} v^{(1)} \\ v^{(2)} \end{bmatrix}$ where $v^{(1)} \in \mathbb{F}_q^{D'}$, $v^{(2)} \in \mathbb{F}_q^{D-D'}$. Then (recalling σ from Definition 3.2)

$$\begin{aligned} H(A\tau) &= H(W\sigma) \\ &= H_{v^{(1)} \sim \mathbb{F}_\ell^{D'}, v^{(2)} \sim \mathbb{F}_\ell^{D-D'}, i \sim [(q-1)/(\ell-1)]} (\alpha_i W^{(1)} v^{(1)} + \alpha_i W^{(2)} v^{(2)}) \\ &\geq H_{v^{(1)} \sim \mathbb{F}_\ell^{D'}, v^{(2)} \sim \mathbb{F}_\ell^{D-D'}, i \sim [(q-1)/(\ell-1)]} (\alpha_i W^{(1)} v^{(1)} + \alpha_i W^{(2)} v^{(2)} | W^{(2)} v^{(2)}) \\ &= \sum_{w \in \mathbb{F}_\ell^{D'}} \Pr_{v^{(2)} \sim \mathbb{F}_\ell^{D-D'}} [W^{(2)} v^{(2)} = w] H_{v^{(1)} \sim \mathbb{F}_\ell^{D'}, i \sim [(q-1)/(\ell-1)]} (\alpha_i W^{(1)} v^{(1)} + \alpha_i w) \\ &\geq \min_{w \in \mathbb{F}_\ell^{D'}} H_{v^{(1)} \sim \mathbb{F}_\ell^{D'}, i \sim [(q-1)/(\ell-1)]} (\alpha_i W^{(1)} v^{(1)} + \alpha_i w), \end{aligned}$$

where above we are using the notation $H_{x \sim X}$ to denote that the randomness in the definition of the entropy is over the choice of a uniformly random x in X . Thus, it suffices to show, for any fixed vector w , we have

$$H_{v^{(1)} \sim \mathbb{F}_\ell^{D'}, i \sim [(q-1)/(\ell-1)]} (\alpha_i W^{(1)} v^{(1)} + \alpha_i w) \geq \log \left(\ell^{D'} \cdot q^{D'/10D} \right) \quad (1)$$

¹⁰Throughout this proof, the output of τ is treated as a column vector.

Before finishing the proof, we first give some intuition for the remaining details. First consider the case $w = 0$. Note that, over all choices of $v^{(1)}$ and i , the vectors $\alpha_i v^{(1)}$ are all distinct, except the all 0s vector. Thus, as $W^{(1)}$ is invertible, the distribution of $\alpha_i W^{(1)} v^{(1)} + \alpha_i w$ is close to the uniform distribution on approximately $\ell^{D'} \cdot \frac{q-1}{\ell-1} \approx \ell^{D'-1} \cdot q$ vectors, so the entropy is at least roughly $\log(\ell^{D'-1} \cdot q)$; this turns out to be enough.

When $w \neq 0$, we do not have the same near-uniform distribution, but we do have the following useful property that carries over from the $w = 0$ case: for a fixed $w \in \mathbb{F}_q^{D'}$ and $\alpha_i \neq \alpha_j$, there exists at most one pair $(v, u) \in (\mathbb{F}_\ell^D)^2$ such that $\alpha_i W^{(1)} v + \alpha_i w = \alpha_j W^{(1)} u + \alpha_j w$. To see this, suppose for contradiction there are two, (v, u) and (v', u') . Then subtracting, we have $\alpha_i W^{(1)}(v - v') = \alpha_j W^{(1)}(u - u')$. Since $W^{(1)}$ is invertible, we have $\alpha_i(v - v') = \alpha_j(u - u') \in \alpha_i \mathbb{F}_\ell^D \cap \alpha_j \mathbb{F}_\ell^D = \{0\}$. Thus, $v = v'$ and $u = u'$, a contradiction. Using this property, we know that, over the randomness of $v^{(1)}$ and i , there are not many ‘‘collisions’’ in $\alpha_i W^{(1)} v^{(1)} + \alpha_i w$, so the entropy should again be close to the entropy of the uniform distribution on $(i, v^{(1)})$, which is $\log(\ell^{D'-1} \cdot q)$. We can bound the entropy of such a distribution with few collisions with a careful application of Jensen’s inequality (Lemma 3.5). We then show the resulting bound is sufficient by some straightforward calculations. We note that our bounds hold for all prime powers ℓ and all $t \geq 2$, rather than simply for sufficiently large ℓ and t ; this requires the argument to be a little more delicate.

We now show the rest of the proof. Fix $w \in \mathbb{F}_q^{D'}$. For $z \in \mathbb{F}_q^{D'}$, let

$$n_z = \left| \left\{ (i, v^{(1)}) : v^{(1)} \in \mathbb{F}_\ell^{D'}, \alpha_i W^{(1)} v^{(1)} + \alpha_i w = z \right\} \right|.$$

Thus, we have

$$\sum_z n_z = \frac{q-1}{\ell-1} \cdot \ell^{D'}$$

since there are $(q-1)/(\ell-1)$ choices for i and $\ell^{D'}$ choices for $v^{(1)}$.

Further, we have that

$$\sum_z \binom{n_z}{2} \leq \binom{(q-1)/(\ell-1)}{2}.$$

This is true because, on one hand, the left side counts the number of pairs $(i, v), (j, u)$ so that

$$\alpha_i W^{(1)} u + \alpha_i w = \alpha_j W^{(1)} v + \alpha_j w,$$

by caseworking on the value $z = \alpha_i W^{(1)} u + \alpha_i w = \alpha_j W^{(1)} v + \alpha_j w$. On the other hand, for any fixed i and j , there is at most one such pair (i, v) and (j, u) , so the total number of pairs is at most $\binom{(q-1)/(\ell-1)}{2}$.

For a uniform i and $v^{(1)} \sim \mathbb{F}_q^{D'}$, the vector $\alpha_i W^{(1)} v^{(1)} + \alpha_i w$ equals a vector $z \in \mathbb{F}_q^{D'}$ with probability proportional to n_z . Thus, by Lemma 3.5 with $N_1 = \frac{q-1}{\ell-1} \cdot \ell^{D'}$ and $N_2 = \binom{(q-1)/(\ell-1)}{2}$, we have

$$\begin{aligned} H_{v^{(1)} \sim \mathbb{F}_q^{D'}, i \sim [(q-1)/(\ell-1)]}(\alpha_i W^{(1)} v^{(1)} + \alpha_i w) &\geq \log \left(\frac{q-1}{\ell-1} \cdot \ell^{D'} \right) - \log \left(1 + \frac{\binom{(q-1)}{\ell-1} \cdot \left(\frac{q-1}{\ell-1} - 1 \right)}{\frac{q-1}{\ell-1} \cdot \ell^{D'}} \right) \\ &= \log \left(\frac{q-1}{\ell-1} \cdot \ell^{D'} \right) - \log \left(1 + \frac{q-\ell}{(\ell-1) \cdot \ell^{D'}} \right). \quad (2) \end{aligned}$$

We now show that (2) implies (1). Recall that $q = \ell^t$. We have

$$\begin{aligned}
1 + \frac{q - \ell}{(\ell - 1) \cdot \ell^{D'}} &< 1 + \frac{q - 1}{(\ell - 1)\ell^{D'}} \\
&= \frac{q - 1}{(\ell - 1)q^{0.1D'/D}} \cdot \left(\frac{(\ell - 1)q^{0.1D'/D}}{q - 1} + \left(\frac{q^{0.1/D}}{\ell} \right)^{D'} \right) \\
&\leq \frac{q - 1}{(\ell - 1)q^{0.1D'/D}} \cdot \left(\frac{(\ell - 1)q^{0.1}}{q - 1} + \left(\frac{\ell^{0.1}}{\ell} \right)^{D'} \right) \\
&\leq \frac{q - 1}{(\ell - 1)q^{0.1D'/D}} \cdot \left(\frac{(\ell - 1)\ell^{0.2}}{\ell^2 - 1} + \frac{1}{\ell^{0.9}} \right) \\
&\leq \frac{q - 1}{(\ell - 1)q^{0.1D'/D}} \cdot (0.4 + 0.6).
\end{aligned}$$

Therefore we conclude that

$$1 + \frac{q - \ell}{(\ell - 1) \cdot \ell^{D'}} < \frac{q - 1}{(\ell - 1)q^{0.1D'/D}}. \quad (3)$$

In the first inequality we used that $q - \ell < q - 1$. In the second inequality, we used that $D' \leq D$ and that $q^{1/D} \leq q^{t/D} \leq \ell$ (recall $t \leq D$). In the third inequality, we used that $\frac{x^{0.1}}{x-1}$ is decreasing for $x > 1$, that $\ell^2 \leq q$, and that $D' \geq 1$. In the fourth inequality we used $\ell \geq 2$ and that $\frac{(x-1)x^{0.2}}{x^2-1}$ and $x^{-0.9}$ are decreasing for $x \geq 2$.

Combining (3) with (2) proves (1), so

$$H(A\tau) > \log(q^{0.1D'/D}\ell^{D'})$$

which gives us

$$H_q(A\tau) > \log_q(q^{0.1D'/D}\ell^{D'}) = D' \left(\log_q(\ell) + \frac{1}{10D} \right)$$

These computations holds for any matrix A of rank $D' \leq D$, so we have that τ is *not* $(\log_q(\ell) + \frac{1}{10D})$ -implicitly rare. \square

This concludes the proof when $\rho = 0$; we continue to the case when $\rho > 0$.

Proof of Lemma 3.4 for $\rho > 0$. We need to show that for any $A \in \mathbb{F}_q^{L' \times L}$ such that $AG \in \mathbb{F}_q^{L' \times D}$ has rank $D' \leq D$, the entropy of $H_q(A\tau)$ is at least $D'(\rho + (1 - \rho)\log_q \ell + \frac{1-\rho}{10D})$. To see this first note that

$$\tau = \begin{cases} G\alpha_i v, & \text{with probability } 1 - \rho, \text{ where } i \sim [1, \dots, (q-1)/(\ell-1)], v \sim \mathbb{F}_\ell^D \\ Gw, & \text{with probability } \rho, \text{ where } w \sim \mathbb{F}_q^D \end{cases}$$

Define τ_1 as the distribution of $G\alpha_i v$ and τ_2 as the distribution of Gw . Now note that $\tau =$

$(1 - \rho)\tau_1 + \rho\tau_2$. Let $g(x) = -x \log x$. Since $g(x)$ is concave, we have

$$\begin{aligned} H_q(A\tau) &= \sum_{v \in \mathbb{F}_q^L} g(\Pr_{A\tau}(v)) \\ &= \sum_{v \in \mathbb{F}_q^L} g((1 - \rho) \cdot \Pr_{A\tau_1}(v) + \rho \cdot \Pr_{A\tau_2}(v)) \\ &\geq \sum_{v \in \mathbb{F}_q^L} (1 - \rho) \cdot g(\Pr_{A\tau_1}(v)) + \rho \cdot g(\Pr_{A\tau_2}(v)) \\ &= (1 - \rho) \cdot H_q(A\tau_1) + \rho \cdot H_q(A\tau_2). \end{aligned}$$

From the $\rho = 0$ case we already know that

$$H_q(A\tau_1) \geq D' \left(\log_q(\ell) + \frac{1}{10D} \right)$$

When w is uniform on \mathbb{F}_q^D , then AGw is uniformly distributed in the \mathbb{F}_q -span of AG which has rank D' , so $H_q(A\tau_2) = D'$. We thus have

$$H_q(A\tau) \geq (1 - \rho)H_q(A\tau_1) + \rho H_q(A\tau_2) \geq (1 - \rho) \cdot D' \left(\log_q(\ell) + \frac{1}{10D} \right) + \rho D',$$

as desired. □

4 Lower bounds for list-decoding with errors

Our main theorem in this section is the following.

Theorem 4.1. *Fix a prime power q , fix $p \in (0, 1 - \frac{1}{q})$, and fix $\delta \in (0, 1)$. There exists $\varepsilon_{p,q,\delta} > 0$ such that for all $\varepsilon \in (0, \varepsilon_{p,q,\delta})$ and n sufficiently large, a random linear code in \mathbb{F}_q^n of rate $1 - h_q(p) - \varepsilon$ is not $\left(p, \left\lfloor \frac{h_q(p)}{\varepsilon} - \delta \right\rfloor \right)$ -list-decodable with probability $1 - q^{-\Omega(n)}$.*

4.1 Proof of Theorem 4.1

Our proof of Theorem 4.1 below follows the same outline as the proof of Theorem 3.1 above. We first define a bad distribution τ in Definition 4.2; then we will show that it is bad in Proposition 4.3; then we will show that it is implicitly rare (and hence abundant by Theorem 2.5) in Lemma 4.4. Finally we will prove Theorem 4.1 from these pieces.

Below, we let $\text{Bernoulli}_q(p)$ be the distribution that returns $0 \in \mathbb{F}_q$ with probability $1 - p$ and any other element of \mathbb{F}_q with probability $\frac{p}{q-1}$.

Definition 4.2 (The bad distribution τ for list-decoding lower bounds). *Let $p \in (0, 1 - \frac{1}{q})$ and $\delta > 0$. Choose $L > 0$. Define the distribution τ on \mathbb{F}_q^L as the distribution of the random vector $u + \alpha \mathbf{1}_L$, where $u \sim \text{Bernoulli}_q(p)^L$, and α is sampled independently and uniformly from \mathbb{F}_q .*

First, we observe that τ is indeed bad, in the sense that it provides a counter-example to list-decodability.

Proposition 4.3 (τ is bad). *Let τ be as in Definition 4.2. Let $\mathcal{C} \subseteq \mathbb{F}_q^n$ and let $M \in \mathcal{M}_{n,\tau}$. If $M \subseteq \mathcal{C}$, then \mathcal{C} is not (p, L) -list-decodable.*

Proof. Let $M \in \mathcal{M}_{n,\tau}$. We want to show that the columns of M all lie in a single ball of radius pn .

By definition of τ and $\mathcal{M}_{n,\tau}$, we may write the j -th row of M as $u^{(j)} + \alpha_j \mathbf{1}_L$, so that the empirical distribution of the pairs $(u^{(j)}, \alpha_j)_{1 \leq j \leq n}$ is $\text{Bernoulli}_q(p)^L \times \text{Uniform}(\mathbb{F}_q)$.¹¹

For any $i \in [L]$, the number of $j \in [n]$ such that $M_{i,j} = u_i^{(j)} + \alpha_j \neq \alpha_j$ is exactly the number of times $u_i^{(j)} \neq 0$, which is pn , since $u_i^{(j)}$ is distributed as $\text{Bernoulli}_q(p)$. Thus, each column $M_{i,*}$ of M has distance at most pn from the word $(\alpha_1, \dots, \alpha_n)$, so that any code containing M has L codewords in a ball of radius pn and hence is not (p, L) -list-decodable. \square

Next, we show that τ is appropriately implicitly rare for large enough L .

Lemma 4.4. *Let $p \in (0, 1 - \frac{1}{q})$ and let $\delta > 0$. There exists $L_{p,q,\delta}$ such that, for $L \geq L_{p,q,\delta}$, the distribution τ given in Definition 4.2 is not $(h_q(p) + \frac{h_q(p)}{L+\delta})$ -implicitly rare.*

We prove Lemma 4.4 in Section 4.2 below. Before we prove Lemma 4.4, we show how to use it to prove Theorem 4.1.

Proof of Theorem 4.1, assuming Lemma 4.4. Let $L_{p,q,\delta/2}$ be as in Lemma 4.4 and choose $\varepsilon_{p,q,\delta} \stackrel{\text{def}}{=} \frac{h_q(p)}{L_{p,q,\delta/2}+1}$. Fix $\varepsilon < \varepsilon_{p,q,\delta}$. Let $L = \lfloor \frac{h_q(p)}{\varepsilon} - \delta \rfloor$. Let τ be as in Definition 4.2 with this choice of L . By Lemma 4.4, as $L \geq L_{p,q,\delta/2}$, τ is not $(h_q(p) + \frac{h_q(p)}{L+\delta/2})$ -implicitly rare. Thus, as $\varepsilon \leq \frac{h_q(p)}{L+\delta} < \frac{h_q(p)}{L+\delta/2}$, there is some constant $c_{p,q,\varepsilon} > 0$ so that τ is not $(h_q(p) + \varepsilon + c_{p,q,\varepsilon})$ -implicitly rare.

Then Theorem 2.5 with $\eta = c_{p,q,\varepsilon}$ tells us that, for n sufficiently large, a random linear code of rate $1 - (h_q(p) + \varepsilon + c_{p,q,\varepsilon}) + c_{p,q,\varepsilon} = 1 - h_q(p) - \varepsilon$ contains L codewords given by some matrix $M \in \mathcal{M}_{n,\tau}$ with probability at least $1 - q^{-\Omega_{p,q,\varepsilon}(n)}$.

Finally, Proposition 4.3 implies that \mathcal{C} is not (p, L) -list-decodable. Our choice of L proves the theorem. \square

4.2 Proof of Lemma 4.4

In this section we prove Lemma 4.4, which completes the proof of Theorem 4.1. To prove Lemma 4.4 we need to prove that $A\tau$ has high entropy for any matrix A . We begin with the following lemma, which essentially shows that this is true when A is either the identity I_L or the identity with an additional all-1s column $[I_L | \mathbf{1}_L]$.

Lemma 4.5. *Let q be a prime power, $p \in (0, 1 - \frac{1}{q})$, $p' \in [p, 1 - \frac{1}{q}]$, and $\delta > 0$. There exists $L_{p,q,\delta}$ such that, for all $L \geq L_{p,q,\delta}$ and $0 \leq d \leq L$, the following holds. Let v be a vector sampled from $\text{Bernoulli}_q(p)^d$ and let w be a vector independently sampled from $\text{Bernoulli}_q(p') \cdot \mathbf{1}_d$. Then*

$$H_q(v + w) \geq d \cdot \left(h_q(p) + \frac{h_q(p)}{L + \delta} \right). \quad (4)$$

¹¹This is without loss of generality: if not, as per Remark 2.3, we can associate pairs with rows so that the empirical distribution is close to $\text{Bernoulli}_q(p)^L \times \text{Uniform}(\mathbb{F}_q)$ up to an additive factors that are $o(1)$ as $n \rightarrow \infty$. After adjusting parameters, this has a negligible effect on the analysis and final result.

Proof. If $d = 0$, the assertion is trivial, so assume $d \geq 1$.

We will bound $H_q(v + w)$ in two cases, one when d is small (relative to L) and one when d is large. (The precise definitions of “small” and “large” will be determined below.)

First we consider the case where d is small. We have (for any d) that

$$\begin{aligned}
H_q(v + w) &= H_q(v_1 + w_1, v_2 + w_2, \dots, v_d + w_d) \\
&= H_q(v_2 + w_1, \dots, v_d + w_1 | v_1 + w_1) + H_q(v_1 + w_1) \\
&\geq H_q(v_2 + w_1, \dots, v_d + w_1 | v_1, w_1) + H_q(v_1 + w_1) \\
&= H_q(v_2, \dots, v_d) + H_q(v_1 + w_1)
\end{aligned} \tag{5}$$

The second equality uses the definition of conditional entropy, and that $w_1 = w_2 = \dots = w_d$. The inequality follows from the data processing inequality.

Now, $v_1 + w_1$ is nonzero if $v_1 = 0$ and $w_1 \neq 0$, if $v_1 \neq 0$ and $w_1 = 0$, or if $v_1, w_1 \neq 0$ and $v_1 + w_1 \neq 0$. This happens with probability $p^* = (1 - p')p + (1 - p)p' + \frac{(q-2)pp'}{q-1}$. In the case that $w_1 + v_1$ is nonzero, then each nonzero element of \mathbb{F}_q has equal probability. Thus $v_1 + w_1$ is distributed as $\text{Bernoulli}_q(p^*)$. One can check that $H_q(\text{Bernoulli}_q(p^*)) = h_q(p^*)$, so from (5) we have

$$\begin{aligned}
H_q(v + w) &\geq H_q(v_2, \dots, v_d) + h_q(p^*) \\
&= (d - 1) \cdot h_q(p) + h_q(p^*).
\end{aligned} \tag{6}$$

Since

$$p < 2p(1 - p) + \frac{(q - 2)p^2}{q - 1} \leq p^* \leq 1 - \frac{1}{q},$$

we have $h_q(p^*) \geq h_q(p) + \varepsilon_{p,q}$ for some $\varepsilon_{p,q} > 0$ depending only on p and q . The second inequality follows from the fact that p^* increases with p' and $p' \geq p$. Thus, when $d \leq \varepsilon_{p,q} \cdot L$, (6) implies that

$$H_q(v + w) \geq d \cdot h_q(p) + \varepsilon_{p,q} \geq d \cdot \left(h_q(p) + \frac{1}{L} \right) > d \cdot \left(h_q(p) + \frac{h_q(p)}{L + \delta} \right),$$

where in the last line we have used that $\delta > 0$ and $h_q(p) < 1$. This bounds $H_q(v + w)$ in the case when d is “small,” meaning that $d < \varepsilon_{p,q} \cdot L$.

Next we handle the case when d is “large.” We have (for any d) that

$$\begin{aligned}
H_q(v + w) &= H_q(v + w | w) + H_q(w) - H_q(w | v + w) \\
&= d \cdot h_q(p) + h_q(p) - H_q(w | v + w).
\end{aligned}$$

Let $w_0 \in \mathbb{F}_q$ be the majority element of $v + w \in \mathbb{F}_q^d$. In expectation a $1 - p > \frac{1}{q}$ fraction of the d coordinates of v are 0. Similarly, for any nonzero $\alpha \in \mathbb{F}_q$, in expectation a $\frac{p}{q-1} < \frac{1}{q}$ fraction of the coordinates of v are α .

By Hoeffding’s inequality, for any nonzero $\alpha \in \mathbb{F}_q$,

$$\Pr \left[|\{i : v_i = \alpha\}| > \frac{d}{q} \right] \leq 2 \exp \left(-2d \left(\frac{1}{q} - \frac{p}{q-1} \right) \right) = \exp(-\Omega_{p,q}(d))$$

and similarly

$$\Pr \left[|\{i : v_i = 0\}| < \frac{d}{q} \right] \leq 2 \exp \left(-2d \left(1 - p - \frac{1}{q} \right) \right) = \exp(-\Omega_{p,q}(d)).$$

If it is the case that there are more than d/q zeros in \vec{v} and fewer than d/q of any other element in \mathbb{F}_q , then w_0 will be the same as the only element that appears in w , and we will have $w = w_0 \mathbf{1}_d$. Thus, the probability that w is not equal to $w_0 \mathbf{1}_d$ is at most

$$p_{err} \stackrel{\text{def}}{=} \Pr [w \neq w_0 \mathbf{1}_d] \leq 2q \cdot \exp(-\Omega_{p,q}(d)) = \exp(-\Omega_{p,q}(d)).$$

By Fano's inequality, as w takes at most q values, we have

$$H_q(w|v+w) = \frac{1}{\log q} H(w|v+w) \leq \frac{1}{\log q} (h(p_{err}) + p_{err} \cdot \log(q-1)) \leq \exp(-\Omega_{p,q}(d)).$$

Thus, there exists some $d_{p,q,\delta}$ such that, for $d \geq d_{p,q,\delta}$, we have $H_q(w|v+w) \leq \frac{\delta h_q(p)}{d+\delta}$, in which case

$$\begin{aligned} H_q(v+w) &= d \cdot h_q(p) + h_q(p) - H_q(w|v+w) \\ &\geq d \cdot h_q(p) + \frac{d}{d+\delta} \cdot h_q(p) \\ &\geq d \cdot \left(h_q(p) + \frac{h_q(p)}{L+\delta} \right). \end{aligned}$$

This completes the case where d is "large," meaning that $d \geq d_{p,q,\delta}$.

We have shown that (4) holds when $d \leq \varepsilon_{p,q} \cdot L$ and when $d \geq d_{p,q,\delta}$. Thus, for

$$L \geq d_{p,q,\delta} / \varepsilon_{p,q} \stackrel{\text{def}}{=} L_{p,q,\delta},$$

we have that (4) always holds, as desired. \square

Using Lemma 4.5, we may now prove Lemma 4.4 which says that τ is not implicitly rare.

Proof of Lemma 4.4. Let $L_{p,q,\delta}$ be as in Lemma 4.5. Let $L \geq L_{p,q,\delta}$, and let τ be the corresponding distribution in the lemma statement.¹² Fix a matrix A of rank L' . As τ is supported on \mathbb{F}_q^L , the rank of $A\tau$ is L' . We show that $H(A\tau) \geq L' \cdot (h_q(p) + \frac{h_q(p)}{L+\delta})$. At a high level, our strategy is to decompose the distribution $A\tau$ into several distributions that each have the set up of Lemma 4.5. Furthermore, this decomposition has enough conditional independence that the entropy of $A\tau$ can be bounded below by the sum of the entropies of the smaller distributions, which we can bound below by Lemma 4.5.

Since removing rows that are linearly dependent on other rows does not change $H_q(A\tau)$, we may assume that A has exactly L' rows. Since permuting the coordinates of τ yields the same distribution τ , permuting the columns of A does not change the entropy $H_q(A\tau)$. Furthermore, if B is invertible, $H_q(BA\tau) = H_q(A\tau)$. Thus, by running Gaussian elimination on the rows of A , we may assume without loss of generality that

$$A = \left[\begin{array}{c|ccc|c} I_{L'} & \begin{array}{c} | \\ v^{(1)} \\ | \end{array} & \begin{array}{c} | \\ v^{(2)} \\ | \end{array} & \begin{array}{c} \cdots \\ \cdots \\ \cdots \end{array} & \begin{array}{c} | \\ v^{(k)} \\ | \end{array} \end{array} \right]$$

¹²As in Lemma 3.4, we treat the output of τ as a column vector.

where $v^{(1)}, \dots, v^{(k)} \in \mathbb{F}_q^{L'}$ and $k = L - L'$. Let a sample from τ be given by

$$\begin{bmatrix} x_1 \\ \vdots \\ x_{L'} \\ y_1 \\ \vdots \\ y_k \end{bmatrix} + \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \cdot y_{k+1}.$$

where $x_1, \dots, x_{L'}, y_1, \dots, y_k \sim \text{Bernoulli}_q(p)$ and $y_{k+1} \sim \text{Bernoulli}_q(1 - \frac{1}{q})$. (Note that this means that y_{k+1} is uniform on \mathbb{F}_q .) Then $A\tau$ is given by

$$\begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix} + \sum_{i=1}^{k+1} \begin{bmatrix} | \\ v^{(i)} \\ | \end{bmatrix} \cdot y_i$$

where we let $v^{(k+1)}$ be the product $A \cdot \mathbf{1}_L \in \mathbb{F}_q^{L'}$.

By definition of $v^{(k+1)}$, for any coordinate $i \notin \bigcup_{j=1}^k \text{supp}(v^{(j)})$, we have $i \in \text{supp}(v^{(k+1)})$. Thus, $\bigcup_{i=1}^{k+1} \text{supp}(v^{(i)}) = [L']$. For $i = 1, \dots, k+1$, let $J_i = \text{supp}(v^{(i)}) \setminus (\bigcup_{j=1}^{i-1} J_j)$ (when $i = 1$ the union is the empty set), so that J_1, \dots, J_{k+1} form a partition of $[L']$. Recall that the notation $x_J \in \mathbb{F}_q^{|J|}$ denotes the vector $(x_i)_{i \in J}$ with coordinates from J in increasing order. We have

$$\begin{aligned} H_q(A\tau) &= H_q(A\tau | x_{J_{k+1}}, y_{k+1}) + I_q(A\tau; x_{J_{k+1}}, y_{k+1}) \\ &= H_q(A\tau | x_{J_k}, x_{J_{k+1}}, y_k, y_{k+1}) + I_q(A\tau; x_{J_k}, y_k | x_{J_{k+1}}, y_{k+1}) + I_q(A\tau; x_{J_{k+1}}, y_{k+1}) \end{aligned}$$

Continuing, we have

$$\begin{aligned} H_q(A\tau) &= H_q(A\tau | x_{J_1}, \dots, x_{J_{k+1}}, y_1, \dots, y_{k+1}) + \sum_{i=1}^{k+1} I_q(A\tau; x_{J_i}, y_i | x_{J_{i+1}}, \dots, x_{J_{k+1}}, y_{i+1}, \dots, y_{k+1}) \\ &= \sum_{i=1}^{k+1} I_q(A\tau; x_{J_i}, y_i | x_{J_{i+1}}, \dots, x_{J_{k+1}}, y_{i+1}, \dots, y_{k+1}), \end{aligned}$$

where the second inequality uses that $A\tau$ is completely determined by x and y and hence by $x_{J_1}, \dots, x_{J_{k+1}}, y_1, \dots, y_{k+1}$. Above, when $i = k+1$, the summand is simply $I_q(A\tau; x_{J_{k+1}}, y_{k+1})$.

Thus, by the data processing inequality, we have

$$\begin{aligned}
H_q(A\tau) &\geq \sum_{i=1}^{k+1} I_q((A\tau)_{J_i}; x_{J_i}, y_i | x_{J_{i+1}}, \dots, x_{J_{k+1}}, y_{i+1}, \dots, y_{k+1}) \\
&= \sum_{i=1}^{k+1} I_q \left(x_{J_i} + \mathbf{1}_{|J_i|} \cdot y_i + \sum_{j>i} (v^{(j)} y_j)_{J_i}; x_{J_i}, y_i \middle| x_{J_{i+1}}, \dots, x_{J_{k+1}}, y_{i+1}, \dots, y_{k+1} \right) \\
&= \sum_{i=1}^{k+1} I_q(x_{J_i} + \mathbf{1}_{|J_i|} \cdot y_i; x_{J_i}, y_i | x_{J_{i+1}}, \dots, x_{J_{k+1}}, y_{i+1}, \dots, y_{k+1}) \\
&= \sum_{i=1}^{k+1} I_q(x_{J_i} + \mathbf{1}_{|J_i|} \cdot y_i; x_{J_i}, y_i) \\
&= \sum_{i=1}^{k+1} H_q(x_{J_i} + \mathbf{1}_{|J_i|} \cdot y_i)
\end{aligned}$$

The first equality uses that $v^{(1)}, \dots, v^{(i-1)}$ have no support in J_i (by definition of J_i). The second equality uses that y_{i+1}, \dots, y_{k+1} have already been conditioned on. The third equality uses that the x_i 's and y_i 's are independent. The last equality uses that $H_q(x_{J_i} + \mathbf{1}_{|J_i|} \cdot y_i | x_{J_i}, y_i) = 0$. As $L \geq L_{p,q,\delta}$, we may apply Lemma 4.5 with $v = x_{J_i}$ and $w = \mathbf{1}_{|J_i|} \cdot y_i$ and $d = |J_i|$. This gives

$$\begin{aligned}
H_q(A\tau) &\geq \sum_{i=1}^{k+1} H_q(x_{J_i} + \mathbf{1}_{|J_i|} \cdot y_i) \\
&\geq \sum_{i=1}^{k+1} |J_i| \cdot \left(h_q(p) + \frac{h_q(p)}{L + \delta} \right) \\
&= L' \cdot \left(h_q(p) + \frac{h_q(p)}{L + \delta} \right),
\end{aligned}$$

as desired. □

5 Upper bounds for average-radius list-decoding over \mathbb{F}_2

In this section we prove the following theorem. Recall that we abbreviate $h(p) = h_2(p)$.

Theorem 5.1. *Let $n \in \mathbb{N}$. Let $p \in (0, \frac{1}{2})$ and $R = 1 - h(p) - \varepsilon$, where $0 < \varepsilon < 1 - h(p)$. Let $L = \lfloor \frac{h(p)}{\varepsilon} + 2 \rfloor$. Then, a random linear code $\mathcal{C} \leq \mathbb{F}_2^n$ of rate R is (p, L) -average-radius list-decodable with probability $1 - 2^{-\Omega_{p,\varepsilon}(n)}$.*

Recall from the introduction that, following the techniques in [GHSZ02] and [LW18], we imagine sampling independent and uniform vectors b_1, \dots, b_k and constructing the “intermediate” random linear codes $\mathcal{C}_i = \text{span}\{b_1, \dots, b_i\}$. A potential function based argument is used to show that, with high probability, each of these intermediate codes is indeed (p, L) -average-radius list-decodable; in particular, this is true for \mathcal{C}_k .

Before discussing our potential function, we first briefly review the techniques of [GHSZ02] and [LW18]; in particular, we describe the potential function they use. First, for a code \mathcal{C} and a vector $x \in \mathbb{F}_2^n$, we define

$$L_{\mathcal{C}}(x) := |\{c \in \mathcal{C} : \delta(x, c) \leq p\}| .$$

In [GHSZ02], the authors define

$$S_{\mathcal{C}} := \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} 2^{\varepsilon n L_{\mathcal{C}}(x)}$$

and observe that, for any $b_1, \dots, b_i \in \mathbb{F}_2^n$,

$$\mathbb{E}_{b_{i+1} \sim \mathbb{F}_2^n} [S_{\mathcal{C}_i + \{0, b_{i+1}\}}] = S_{\mathcal{C}_i}^2 ,$$

where we recall $\mathcal{C}_i = \text{span}\{b_1, \dots, b_i\}$. That is, the potential function squares in expectation, so the probabilistic method guarantees that we can choose some b_{i+1} for which $S_{\mathcal{C}_i + \{0, b_{i+1}\}} \leq S_{\mathcal{C}_i}^2$. Thus, for some choice of b_1, \dots, b_k , one has $S_{\mathcal{C}_k} \leq (S_{\{0\}})^{2^k}$.

In [LW18], the definition of $S_{\mathcal{C}}$ is slightly modified:

$$S_{\mathcal{C}} := \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} 2^{\frac{\varepsilon n L_{\mathcal{C}}(x)}{1+\varepsilon}} .$$

This little bit of extra room allows to show that, in fact, with high probability over the choice of b_{i+1} , $S_{\mathcal{C}_i + \{0, b_{i+1}\}} \leq S_{\mathcal{C}_i}^2$. By a union bound, it follows that with high probability, $S_{\mathcal{C}_k} \leq (S_{\{0\}})^{2^k}$.

In either case, to conclude the proof, one observes the bound¹³ $S_{\{0\}} \leq 1 + 2^{-n(1-h(p)-\varepsilon)}$ and then uses

$$S_{\mathcal{C}_k} \leq (S_{\{0\}})^{2^k} \leq (1 + 2^{-n(1-h(p)-\varepsilon)})^{2^k} \leq \exp\left(2^{k-n(1-h(p)-\varepsilon)}\right) \leq O(1)$$

for k chosen as above.

5.1 Alterations for average-radius list-decoding

While this argument analyzes the (absolute-radius) list-decodability of random linear codes very effectively, it is not immediately clear how to generalize the argument to study average-radius list-decodability. We now introduce the additional ideas we need to derive Theorem 5.1. We will fix a threshold parameter $\lambda \in (p, \frac{1}{2})$ for which $h(\lambda) < 1 - R = h(p) + \varepsilon$, to be determined later, and define

$$\eta \stackrel{\text{def}}{=} 1 - R - h(\lambda).$$

We define the function $M_{R,\lambda} : [0, 1] \rightarrow \mathbb{R}$ by

$$M_{R,\lambda}(\gamma) := \begin{cases} 1 - R - h(\gamma) & \text{if } \gamma < \lambda \\ 0 & \text{if } \gamma \geq \lambda \end{cases} .$$

¹³Actually, for the potential function in [LW18], one has $S_{\{0\}} \leq 1 + 2^{-n(1-h(p)-\frac{\varepsilon}{1+\varepsilon})}$, but this difference does not matter for the conclusion.

Remark 5.2. One can think of this quantity as a sort of “normalized entropy change” up to the threshold λ . Recalling that $1 - R = h(p) + \varepsilon$, if $\gamma < \lambda$, then

$$M_{R,\lambda}(\gamma) \approx \frac{1}{n} (h(p) - h(\lambda)) \approx \log \left(\frac{|B^n(p)|}{|B^n(\gamma)|} \right),$$

where $B^n(p)$ denotes the Hamming ball in \mathbb{F}_2^n of radius p . Hence, $M_{R,\lambda}(\gamma)$ is something like a normalized “surprise” an observer would experience if they are expecting a random vector of weight $\leq p$ and see a vector of weight $\leq \gamma$.

For a linear code $\mathcal{C} \leq \mathbb{F}_2^n$ and $x \in \mathbb{F}_2^n$ we define

$$L_{\mathcal{C},R,\lambda}(x) := \sum_{y \in \mathcal{C}} M_{R,\lambda}(\delta(x, y)).$$

This is intuitively the “smoothed-out” list-size of x , where nearby codewords are weighted more heavily than far away codewords, and the weighting is given by the “entropy change” implied by the distance from x to y .

Next, we define

$$A_{\mathcal{C},R,\lambda}(x) := 2^{\frac{nL_{\mathcal{C},R,\lambda}(x)}{1+\eta}}$$

and

$$S_{\mathcal{C},R,\lambda} := \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} A_{\mathcal{C},R,\lambda}(x).$$

The quantity $S_{\mathcal{C},R,\lambda}$ is the potential function we will analyze.

5.2 Proof of Theorem 5.1

In this subsection we prove Theorem 5.1. The quantities R and λ (and hence $\eta = 1 - R - h(\lambda)$) will be fixed throughout—although the precise value of λ will be determined later—and so we will suppress their dependence and simply write $M(x)$, $L_{\mathcal{C}}(x)$, $A_{\mathcal{C}}(x)$ and $S_{\mathcal{C}}$.

First, we observe that the following analog of [LW18, Lemma 3.2] holds. The proof is a simple adaptation of theirs (which in turn follows [GHSZ02]).

Lemma 5.3. For all $\mathcal{C} \leq \mathbb{F}_2^n$ and $b \in \mathbb{F}_2^n$,

$$\begin{aligned} L_{\mathcal{C}+\{0,b\}}(x) &\leq L_{\mathcal{C}}(x) + L_{\mathcal{C}}(x+b), \\ A_{\mathcal{C}+\{0,b\}}(x) &\leq A_{\mathcal{C}}(x) \cdot A_{\mathcal{C}}(x+b). \end{aligned}$$

Moreover, equality holds if and only if $b \notin \mathcal{C}$.

Proof. We have

$$\begin{aligned} L_{\mathcal{C}+\{0,b\}}(x) &\leq \sum_{y \in \mathcal{C}} M(\delta(x, y)) + M(\delta(x, y+b)) \\ &= \sum_{y \in \mathcal{C}} M(\delta(x, y)) + M(\delta(x+b, y)) = L_{\mathcal{C}}(x) + L_{\mathcal{C}}(x+b), \end{aligned}$$

and equality holds in the first line if and only if $\mathcal{C} \cap (b + \mathcal{C}) = \emptyset$, or, equivalently, $b \notin \mathcal{C}$. The second inequality of the lemma statement follows from the first. \square

Next, we bound $S_{\{0\}}$. We have

$$\begin{aligned} S_{\{0\}} &\leq 1 + 2^{-n} \sum_{\substack{x \in \mathbb{F}_2^n \\ \text{wt}(x) \leq \lambda}} 2^{\frac{n \cdot (1-R-h(\text{wt}(x)))}{1+\eta}} \\ &\leq 1 + \sum_{i=0}^{\lfloor \lambda n \rfloor} 2^{-n \left(1-h(i/n) - \frac{h(\lambda)+\eta-h(i/n)}{1+\eta}\right)}. \end{aligned}$$

As this sum is dominated by its last term, we deduce

$$S_{\{0\}} \leq 1 + (\lambda n) 2^{-n \left(1-h(\lambda) - \frac{\eta}{1+\eta}\right)}. \quad (7)$$

From here, we can combine Lemma 5.3 and (7) to deduce

Lemma 5.4. *Let $p \in (0, \frac{1}{2})$ and $R = 1 - h(p) - \varepsilon$ for $0 < \varepsilon < 1 - h(p)$. Let $\mathcal{C}_{Rn} \leq \mathbb{F}_2^n$ be a random linear code of rate R . Then $S_{\mathcal{C}_{Rn}} \leq 2$ with probability at least $1 - \exp(-\Omega_\eta(n))$.*

The proof of this lemma is completely analogous to that of [LW18, Lemma 3.3]. One only needs to be careful about the growth rate of $S_{\mathcal{C}}$. In particular, this proof crucially uses that η is positive. We again choose vectors b_1, \dots, b_{Rn} independently and uniformly at random. If $\mathcal{C}_i = \text{span}\{b_1, \dots, b_i\}$, we need “in expectation” that $S_{\mathcal{C}_i} \leq 1 + 2^{-\Omega_\eta(n)}$ for all i for the error bounds to succeed. As we expect the $o(1)$ term to roughly double, we need $2^{Rn} \cdot (S_{(0)} - 1) = 2^{-n(\eta - \frac{\eta}{1+\eta})} \leq 2^{-\Omega_\eta(n)}$. For completeness, we provide the proof of Lemma 5.4 in Appendix A.

Thus, in order to conclude Theorem 5.1, we are simply required to demonstrate that $S_{\mathcal{C}} \leq 2$ implies that \mathcal{C} is (p, L) -average-radius list-decodable: this is the crux of our contribution. The main lemma we require is the following.

Lemma 5.5. *Let $\mathcal{C} \leq \mathbb{F}_2^n$ be a linear code of rate R such that $S_{\mathcal{C}} \leq 2$. Then, for all $x \in \mathbb{F}_2^n$ and $D \subseteq \mathcal{C} \cap B(x, \lambda)$, it holds that*

$$\sum_{y \in D} h(\delta(x, y)) \geq (|D| - 1 - \eta)(1 - R) - \frac{1 + \eta}{n}.$$

Proof. First, observe that for any $x \in \mathbb{F}_2^n$,

$$\begin{aligned} L_{\mathcal{C}}(x) &\geq \sum_{y \in D} ((1 - R) - h(\delta(x, y))) = |D|(1 - R) - \sum_{y \in D} h(\delta(x, y)) \\ \text{so } \log A_{\mathcal{C}}(x) &\geq n \frac{|D|(1 - R) - \sum_{y \in D} h(\delta(x, y))}{1 + \eta}. \end{aligned} \quad (8)$$

Next, as $\delta(x, y) = \delta(x + z, y + z)$ for any $z \in \mathbb{F}_2^n$, we have, for any $x \in \mathbb{F}_2^n$ and $c \in \mathcal{C}$, that $L_{\mathcal{C}}(x) = L_{\mathcal{C}}(x + c)$ and hence $A_{\mathcal{C}}(x) = A_{\mathcal{C}}(x + c)$. Thus, $\max_{x \in \mathbb{F}_2^n} A_{\mathcal{C}}(x)$ is attained at at least $|\mathcal{C}|$ different values of x , so

$$S_{\mathcal{C}} = \frac{1}{2^n} \sum_{x \in \mathbb{F}_2^n} A_{\mathcal{C}}(x) \geq \frac{1}{2^n} \cdot |\mathcal{C}| \cdot \max_{x \in \mathbb{F}_2^n} A_{\mathcal{C}}(x) = 2^{-(1-R)n} \cdot \max_{x \in \mathbb{F}_2^n} A_{\mathcal{C}}(x).$$

Combining this with (8), we have, for any $x \in \mathbb{F}_2^n$,

$$\begin{aligned}
1 &\geq \log S_{\mathcal{C}} \geq -(1-R)n + \log(A_{\mathcal{C}}(x)) \\
&\geq n \cdot \left(-(1-R) + \frac{|D|(1-R) - \sum_{y \in D} h(\delta(x, y))}{1+\eta} \right) \\
&= n \cdot \frac{(|D| - 1 - \eta)(1-R) - \sum_{y \in D} h(\delta(x, y))}{1+\eta}.
\end{aligned}$$

Rearranging yields the lemma. \square

We may now conclude Theorem 5.1.

Proof of Theorem 5.1. Since $L = \lfloor \frac{h(p)}{\varepsilon} + 2 \rfloor > \frac{h(p)}{\varepsilon} + 1 = \frac{1-R}{\varepsilon}$, there exists $\eta > 0$ small enough so that for all sufficiently large n

$$L > \frac{1-R + \eta + \frac{1+\eta}{n}}{\varepsilon - \eta}. \quad (9)$$

Thus, we define λ so that η (which we defined as $\eta = 1 - R - h(\lambda)$) satisfies (9). Let \mathcal{C} be a random linear code of rate R . Due to Lemma 5.4, the conclusion of Lemma 5.5, holds with probability $1 - 2^{-\Omega_{\eta}(n)}$ for \mathcal{C} . It remains to show that, assuming n is sufficiently large, any code \mathcal{C} satisfying the conclusion of Lemma 5.5 is (p, L) -average-radius list-decodable.

Let $x \in \mathbb{F}_2^n$ and $\Lambda \subseteq \mathcal{C}$ such that $|\Lambda| = L$; our goal is to show that, for all such x and Λ ,

$$\frac{1}{L} \sum_{y \in \Lambda} \delta(x, y) > p. \quad (10)$$

Let

$$D = \{y \in \Lambda : \delta(x, y) \leq \lambda\}$$

and define

$$h^*(\alpha) = \begin{cases} h(\alpha) & \text{if } \alpha \leq \frac{1}{2} \\ 1 & \text{if } \alpha > \frac{1}{2} \end{cases}.$$

Now,

$$\sum_{y \in \Lambda} h^*(\delta(x, y)) \geq \sum_{y \in D} h(\delta(x, y)) + (L - |D|)h(\lambda) \quad (11)$$

$$\geq (|D| - 1 - \eta)(1-R) + (L - |D|)(1-R - \eta) - \frac{1+\eta}{n} \quad (12)$$

$$= (1-R) \cdot (L-1) - \eta \cdot (1-R) - \eta \cdot (L-|D|) - \frac{1+\eta}{n}$$

$$\geq (1-R) \cdot (L-1) - \eta \cdot (L+1) - \frac{1+\eta}{n}$$

$$= (1-R)L - (1-R) - \eta \cdot (L+1) - \frac{1+\eta}{n}$$

$$= Lh(p) - (1-R) - (L+1)\eta + L\varepsilon - \frac{1+\eta}{n} \quad (13)$$

$$= Lh(p) - (1-R + \eta) + (L-1)(\varepsilon - \eta) - \frac{1+\eta}{n}$$

$$> Lh(p). \quad (14)$$

Here, the Inequality (11) holds because $h^*(\alpha) > h(\lambda)$ for all $\alpha > \lambda$; Inequality (12) is the conclusion of Lemma 5.5; Equality (13) follows from the fact that $R = 1 - h(p) - \varepsilon$; and Inequality (14) follows from (9). Thus, we deduce

$$\frac{1}{L} \sum_{y \in \Lambda} h^*(\delta(x, y)) > h(p). \quad (15)$$

Since h^* is concave,

$$h^* \left(\frac{1}{L} \sum_{y \in \Lambda} (\delta(x, y)) \right) \geq \frac{1}{L} \sum_{y \in \Lambda} h^*(\delta(x, y)),$$

and so (10) follows from (15), the monotonicity of h^* and the fact that $h^*(p) = h(p)$. \square

Remark 5.6. *Just as the argument in [LW18] generalizes easily to the case of rank-metric codes, the same holds for the argument given above. Briefly, a rank-metric code is a set of matrices $\mathcal{C} \subseteq \mathbb{F}^{m \times n}$, and the rank-distance between two matrices X and Y is $\delta_R(X, Y) = \frac{1}{n} \text{rank}(X - Y)$ (where we assume without loss of generality that $n \leq m$). Using this notion of distance, one can again obtain a notion of list-decodability, and moreover average-radius list-decodability. There is a “rank-metric” list-decoding capacity $R^*(p)$. [LW18] showed that random linear rank-metric codes over the binary field \mathbb{F}_2 of rate $R^*(p) - \varepsilon$ are with high probability $(p, (1 - R^*(p))/\varepsilon + 2)$ -list-decodable, and one can adapt the argument above to show that such codes are with high probability $(p, (1 - R^*(p))/\varepsilon + 2)$ -average-radius list-decodable.*

References

- [AK97] Noga Alon and Michael Krivelevich. The concentration of the chromatic number of random graphs. *Combinatorica*, 17(3):303–313, 1997.
- [AN05] Dimitris Achlioptas and Assaf Naor. The two possible values of the chromatic number of a random graph. *Annals of Mathematics*, 162(3):1335–1351, 2005.
- [BE76] Béla Bollobás and Paul Erdős. Cliques in random graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 80, pages 419–427. Cambridge University Press, 1976.
- [Bli86] Volodia M. Blinovsky. Bounds for codes in the case of list decoding of finite volume. *Problems of Information Transmission*, 22(1):7–19, 1986.
- [Bli05] Vladimir M Blinovsky. Code bounds for multiple packings over a nonbinary finite alphabet. *Problems of Information Transmission*, 41(1):23–32, 2005.
- [BSKR09] Eli Ben-Sasson, Swastik Kopparty, and Jaikumar Radhakrishnan. Subspace polynomials and limits to list decoding of reed-solomon codes. *IEEE Transactions on Information Theory*, 56(1):113–120, 2009.
- [CGV13] Mahdi Cheraghchi, Venkatesan Guruswami, and Ameya Velingker. Restricted isometry of fourier matrices and list decodability of random linear codes. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 432–442, 2013.

- [DL12] Zeev Dvir and Shachar Lovett. Subspace evasive sets. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 351–358. ACM, 2012.
- [DMOZ19] Dean Doron, Dana Moshkovitz, Justin Oh, and David Zuckerman. Nearly optimal pseudorandomness from hardness. Technical report, ECCC preprint TR19-099, 2019.
- [Eli57] Peter Elias. List decoding for noisy channels. *Wescon Convention Record, Part 2*, pages 94–104, 1957.
- [Eli91] Peter Elias. Error-correcting codes for list decoding. *IEEE Transactions on Information Theory*, 37(1):5–12, 1991.
- [Gal62] Robert G. Gallager. Low-density parity-check codes. *IRE Trans. Information Theory*, 8(1):21–28, 1962.
- [GHK11] Venkatesan Guruswami, Johan Håstad, and Swastik Kopparty. On the list-decodability of random linear codes. *IEEE Trans. Information Theory*, 57(2):718–725, 2011.
- [GHSZ02] Venkatesan Guruswami, Johan Håstad, Madhu Sudan, and David Zuckerman. Combinatorial bounds for list decoding. *IEEE Trans. Information Theory*, 48(5):1021–1034, 2002.
- [GI01] Venkatesan Guruswami and Piotr Indyk. Expander-based constructions of efficiently decodable codes. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 658–667, 2001.
- [GI02] Venkatesan Guruswami and Piotr Indyk. Near-optimal linear-time codes for unique decoding and new list-decodable codes over smaller alphabets. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 812–821, 2002.
- [GI03] Venkatesan Guruswami and Piotr Indyk. Linear time encodable and list decodable codes. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 126–135, 2003.
- [GI04] Venkatesan Guruswami and Piotr Indyk. Efficiently decodable codes meeting gilbert-varshamov bound for low rates. In *SODA*, volume 4, pages 756–757. Citeseer, 2004.
- [GK16] Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. *Combinatorica*, 36(2):161–185, 2016.
- [GN14] Venkatesan Guruswami and Srivatsan Narayanan. Combinatorial limitations of average-radius list-decoding. *IEEE Trans. Information Theory*, 60(10):5827–5842, 2014.
- [GNP⁺13] Anna C Gilbert, Hung Q Ngo, Ely Porat, Atri Rudra, and Martin J Strauss. 12/12-foreach sparse recovery with low risk. In *International Colloquium on Automata, Languages, and Programming*, pages 461–472. Springer, 2013.

- [GR05] Venkatesan Guruswami and Atri Rudra. Limits to list decoding reed-solomon codes. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 602–609, 2005.
- [GR08] Venkatesan Guruswami and Atri Rudra. Concatenated codes can achieve list-decoding capacity. *Electronic Colloquium on Computational Complexity (ECCC)*, 15(054), 2008.
- [Gur03] Venkatesan Guruswami. List decoding from erasures: Bounds and code constructions. *IEEE Transactions on Information Theory*, 49(11):2826–2833, 2003.
- [GV05] Venkatesan Guruswami and Salil P. Vadhan. A lower bound on list size for list decoding. In *Approximation, Randomization and Combinatorial Optimization, Algorithms and Techniques, 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2005 and 9th International Workshop on Randomization and Computation, RANDOM 2005, Berkeley, CA, USA, August 22-24, 2005, Proceedings*, pages 318–329, 2005.
- [GX12] Venkatesan Guruswami and Chaoping Xing. Folded codes from function field towers and improved optimal rate list decoding. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 339–350. ACM, 2012.
- [GX13] Venkatesan Guruswami and Chaoping Xing. List decoding Reed-Solomon, Algebraic-Geometric, and Gabidulin subcodes up to the Singleton bound. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 843–852. ACM, 2013.
- [HIOS15] Iftach Haitner, Yuval Ishai, Eran Omri, and Ronen Shaltiel. Parallel hashing via list recoverability. In *Annual Cryptology Conference*, pages 173–190. Springer, 2015.
- [HRW17] Brett Hemenway, Noga Ron-Zewi, and Mary Wootters. Local list recovery of high-rate tensor codes & applications. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 204–215. IEEE, 2017.
- [HW18] Brett Hemenway and Mary Wootters. Linear-time list recovery of high-rate expander codes. *Information and Computation*, 261:202–218, 2018.
- [INR10] Piotr Indyk, Hung Q Ngo, and Atri Rudra. Efficiently decodable non-adaptive group testing. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1126–1142. SIAM, 2010.
- [KRRZ⁺19] Swastik Kopparty, Nicolas Resch, Noga Ron-Zewi, Shubhangi Saraf, and Shashwat Silas. On list recovery of high-rate tensor codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 2019.
- [KRSW18] Swastik Kopparty, Noga Ron-Zewi, Shubhangi Saraf, and Mary Wootters. Improved decoding of folded reed-solomon and multiplicity codes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 212–223. IEEE, 2018.
- [Luc91] Tomasz Luczak. A note on the sharp concentration of the chromatic number of random graphs. *Combinatorica*, 11(3):295–297, 1991.

- [LW18] Ray Li and Mary Wootters. Improved list-decodability of random linear binary codes. *arXiv preprint arXiv:1801.07839*, 2018.
- [Mat72] David W Matula. Employee party problem. In *Notices of the American Mathematical Society*, volume 19, pages A382–A382. AMER MATHEMATICAL SOC 201 CHARLES ST, PROVIDENCE, RI 02940-2213, 1972.
- [MRRZ⁺19] Jonathan Mosheiff, Nicolas Resch, Noga Ron-Zewi, Shashwat Silas, and Mary Wootters. Ldpc codes achieve list decoding capacity. *arXiv preprint arXiv:1909.06430*, 2019.
- [NPR11] Hung Q Ngo, Ely Porat, and Atri Rudra. Efficiently decodable error-correcting list disjunct matrices and applications. In *International Colloquium on Automata, Languages, and Programming*, pages 557–568. Springer, 2011.
- [RW10] Oliver Riordan and Nicholas Wormald. The diameter of sparse random graphs. *Combinatorics, Probability and Computing*, 19(5-6):835–926, 2010.
- [RW14a] Atri Rudra and Mary Wootters. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 764–773. ACM, 2014.
- [RW14b] Atri Rudra and Mary Wootters. It’ll probably work out: improved list-decoding through random operations. *Electronic Colloquium on Computational Complexity (ECCC)*, 21:104, 2014.
- [RW18] Atri Rudra and Mary Wootters. Average-radius list-recovery of random linear codes. In *Proceedings of the 2018 ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2018.
- [Woo13] Mary Wootters. On the list decodability of random linear codes with large error rates. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 853–860, 2013.
- [Woz58] Jack Wozencraft. List decoding. *Quarter Progress Report*, 48:90–95, 1958.
- [ZP81] Victor Vasilievich Zyablov and Mark Semenovich Pinsker. List concatenated decoding. *Problemy Peredachi Informatsii*, 17(4):29–33, 1981.

A Proof of Lemma 5.4

First, we restate Lemma 5.4 for the reader’s convenience.

Lemma 5.4. *Let $p \in (0, \frac{1}{2})$ and $R = 1 - h(p) - \varepsilon$ for $0 < \varepsilon < 1 - h(p)$. Let $\mathcal{C}_{Rn} \leq \mathbb{F}_2^n$ be a random linear code of rate R . Then $S_{\mathcal{C}_{Rn}} \leq 2$ with probability at least $1 - \exp(-\Omega_\eta(n))$.*

To prove Lemma 5.4, we introduce the notation $T_{\mathcal{C}} = S_{\mathcal{C}} - 1$ and show that if $T_{\mathcal{C}}$ bounded away from 1, it doubles with sufficiently large probability whenever we add a uniformly random vector to \mathcal{C} .

Lemma A.1. *If $\mathcal{C} \leq \mathbb{F}_2^n$ is a fixed linear code,*

$$\Pr_{b \sim \mathbb{F}_2^n} [S_{\mathcal{C}+\{0+b\}} > 1 + 2T_{\mathcal{C}} + T_{\mathcal{C}}^{1.5}] < T_{\mathcal{C}}^{0.5}.$$

Proof. Applying Lemma 5.3, for any fixed $b \in \mathbb{F}_2^n$,

$$\begin{aligned} S_{\mathcal{C}} &= \mathbb{E}_{x \sim \mathbb{F}_2^n} [A_{\mathcal{C}+\{0,b\}}(x)] \\ &\leq \mathbb{E}_{x \sim \mathbb{F}_2^n} [A_{\mathcal{C}}(x)A_{\mathcal{C}+b}(x)] \\ &= \mathbb{E}_{x \sim \mathbb{F}_2^n} [-1 + A_{\mathcal{C}}(x) + A_{\mathcal{C}}(x+b) + (A_{\mathcal{C}}(x) - 1)(A_{\mathcal{C}}(x+b) - 1)] \\ &= 1 + 2T_{\mathcal{C}} + \mathbb{E}_{x \sim \mathbb{F}_2^n} [(A_{\mathcal{C}}(x) - 1)(A_{\mathcal{C}}(x+b) - 1)]. \end{aligned}$$

Now, if x and b are independent and uniformly random over \mathbb{F}_2^n , then so are x and $x+b$, so we conclude

$$\mathbb{E}_b \mathbb{E}_x [(A_{\mathcal{C}}(x) - 1)(A_{\mathcal{C}}(x+b) - 1)] = \mathbb{E}_x [(A_{\mathcal{C}}(x) - 1)] \mathbb{E}_b [(A_{\mathcal{C}}(x+b) - 1)] = T_{\mathcal{C}}^2.$$

Hence, applying Markov's inequality (which is justified as $A_{\mathcal{C}}(x) - 1 > 0$ for all x),

$$\begin{aligned} \Pr_{b \sim \mathbb{F}_2^n} [S_{\mathcal{C}+\{0+b\}} > 1 + 2T_{\mathcal{C}} + T_{\mathcal{C}}^{1.5}] &\leq \Pr_{b \sim \mathbb{F}_2^n} \left[\mathbb{E}_x [(A_{\mathcal{C}}(x) - 1)] \mathbb{E}_b [(A_{\mathcal{C}}(x+b) - 1)] > T_{\mathcal{C}}^{1.5} \right] \\ &< \frac{T_{\mathcal{C}}^{1.5}}{T_{\mathcal{C}}^2} = T_{\mathcal{C}}^{0.5}. \quad \square \end{aligned}$$

We can now iteratively apply Lemma A.1 to conclude Lemma 5.4.

Proof of Lemma 5.4. Throughout the argument, we may assume n is sufficiently large compared to η . First, for $i = 0, 1, \dots, k$, consider

$$\begin{aligned} \delta_0 &:= \lambda n 2^{-n(1-h(\lambda)-\frac{\eta}{1+\eta})}, \\ \delta_i &:= 2\delta_{i-1} + \delta_{i-1}^{1.5}. \end{aligned}$$

By induction, we claim that for all $i \leq n(1-h(p)-\varepsilon)$, we have $\delta_i < 2^{i+1}\delta_0 < 2^{-\frac{\eta^2 n}{3}}$. First, we note that for $k = (1-h(p)-\varepsilon)n$,

$$\begin{aligned} 2^{k+1}\delta_0 &= (2\lambda n) \cdot 2^{n(1-h(p)-\varepsilon)} \cdot 2^{-n(1-h(\lambda)-\frac{\eta}{1+\eta})} \\ &\leq (2\lambda n) \cdot 2^{n[(1-h(p)-\varepsilon)-(1-h(\lambda)-\frac{\eta}{1+\eta})-\frac{\eta^2}{2}]} \end{aligned} \tag{16}$$

$$= (2\lambda n) \cdot 2^{-\frac{\eta^2 n}{2}} \tag{17}$$

$$< 2^{-\frac{\eta^2 n}{3}}. \tag{18}$$

In the above, Inequality (16) follows from the inequality $\frac{\eta}{1+\eta} \geq \eta + \eta/2$, valid for $\eta \in (0, 1)$. We used the equality $h(p) + \varepsilon = h(\lambda) + \eta$ to obtain (17). The last line, (18), holds for sufficiently large

n . Hence, $2^{i+1}\delta_0 < 2^{-\frac{\eta^2 n}{3}}$ for all $i \leq k$, so we may assume this in inductively proving $\delta_i < 2^{i+1}\delta_0$ for all $i = 0, 1, \dots, k$.

Now, we clearly have $\delta_0 < 2\delta_0$ (so the base case of the induction holds), while for $i \geq 1$ we bound

$$\delta_i = 2\delta_{i-1} \left(1 + \frac{\sqrt{\delta_{i-1}}}{2}\right) = 2^i \delta_0 \prod_{j=0}^{i-1} \left(1 + \frac{\sqrt{\delta_j}}{2}\right) \leq 2^i \delta_0 \cdot \exp\left(\frac{1}{2} \cdot \sum_{j=0}^{i-1} \sqrt{\delta_j}\right) < 2^{i+1} \delta_0.$$

In the first two equalities, we applied the definitions of the δ_i 's. The first inequality applies the estimate $1 + z \leq e^z$, while the second uses the induction hypothesis $\delta_j < 2^{-\frac{\eta^2 n}{3}}$ for $j < i$ and by ensuring n is sufficiently large.

Now, let $b_1, \dots, b_k \sim \mathbb{F}_2^n$ be i.i.d. uniform random vectors, and let $\mathcal{C}_i = \text{span}\{b_1, \dots, b_i\}$ denote the ‘‘intermediate’’ random linear codes. Call \mathcal{C}_i *good* if $T_{\mathcal{C}_i} \leq \delta_i$; we wish to show that with high probability, \mathcal{C}_i is good for all $i \leq k$. For $i = 0$, we apply (7) and obtain

$$T_{\mathcal{C}_0} = S_{\{0\}} - 1 \leq 1 + \lambda n 2^{-n(1-h(\lambda)-\frac{\eta}{1+\eta})} - 1 = \delta_0.$$

Now, let $i \geq 1$ and assume \mathcal{C}_i is good. By Lemma A.1,

$$\begin{aligned} \Pr[\mathcal{C}_{i+1} \text{ is not good}] &= \Pr[T_{\mathcal{C}_{i+1}} > \delta_{i+1}] \leq \Pr[T_{\mathcal{C}_{i+1}} > 2T_{\mathcal{C}_i} + T_{\mathcal{C}_i}^{1.5}] \\ &< T_{\mathcal{C}_i}^{0.5} \leq \delta_i^{0.5}. \end{aligned}$$

Thus, with probability at least

$$1 - (\delta_0^{0.5} + \delta_1^{0.5} + \dots + \delta_k^{0.5}) > 1 - k 2^{-\frac{\eta^2 n}{6}} \geq 1 - 2^{-\Omega_\eta(n)}$$

we have $T_{\mathcal{C}_i} \leq \delta_i$ for all $i = 0, 1, \dots, k$, as desired. In particular, we conclude $S_{\mathcal{C}_k} = 1 + T_{\mathcal{C}_k} \leq 1 + 2^{-\frac{\eta^2 n}{3}} \leq 2$ with probability $1 - 2^{-\Omega_\eta(n)}$. \square