

A New Coding-based Algorithm for Finding Closest Pair of Vectors*

Ning Xie[†]

Shuai Xu[‡]

Yekun Xu[§]

Abstract

Given n vectors x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$, how to find two vectors whose pairwise Hamming distance is minimum? This problem is known as the *Closest Pair Problem*. If these vectors are generated uniformly at random except two of them are correlated with Pearson-correlation coefficient ρ , then the problem is called the *Light Bulb Problem*. In this work, we propose a novel coding-based scheme for the Closest Pair Problem. We design both randomized and deterministic algorithms, which achieve the best-known running time when the length of input vectors m is small and the minimum distance is very small compared to m . Specifically, the running time of our randomized algorithm is $O(n \log^2 n \cdot 2^{cm} \cdot \text{poly}(m))$ and the running time of our deterministic algorithm is $O(n \log n \cdot 2^{c'm} \cdot \text{poly}(m))$, where c and c' are constants depending only on the (relative) distance of the closest pair. When applied to the Light Bulb Problem, our result yields state-of-the-art deterministic running time when the Pearson-correlation coefficient ρ is very large. Specifically, when $\rho \geq 0.9933$, our deterministic algorithm runs faster than the previously best deterministic algorithm (Alman, SOSA 2019).

keywords: Closest Pair Problem, Light Bulb Problem, Error Correcting Codes

*An extended abstract of this article appeared in Proceedings of the 13th International Computer Science Symposium in Russia (CSR'18), pages 321–333.

[†]Florida International University, Miami, FL 33199, USA. Email: nxie@cis.fiu.edu. Research supported in part by NSF grant 1423034.

[‡]Florida International University, Miami, FL 33199, USA. Email: sxu010@fiu.edu. Research supported in part by NSF grant 1423034.

[§]Florida International University, Miami, FL 33199, USA. Email: yxu040@fiu.edu. Research supported in part by NSF grant 1423034.

1 Introduction

We consider the following classic *Closest Pair Problem*: given n vectors x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$, how to find the two vectors with the minimum pairwise distance? Here the distance is the usual Hamming distance: $\text{dist}(x_i, x_j) = |\{k \in [m] : (x_i)_k \neq (x_j)_k\}|$, where $(x_i)_k$ denotes the k^{th} component of vector x_i . Without loss of generality, we assume that $d_{\min} = \text{dist}(x_0, x_1)$ is the unique minimum distance and all other pairwise distances are greater than d_{\min} .

The Closest Pair Problem is one of the most fundamental and well-studied problems in many science disciplines, having a wide spectrum of applications in computational finance, DNA detection, weather prediction, etc. For instance, the Closest Pair Problem has the following interesting application in bioinformatics. Scientists wish to find connections between Single Nucleotide Polymorphisms (SNPs) and phenotypic traits. SNPs are one of the most common types of genetic differences among people, with each SNP representing a variation in a single DNA block called *nucleotide* [22]. Screening for most correlated pairs of SNPs has been applied to study such connections [11, 15, 17, 38]. As the number of SNPs in humans is estimated to be around 10 to 11 million, for problem size n of this size, any improvement in running time for solving the Closest Pair Problem would have huge impacts on genetics and computational biology [38].

In theoretical computer science, the Closest Pair Problem has a long history in computational geometry, see e.g. [43] for a survey of many classic algorithms for the problem. The naive algorithm for the Closest Pair Problem takes $O(mn^2)$ time. When the dimension m is a constant, either in the Euclidean space or ℓ_p space, the classic divide-and-conquer based algorithm runs in $O(n \log n)$ time [14]. Rabin [42] combined the floor function with randomization to devise a linear time algorithm. In 1995, Khuller and Matias [31] simplified Rabin’s algorithm to achieve the same running time $O(n)$ and space complexity $O(n)$. Golin *et al.* [24] used dynamic perfect hashing to implement a dictionary and obtained the same linear time and space bounds.

When the dimension m is not a constant, due to a well-know phenomenon called *curse of dimensionality*, this problem becomes much harder. The first subquadratic time algorithm for the Closest Pair Problem is due to Alman and Williams [5] for m as large as $\log^{2-o(1)} n$. The algorithm is built on a newly developed framework called *polynomial method* [49, 50, 2]. In particular, Alman and Williams first constructed a probabilistic polynomial of degree $O(\sqrt{n \log 1/\epsilon})$ which computes the MAJORITY function on n variables with error at most ϵ , then applied the polynomial method to design an algorithm which runs in $n^{2-1/O(s(n) \log^2 s(n))}$ time where $m = s(n) \log n$, and computed the minimum Hamming distance among all red-blue vector pairs¹ through polynomial evaluations. In a more recent work, Alman *et al.* [4] unified Valiant’s fast matrix multiplication approach [45] with that by Alman and Williams [5]. They constructed probabilistic *polynomial threshold functions* (PTFs) to obtain a simpler algorithm which improved to randomized time $n^{2-1/O(\sqrt{s(n)} \log^{3/2} s(n))}$ or deterministic time $n^{2-1/O(s(n) \log^2 s(n))}$.

The Light Bulb Problem. A special case of the Closest Pair Problem, the so-called *Light Bulb Problem*, was first posed by L. Valiant in 1988 [46]. In this problem, we are given a set of n vectors in $\{0, 1\}^m$ chosen uniformly at random from the Boolean hypercube, except that two of them are non-trivially correlated (specifically, have Pearson-correlation coefficient ρ , which is equivalent to that the expected Hamming distance between the correlated pair is $\frac{1-\rho}{2}m$), the problem then is to find the correlated pair.

Paturi *et al.* [41] gave the first non-trivial algorithm, which runs² in $\tilde{O}(n^{2-\log(1+\rho)})$. In 2010, Dubiner [19] proposed a Bucketing Coding algorithm which runs in time $\tilde{O}(n^{\frac{2}{1+\rho}})$. The well-known *locality*

¹The actual problem solved in [5] is the so-called *Bichromatic Hamming Closest Pair Problem*; see discussion in Section 1.3 below.

²We adopt the common notation $\tilde{O}(n^k)$ to denote $n^k \cdot \text{polylog}(n)$.

sensitive hashing scheme of Indyk and Motwani [27] performs slightly worse than Paturi *et al.*'s hash-based algorithm but recent *data-dependent LSH* [9] matches the running time of Dubiner's. Roughly speaking, a family of hash functions \mathcal{H} is called (r, cr, p_1, p_2) -sensitive if, for any two points p and q in a metric space (X, d) , a randomly chosen hash function from \mathcal{H} hashes p and q into the same bucket with probability at least p_1 if they are close (i.e., when $d(p, q) \leq r$) and with probability at most p_2 if they are far apart (i.e., when $d(p, q) \geq cr$), where $c > 1$ is the *approximation* factor and $p_1 > p_2$. Indyk and Motwani [27] proved that such a family of LSH can be used to construct a data structure solving the c -approximate Nearest Neighbor Search problem. Specifically, for a data set consisting of at most n points from X , the data structure uses $\tilde{O}(n^{1+\varrho})$ space (and $\tilde{O}(n^{1+\varrho})$ preprocessing time) and supports $\tilde{O}(m \cdot n^\varrho)$ query time, where m is the dimension of the space and $\varrho := \frac{\log 1/p_1}{\log 1/p_2}$ basically quantifies the quality of the LSH. When (X, d) is the Hamming space, the original work of Indyk and Motwani [27] achieved $\varrho \leq 1/c$, while the current best result is $\varrho = 1/2c - 1$ by Andoni *et al.* [9], under the framework of *data-dependent LSH*. Applying LSH to the Light Bulb Problem, we have $m = O(\log n)$, $c \geq \frac{1}{1-\rho}$ with high probability, and we need to pay the one-time preprocessing time and n queries for each vector to search for its nearest neighbor in the data set. Therefore LSH solves the Light Bulb Problem in time $\tilde{O}(n^{2-\rho})$ using the original data-independent scheme of Indyk and Motwani, and can be improved to $\tilde{O}(n^{\frac{2}{1+\rho}})$ using the data-dependent scheme in [9]. As ρ gets small, all these three algorithms have running time $\tilde{O}(n^{2-c\rho})$ for various constants c .³ Comparing the constants in these three algorithms, Dubiner and data-dependent LSH achieve the best constant, which is $\tilde{O}(n^{2-2\rho})$, in the limit of $\rho \rightarrow 0$. Asymptotically the same bound was also achieved by May and Ozerov [35], in which the authors used algorithms that find Hamming closest pairs to improve the running time of decoding random binary linear codes.

The breakthrough result of Valiant [45] is a fast matrix multiplication based algorithm which finds the "planted" closest pair in time $O(n^{\frac{5-\omega}{4-\omega} + \epsilon}) < n^{1.62} \cdot \text{poly}(1/\rho)$ with high probability for any constant $\epsilon, \rho > 0$ and $m > n^{\frac{1}{4-\omega}}/\rho^2$, where $\omega < 2.373$ is the exponent of fast matrix multiplications. The most striking feature of Valiant's algorithm is that ρ does not appear in the exponent of n in the running time of the algorithm. Karppa *et al.* [29] further improved Valiant's algorithm to $n^{1.582}$. Very recently, Alman [3] combined techniques in [45] with the polynomial method to give a very elegant and simple algorithm which matches Karppa *et al.*' bound. Moreover, Alman derandomized his algorithm and improved on the previously best deterministic running time by Karppa *et al.* [30]. Note that Valiant, Karppa *et al.* and Alman achieved runtimes of $n^{2-\Omega(1)}(m/\epsilon)^{O(1)}$ for the Light Bulb Problem, which improved upon previous algorithms that rely on the Locality Sensitive Hashing (LSH) schemes. The LSH based algorithms only achieved runtime of $n^{2-O(\epsilon)}$ for the Light Bulb Problem.

We remark that all the above-mentioned algorithms (except May and Ozerov's work) that achieve state-of-the-art running time are based on either involved probabilistic polynomial constructions or impractical $O(n^\omega)$ fast matrix multiplications⁴, or both.

Overview of our main results. In this work, we propose a new coding-based scheme for the Closest Pair Problem. We design both randomized and deterministic algorithms, which achieve the best-known running time when the length of input vectors m is small ($m = O(\log n)$) and the minimum distance is very small compared to m . Specifically, the running time of our randomized algorithm is $O(n \log^2 n \cdot 2^{cm} \cdot \text{poly}(m))$ and the running time of our deterministic algorithm is $O(n \log n \cdot 2^{c'm} \cdot \text{poly}(m))$, where c and c' are constants

³When ρ goes to zero, the exponent in the running time of Paturi *et al.* [41] is $2 - \log(e) \cdot \rho + O(\rho^2)$.

⁴Subcubic fast matrix multiplication algorithms are practical for Strassen-based ones [13, 26] and are practical for very large input sizes up to $\omega = 2.7734$ (see e.g. the survey [40]). However, all other theoretically more efficient algorithms, such as recent developments [44, 52, 33], are superior to the trivial cubic algorithm only for matrices of colossal sizes.

depending only on the (relative) distance of the closest pair; see Section 1.2 for precise statements. Since the running time of our algorithms are exponential in m , they are subquadratic-time algorithms only when $m \leq \alpha \log n$ for some constant $\alpha > 0$. When applied to the Light Bulb Problem, our deterministic algorithm achieves state-of-the-art running time when the Pearson-correlation coefficient ρ is very large.

1.1 Our approach

Algorithm 1: General Idea of Main Algorithm
<p>input : A set of n vectors x_0, \dots, x_{n-1} in $\{0, 1\}^m$ and d_{\min}</p> <p>output: Two vectors and their distance</p> <ol style="list-style-type: none"> 1 generate a binary code $C \subseteq \{0, 1\}^m$ 2 pick a random $y \in \{0, 1\}^m$ 3 for $j \leftarrow 0$ to $n - 1$ do <li style="padding-left: 20px;">4 decode $y + x_j$ in C, and denote the resulting vector by \tilde{x}_j 5 end 6 sort $\tilde{x}_0, \dots, \tilde{x}_{n-1}$ 7 for each of the $n - 1$ pairs of adjacent vectors in the sorted list do <li style="padding-left: 20px;">8 compute the distance between the two <i>original</i> vectors. 9 end 10 output the pair of vectors with the minimum distance and their distance

We propose a simple, error-correcting code based scheme for the Closest Pair Problem. Apart from achieving the best running time for certain range of parameters, we believe that our new approach has the merit of being simple, and hence more likely being practical as well. In particular, neither complicated data structure nor fast matrix multiplication is employed in our algorithms.

The basic idea of our algorithms is very simple. Suppose for concreteness that x_0 and x_1 are the unique pair of vectors that achieve the minimum distance. Our scheme is inspired by the extreme case when x_0 and x_1 are identical vectors. In this case, a simple *sort and check* approach solves the problem in $O(mn \log n)$ time: sort all n vectors and then compute only the $n - 1$ pairwise distances (instead of all $\binom{n}{2}$ distances) of adjacent vectors in the sorted list. Since the two closest vectors are identical, they must be adjacent in the sorted list and thus the algorithm would compute their distance and find them. This motivates us to view the input vectors as received messages that were encoded by an error correction code and have been transmitted through a noisy channel. As a result, the originally identical vectors are no longer the same, nevertheless are still very close. Directly applying the *sort and check* approach would fail but a natural remedy is to decode these received messages into codewords first. Indeed, if the distance between x_0 and x_1 is small and we are lucky to have a codeword c that is very close to both of them, then a unique decoding algorithm would decode both of these two vectors into c . Now if we “sort” the *decoded* vectors and then “check” the corresponding *original* vectors of each adjacent pair of vectors⁵, the algorithm would successfully find the closest pair. How to turn this “good luck” into a working algorithm? Simply try different shift vectors y and view $y + x_i$ as the input vectors, since the Hamming distances are invariant under any shift. The basic idea of our approach is summarized in Algorithm 1.

Figure 1 illustrates the effects “bad” shift vectors and “good” shift vectors on the decoding part of our

⁵Actually, we only need to “check” when the two adjacent decoded vectors are identical.

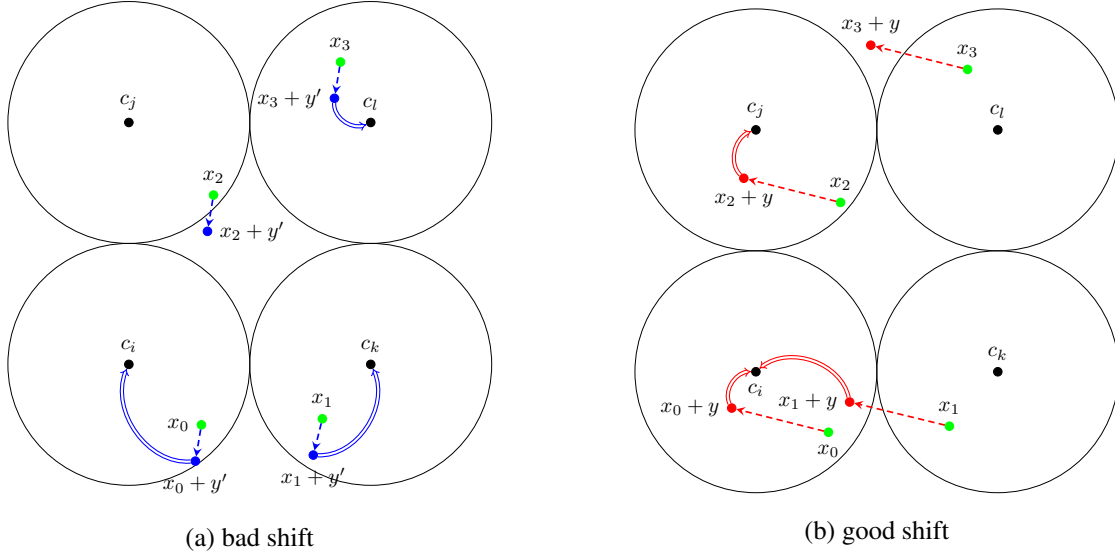


Figure 1: Decoding with good and bad shift vectors

algorithm; here single arrows with dotted lines point from original vectors to shifted target vectors, while double arrows indicate into which codewords are shifted target vectors decoded. In Figure 1a, our shifted target vectors $x_0 + y'$ and $x_1 + y'$ are decoded into two different codewords, so y' is a bad shift. In Figure 1b, our shifted target vectors $x_0 + y$ and $x_1 + y$ are decoded into the same codeword, therefore we can apply the sort-and-check approach to find the closest pair.

Figure 2 illustrates what happens if we sort the vectors directly and why sorting decoded vectors works.

Making the idea of decoding work for larger minimum pairwise distance involves balancing the parameters of the error-correcting code so that it is efficiently decodable as well as having appropriate decoding radius. The decoding radius r should have the following properties. On one hand, r should be small to ensure that there is a codeword c such that only x_0 and x_1 will be decoded into c (therefore x_0 and x_1 will be adjacent in the sorted array and hence will be compared with each other). On the other hand, we would like r to be large so as to maximize the number of “good” shift vectors which enable both x_0 and x_1 decoding to the same codeword. As a result, our algorithms generally perform best when the closest pair distance is very small.

1.2 Our results

Our simple error-correcting code based algorithm can be applied to solve the Closest Pair Problem and the Light Bulb Problem.

1.2.1 The Closest Pair Problem

Our main result is the following simple randomized algorithm for the Closest Pair Problem.

Theorem 1.1 (Main). *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} (and the second smallest distance can be as small as*

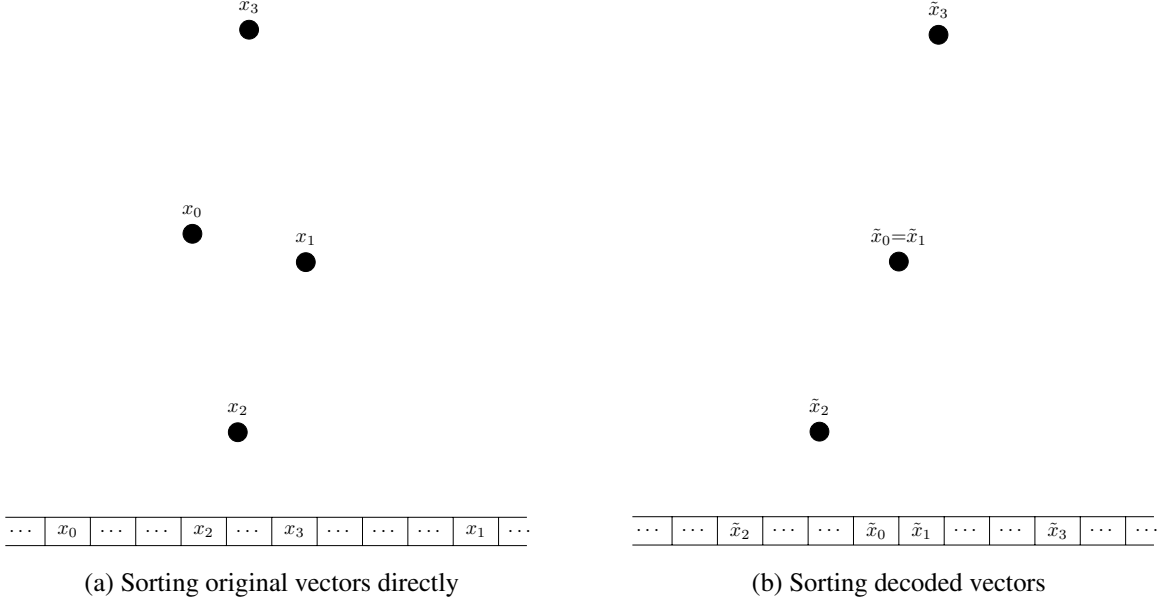


Figure 2: Difference between sorting input vectors directly and sorting decoded vectors.

$d_{\min} + 1$). Suppose⁶ we are given the value of d_{\min} and let $\delta := d_{\min}/m$. Then there is a randomized algorithm running in $O(n \log^2 n \cdot 2^{(1-\kappa_Z(\delta)-\delta)m} \cdot \text{poly}(m))$ which finds the closest pair x_0 and x_1 with probability at least $1 - 1/n^2$. The running time can be improved to $O(n \log^2 n \cdot 2^{(H_2(\delta)-\delta)m} \cdot \text{poly}(m))$, if we are given black-box decoding algorithms for an ensemble of $O(\log m/\epsilon)$ binary error-correcting codes that meet the Gilbert-Varshamov bound.

Here $\kappa_{GV}(\delta)$ and $\kappa_Z(\delta)$ are functions derived from the Gilbert-Varshamov (GV) bound and the Zyablov bound respectively (see Section 2.1.5 for details). Specifically, $\kappa_{GV}(\delta) = 1 - H_2(\delta)$, and both $\kappa_{GV}(\delta)$ and $\kappa_Z(\delta)$ are monotone decreasing functions for $\delta \in [0, 1/2]$, with function values ranging from 1 to 0; see e.g. Figure 9.2 in [25] for an illustration.

The running time of our algorithm depends on — in addition to the number of vectors n — both dimension m and $\delta := d_{\min}/m$. To illustrate its performance we choose two typical vector lengths m , namely those corresponding to the Hamming bound⁷ and the Gilbert-Varshamov (GV) bound⁸, and list the exponents γ' in the running time of the GV-code version of our algorithm as a function of d_{\min} (in fact δ) in Table 1. Here, we write the running of the algorithm as $\tilde{O}(n^{\gamma'})$, where \tilde{O} suppresses any polylogarithmic factor of n . One can see that our algorithm runs in subquadratic time when δ is small, or equivalently when the Hamming distance between the closest pair is small. For instance, when $\delta = 0.05$, and the length $m = 1.4013 \log n$, then the running time is $O(n^{1.3313})$ if we use GV bound.

In the setting of $m = c \log n$ for some not too large constant c , the current best result is the randomized algorithm of Alman *et al.* [4], which runs in $n^{2-1/O(\sqrt{c} \log^{3/2} c)}$ time for the Closest Pair Problem. As it is very hard to calculate the hidden constant in the exponent of their running time, it is impossible to compare our running time with theirs quantitatively.

⁶In fact this assumption can be easily removed with a small overhead in the running time; see the discussion below and Section 4.

⁷The Hamming bound, also known as the sphere packing bound, specifies an upper bound on the number of codewords a code can have given the block length and the minimum distance of the code.

⁸The GV bound is known to be attainable by random codes.

Table 1: Running time of our algorithm when vector length m and relative distance δ meets the Hamming bound and GV bound

δ	Hamming bound		GV bound	
	length of vector ($m/\log n$)	exponent (γ')	length of vector ($m/\log n$)	exponent (γ')
0.01	1.0476	1.0742	1.0879	1.0770
0.025	1.1074	1.1591	1.2029	1.1728
0.05	1.2029	1.2844	1.4013	1.3313
0.075	1.2999	1.4021	1.6242	1.5024
0.1	1.4013	1.5171	1.8832	1.6949
0.125	1.5090	1.6316	2.1909	1.9170
0.133	1.5449	1.6684	2.3064	1.9989

Deterministic algorithm. By checking all shift vectors up to certain Hamming weight, our randomized algorithm can be easily derandomized to yield the following theorem.

Theorem 1.2. *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} (and the second smallest distance can be as small as $d_{\min} + 1$). Suppose we are given the value of d_{\min} and let $\delta := d_{\min}/m$. Then there is a deterministic algorithm that finds the closest pair x_0 and x_1 with running time $O(n \log n \cdot 2^{H_2(1-\kappa_Z(\delta))m} \cdot \text{poly}(m))$, where $H_2(\cdot)$ is the binary entropy function. Moreover, if we are given as black box the decoding algorithm of a random Varshamov linear code with block length m and minimum distance $d_{\min} + 1$, then the running time is $O(n \log n \cdot 2^{H_2(H_2(\delta))m} \cdot \text{poly}(m))$.*

Searching for d_{\min} . If we remove the assumption that d_{\min} is given, our algorithm can be modified to search for d_{\min} first without too much slowdown; more details appear in Section 4.

Theorem 1.3. *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} . Then for any $\epsilon > 0$, there is a randomized algorithm that runs in $O(\epsilon^{-1} n \log^2 n \cdot 2^{(1-\kappa_Z((1+\epsilon)\delta)-\delta H_2(\frac{1-\epsilon}{2}))m} \cdot \text{poly}(m))$ which finds the d_{\min} (and the pair x_0 and x_1) with probability at least $1 - 1/n$. The running time can be improved to $O(\epsilon^{-1} n \log^2 n \cdot 2^{(H_2((1+\epsilon)\delta)-\delta H_2(\frac{1-\epsilon}{2}))m} \cdot \text{poly}(m))$, if we are given black-box decoding algorithms for an ensemble of $O(\log m/\epsilon)$ binary error-correcting codes that meet the Gilbert-Varshamov bound.*

Gapped version. Intuitively, if there is a gap between d_{\min} and the second minimum distance, the Closest Pair Problem should be easier. This is reminiscent of the case of the $(1 + \epsilon)$ -Approximate NNS Problem versus the NNS Problem. However, as we still need to find the *exact* solution to the Closest Pair Problem, the situation here is different.

Theorem 1.4 (Gapped version). *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} . Suppose we are given the values of d_{\min} as well as the second minimum distance d_2 . Let $\delta := d_{\min}/m$ and $\delta' := d_2/m$. Then there is a randomized algorithm running in $O(n \log^2 n \cdot 2^{(1-\kappa_Z(\delta'))-\delta-(1-\delta)H_2(\frac{\delta'-\delta}{2(1-\delta)})m} \cdot \text{poly}(m))$ which finds the closest pair x_0 and x_1 with probability at least $1 - 1/n^2$. Moreover, the running time can be further improved to*




Figure 3: The range of ϵ in which gapped version outperforms non-gapped version

$O(n \log^2 n \cdot 2^{(H_2(\delta') - \delta - (1-\delta)H_2(\frac{\delta' - \delta}{2(1-\delta)}))m} \cdot \text{poly}(m))$, if we are given the black box access to the decoding algorithm of an (m, K, d) -code which meets the Gilbert-Varshamov bound.

Our gapped version algorithm uses $d_2/2$ instead of $d_{\min}/2$ as the decoding radius. This, however, does not always give improved running time as illustrated in Figure 3. In Figure 3, we set $\delta' = (1 + \epsilon)\delta$ and write the running time as $O(n \log^2 n \cdot 2^{\gamma m} \cdot \text{poly}(m))$ for both the gapped version (the blue line) and the non-gapped version (the green line). One can see that using $d_2/2$ as the decoding radius does not always yield the best running time. Indeed, this is the case only when ϵ is small enough. Our numerical calculations show that there exists an optimal decoding radius $d_{\text{opt}}/2$ (which corresponds to the minimum point in the blue line) slightly larger than $d_{\min}/2$ such that whenever $d_2 \geq d_{\text{opt}}$ using $d_{\text{opt}}/2$ as the decoding radius will achieve the fastest running time. Unfortunately we do not know how to calculate this $d_{\text{opt}}/2$ analytically.

1.2.2 The Light Bulb Problem

Applying our algorithms for the Closest Pair Problem to the Light Bulb Problem easily yields the following theorem.

Theorem 1.5. *There is a randomized algorithm for the Light Bulb Problem which runs in time*

$$O(n \cdot \text{poly}(\log n)) \cdot 2^{(1 - \kappa_Z(\frac{1-\rho}{2}) - \frac{1-\rho}{2}) \frac{4 \ln 2 \cdot \log n}{\rho^2} (1+o(1))}$$

and succeeds with probability at least $1 - 1/n^2$. The running time can be further improved to

$$O(n \cdot \text{poly}(\log n)) \cdot 2^{(H_2(\frac{1-\rho}{2}) - \frac{1-\rho}{2}) \frac{4 \ln 2 \cdot \log n}{\rho^2} (1+o(1))},$$

if we are allowed a one-time preprocessing time⁹ of $n^{2.773/\rho^2}$ to generate the decoding lookup table of a random Gilbert's $(m, K, (1 - \rho)m/2)$ -code. Similar results can also be obtained for deterministic algorithms.

Our deterministic algorithm for the Light Bulb Problem performs faster than Alman's deterministic algorithm [3] when the Pearson-correlation coefficient ρ is very large. Moreover, we believe that our algorithms are very simple and therefore are likely to outperform other complicated ones for at least not too large input sizes.

⁹This is because the block length of the code is $m = 4 \ln 2 \log n / \rho^2 < 2.773 \log n / \rho^2$ and preprocessing the code requires $O(2^m) = O(n^{2.773/\rho^2})$ time.

1.3 Related work

The Nearest Neighbor Search problem. The Closest Pair Problem is a special case of the more general *Nearest Neighbor Search* (NNS) problem, defined as follows. Given a set S of n vectors in $\{0, 1\}^m$, and a query point $q \in \{0, 1\}^m$ as input, the problem is to find a point in S which is closest to q . The performance of an NNS algorithm is usually measured by two parameters: the space (which is usually proportional to the preprocessing time) and the query time. It is easy to see that any algorithms for NNS can also be used to solve the Closest Pair problem, as we can try each vector in S as the query vector against the remaining vectors in S , and output the pair with minimum distance.

Most early work on this problem is for fixed dimension. Indeed, when $m = 1$ the problem is easy, as we can just sort the input vectors (which in this case are numbers), then perform a binary search to find the closest vector to the input query. For $m \geq 2$, Clarkson [16] gave an algorithm with query time polynomial in $m \log n$, and space complexity $O(n^{\lceil m/2 \rceil})$. Meiser [36] designed an algorithm which runs in $O(m^5 \log n)$ time and uses $O(n^{m+\epsilon})$ space for arbitrary $\epsilon > 0$. By far, all efficient data structures for NNS have dimension m appear in the exponent of the space complexity, due to the *curse of dimensionality*.

This motivated people to introduce a relaxed version of Nearest Neighbor Search called the $(1 + \epsilon)$ -*Approximate Nearest Neighbor Search* ($(1 + \epsilon)$ -*Approximate NNS*) Problem in the 1990s. The problem now is, for an input query point q , find a point p in S such that the Hamming distance is:

$$\text{dist}(p, q) \leq (1 + \epsilon) \min_{p' \in S} \text{dist}(p', q).$$

We call such a p as a $(1 + \epsilon)$ -approximate nearest neighbor of input query q .

The $(1 + \epsilon)$ -Approximate NNS Problem has been studied extensively in the last two decades. In 1998, Indyk and Motwani [27] used a set of hash functions to store the dataset such that if two points are close enough, they will have a very high probability to be hashed into the same buckets. As a pair of close points have higher probability than a pair of far-apart points to fall into the same bucket, the scheme is called *locality sensitive hashing* (LSH). The query time of LSH is $O(n^{\frac{1}{1+\epsilon}})$, which is sublinear, and the space complexity of LSH is $O(n^{1+\frac{1}{1+\epsilon}})$, which is subquadratic. After Indyk and Motwani introducing the locality sensitive hashing, there have been many improvements on the parameters under different metric spaces, such as ℓ_p metric [32, 18, 7, 39, 37]. Recently, Andoni *et al.* [9] gave tight upper and lower bounds on the time-space trade-offs of (data-dependent) hashing based algorithms for the $(1 + \epsilon)$ -Approximate NNS Problem. This is the first algorithm that achieves sublinear query time and near-linear space, for any $\epsilon > 0$. For many results on the Approximate NNS problem in high dimension, see e.g. [8] for a survey. Some algorithms for the low dimension problem are surveyed in [10].

In 2012 Valiant [45] leveraged fast matrix multiplication to obtain a new algorithm for the $(1 + \epsilon)$ -Approximate NNS Problem that is not based on LSH.¹⁰ The general setting of Valiant’s results is the following. Suppose there is a set of points S in m -dimensional Euclidean (or Hamming) space, and we are promised that for any $a \in S$ and $b \in S$, $\langle a, b \rangle < \alpha$, except for only one pair which has $\langle a, b \rangle \geq \beta$ (which corresponds to the closest pair, and β is known as the Pearson-correlation coefficient), for some $0 < \alpha < \beta < 1$. Valiant’s algorithm finds the closest pair in $n^{\frac{5-\omega}{4-\omega} + \omega \frac{\log \beta}{\log \alpha}} m^{O(1)}$ time, where ω is the exponent for fast matrix multiplication ($\omega < 2.373$). Notice that, if the Pearson-correlation coefficient β is some fixed constant, then when α approaches 0 the running time tends to $n^{\frac{5-\omega}{4-\omega}}$, which is less than $n^{1.62}$. Valiant applied his algorithms to get improved bounds¹¹ for the Learning Sparse Parities with Noise

¹⁰In fact, Valiant’s algorithm can handle polynomially many “outlier” pairs.

¹¹All these results are due to the fact that Valiant’s algorithms are much more robust to weak correlations than other algorithms.

Problem, the Learning k -Juntas with Noise Problem, the Learning k -Juntas without Noise Problem, and so on. More recently, Karppa *et al.* [29] improved upon Valiant’s algorithm and obtained an algorithm that runs in $n^{\frac{2\omega}{3} + O(\frac{\log \beta}{\log \alpha})} m^{O(1)}$ time.

Note that, in general, algorithms for the $(1 + \epsilon)$ -Approximate NNS can only be applied to the gapped version of the Closest Pair Problem; for non-gapped version, as the minimum distant and the second minimum distant can differ by 1, which means that the approximation parameter ϵ tends to zero if the minimum distance is large, the running time will approach to quadratic. However, our non-gapped version algorithm still runs in truly subquadratic time in this case.

Decoding Random Binary Linear Codes. In 2015, May and Ozerov [35] observed that algorithms for high dimensional Nearest Neighbor Search Problem can be used to speedup the approximate matching part of the *information set decoding* algorithm. They designed a new algorithm for the *Bichromatic Hamming Closest Pair* problem when the two input lists of vectors are pairwise independent, and consequently obtained a decoding algorithm for random binary linear codes with time complexity $2^{0.097n}$. This improved upon the previously best result of Becker *et al.* [12] which runs in $2^{0.102n}$.

The Bichromatic Hamming Closest Pair problem. In fact, the problem studied in [5, 4, 35] is the following *Bichromatic Hamming Closest Pair Problem*: we are given n red vectors $R = \{r_0, r_1, \dots, r_{n-1}\}$ and n blue vectors $B = \{b_0, b_1, \dots, b_{n-1}\}$ from $\{0, 1\}^m$, and the goal is to find a red-blue pair with minimum Hamming distance. It is easy to see that the Closest Pair Problem is reducible to the Bichromatic Hamming Closest Pair Problem via a random reduction. In fact, our algorithm for the Closest Pair Problem can also be easily adapted to solve the Bichromatic Hamming Closest Pair Problem as follows. Run the decoding part of our algorithm on both sets R and B to get $\tilde{R} = \{\tilde{r}_0, \tilde{r}_1, \dots, \tilde{r}_{n-1}\}$ and $\tilde{B} = \{\tilde{b}_0, \tilde{b}_1, \dots, \tilde{b}_{n-1}\}$, sort \tilde{R} and \tilde{B} separately (without comparing the original vectors for adjacent pairs in the sorted lists), then merge the two sorted lists into one, and compute the distance between the original vectors for each red-blue pair of vectors that are compared during the merging process. On the other hand, the Bichromatic Closest Pair Problem is unlikely to have *truly* subquadratic algorithms under some mild conditions. Assuming the Strong Exponential Time Hypothesis (SETH), for any $\epsilon > 0$, there exists a constant c such that when the dimension $m = c \log n$, then there is no $2^{o(m)} \cdot n^{2-\epsilon}$ -time algorithm for the Bichromatic Closest Pair Problem [5, 1, 51].

1.4 Organization

The rest of the paper is organized as follows. Preliminaries and notations that we use throughout the paper are summarized in Section 2. In Section 3 we present our main decoding-based algorithms for the Closest Pair Problem, assuming the minimum pairwise distance is given. We then show how to get rid of this assumption in Section 4. In Section 5, we apply our new algorithms to study the Light Bulb Problem. Finally, we conclude with several open problems in Section 6.

2 Preliminaries

Let $m \geq 1$ be a natural number, we use $[m]$ to denote the set $\{1, \dots, m\}$. All logarithms in this paper are base 2 unless specified otherwise.

The binary entropy function, denoted $H_2(p)$, is defined as $H_2(p) := -p \log p - (1 - p) \log(1 - p)$ for $0 \leq p \leq 1$.

Our algorithms therefore do not give improved bounds for these learning problems in the general settings.

Let \mathbb{F}_q be a finite field with q elements¹² and $m \geq 1$ be a natural number. If $x \in \mathbb{F}_q^m$ is an m -dimensional vector over \mathbb{F}_q and $i \in [m]$, then we use $(x)_i$ to denote the i^{th} coordinate of x . The *Hamming distance* between two vectors $x, y \in \mathbb{F}_q^m$ is the number of coordinates at which they differ: $\text{dist}(x, y) = |\{i \in [m] : (x)_i \neq (y)_i\}|$. For a vector $x \in \mathbb{F}_q^m$ and a real number $r \geq 0$, the Hamming ball of radius r around x is $B(x, r) = \{y \in \mathbb{F}_q^m : \text{dist}(x, y) \leq r\}$. The *weight* of a vector x , denoted $\text{wt}(x)$, is the number of coordinates at which $(x)_i \neq 0$. The distance between two vectors x and y is easily seen to be equal to $\text{wt}(x - y)$.

We also need the following bounds on binomial coefficients, see e.g. [34, p. 309].

Lemma 2.1. *Let n be a natural number and λn be an integer, where $0 < \lambda < 1$. Then*

$$\frac{1}{\sqrt{8n\lambda(1-\lambda)}} 2^{nH_2(\lambda)} \leq \binom{n}{\lambda n} \leq \frac{1}{\sqrt{2\pi n\lambda(1-\lambda)}} 2^{nH_2(\lambda)}.$$

2.1 Error correcting codes

Definition 2.2 (Error correcting codes). *Let \mathbb{F}_q be a finite field with q elements¹³ and let $m \geq 1$ be a natural number. A subset C of \mathbb{F}_q^m is called an $(m, K, d)_q$ -code if $|C| = K$ and for any two distinct vectors $x, y \in C$, $\text{dist}(x, y) \geq d$. The vectors in C are called codewords of C , m the block length of C , and d the minimum distance of C .*

Normalized by the block length m , $\kappa(C) := (\log_q K)/m$ is known as the *rate* of C and $\delta(C) := d/m$ is known as the *relative distance* of C . If C is a linear subspace of \mathbb{F}_q^m of dimension k , the code is called a *linear code* and denoted by $[m, k, d]_q$. It is convenient to view such a linear code as the image of an *encoding function* $E : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^m$, and k is called *message length* of C . This can be generalized to non-linear codes as well where we view $\lfloor \log_q K \rfloor$ as the effective message length. We usually drop the subscript q when $q = 2$.

Definition 2.3 (Covering radius). *Let $C \subseteq \mathbb{F}_q^m$ be a code. For any $x \in \mathbb{F}_q^m$, define the distance between x and C to be $\text{dist}(x, C) := \min_{y \in C} \text{dist}(x, y)$ (clearly, $\text{dist}(x, C) = 0$ if and only if x is a codeword of C). The covering radius of a code C , denoted $R(C)$, is defined to be the maximum distance of any vector in \mathbb{F}_q^m from C , i.e., $R(C) = \max_{x \in \mathbb{F}_q^m} \text{dist}(x, C)$.*

2.1.1 Unique decoding

Given an (m, K, d) -code C , if a vector (aka received word) $x \in \mathbb{F}_q^m$ is at a distance $r \leq \lfloor \frac{d-1}{2} \rfloor$ from some codeword w in C , then by triangle inequality, x is closer to w than any other codewords in C . Therefore x can be *uniquely* decoded to the codeword $c \in C$. Such a decoding scheme¹⁴ is called *unique decoding* (or minimum distance decoding) of code C , and we shall call $\lfloor \frac{d-1}{2} \rfloor$ the (unique) *decoding radius* of C .

2.1.2 Gilbert-Varshamov bound and Gilbert's greedy code

The Gilbert-Varshamov bound asserts that there is an infinite family of codes C (essentially random codes or even random linear codes meet this bound almost surely) that satisfy $\kappa(C) \geq 1 - H_2(\delta(C))$. In particular,

¹²When $q = 2$, we use \mathbb{F}_2 and $\{0, 1\}$ interchangeably throughout the paper.

¹³In fact, error correcting codes, as well as constructing new codes out of existing codes by concatenations to be discussed shortly, can be defined more generally over an arbitrary set of q distinct elements called *alphabet* of the code. For the purpose of designing algorithms in this paper, restricting to finite fields is simpler and sufficient.

¹⁴Strictly speaking, the procedure described here is *error correcting* instead of *decoding*, where the latter should return the inverse of codeword c of the encoding function.

the following greedy algorithm of Gilbert [23] finds a (non-linear) binary code C of block length m and minimum distance d and satisfies that $\frac{1}{m} \log K \geq 1 - H_2(d/m) - \epsilon$ for any $\epsilon > 0$ for all sufficiently large m . Start with $S = \mathbb{F}_2^m$ and $C = \emptyset$; while $S \neq \emptyset$, pick any element $x \in S$, add it to C and remove all the elements in $B(x, d)$ from S . We denote such a code by $\text{GV}_{m,d}$.

We will need the following simple facts about $\text{GV}_{m,d}$

Lemma 2.4. *The greedy algorithm of Gilbert can be implemented to run in $O(2^m)$ time, and produces a decoding lookup table that supports constant time unique decoding. That is, for any $x \in \mathbb{F}_2^m$, if there is a codeword $w \in \text{GV}_{m,d}$ with $\text{dist}(x, w) \leq \lfloor \frac{d-1}{2} \rfloor$, then the lookup entry of x is w ; otherwise the entry is a special symbol, say, \perp . Moreover, the code $\text{GV}_{m,d}$ constructed by Gilbert's greedy algorithm satisfies that $R(\text{GV}_{m,d}) \leq d$.*

2.1.3 Reed-Solomon codes

Definition 2.5 (Reed-Solomon codes). *Let \mathbb{F}_q be finite field, k and m be integers satisfying $k \leq m \leq q$. The encoding function for Reed-Solomon code from \mathbb{F}_k to \mathbb{F}_m is the following: First pick m distinct elements $\alpha_1, \dots, \alpha_m \in \mathbb{F}_q$; on input $(a_0, a_1, \dots, a_{k-1}) \in \mathbb{F}_q^k$, define a degree- $k-1$ polynomial $P : \mathbb{F}_q \rightarrow \mathbb{F}_q$ as $P(x) = \sum_{i=0}^{k-1} a_i x^i$; finally output the evaluations of $P(x)$ at $\alpha_1, \dots, \alpha_m$, i.e. the codeword is $(P(\alpha_1), \dots, P(\alpha_m))$. We will denote such a code by $\text{RS}_{q,m,k}$.*

Theorem 2.6. *The Reed-Solomon code defined above is an $[m, k, m - k + 1]_q$ linear code.*

Theorem 2.7 ([48]). *There exists an efficient unique decoding algorithm for Reed-Solomon codes which runs in time $\text{poly}(m, \log q)$.*

Reed-Solomon codes are optimal in the sense that they meet the Singleton bound, which states that for any linear $[m, k, d]_q$ -code, $k \leq m - d + 1$.

2.1.4 Concatenated codes

The most commonly used way to transform a nice code which has constant rate and constant relative distance over a large alphabet to a similarly nice code over binary is *concatenation*, which was first introduced by Forney [21].

Definition 2.8 (Concatenated codes). *Let C_1 be an $(m_1, K_1, d_1)_{\mathbb{F}_Q}$ -code and let C_2 be an $(m_2, K_2, d_2)_{\mathbb{F}_q}$ -code with $K_2 \geq Q$. Then the code obtained by concatenating C_1 with C_2 , denoted by $C = C_1 \diamond C_2$, is an $(m, K, d)_{\mathbb{F}_q}$ -code defined as follows. Let ϕ be any mapping from \mathbb{F}_Q onto \mathbb{F}_q . Then the codewords of $C_1 \diamond C_2$ are obtained by replacing each element in \mathbb{F}_Q of any codeword $w = ((w)_1, \dots, (w)_{m_1}) \in C_1$ with the corresponding codeword in C_2 defined by ϕ ; namely $C = \{\phi((w)_1) \circ \dots \circ \phi((w)_{m_1}) : w \in C_1\}$, where each $\phi((w)_j)$ consists of m_2 elements in \mathbb{F}_q and \circ denotes string concatenation. Note that each codeword in C is an element in $\mathbb{F}_q^{m_1 m_2}$ and there are K_1 such codewords, therefore $m = m_1 m_2$ and $K = K_1$. Usually C_1 is called the outer code and C_2 is called the inner code.*

It is well-known that the minimum distance of C is $d_1 d_2$, and the rate of C is $\kappa(C) = \kappa(C_1) \kappa(C_2)$. Another useful fact is that C can be efficiently decoded as long as both C_1 and C_2 can be efficiently decoded.

Fact 2.9. *Suppose C_1 is an $(m_1, K_1, d_1)_{\mathbb{F}_Q}$ -code with a decoding algorithm A_1 running in $p_1(m_1, \log Q)$ time, C_2 is an $(m_2, K_2, d_2)_{\mathbb{F}_q}$ -code, where $K_2 \geq Q$, and a decoding algorithm A_2 running in $p_2(m_2, \log q)$ time. If C is the concatenated code $C = C_1 \diamond C_2$, and then there is a decoding algorithm A for C which run in time $p_1(m_1, \log Q) + m_1 p_2(m_2, \log q)$ by first decoding m_1 received words of C_2 each consisting of m_2 elements in \mathbb{F}_q , and then decode the m_1 concatenated elements in \mathbb{F}_Q as a received word of C_1 .*

2.1.5 Codes used in our algorithms

Some of the codes to be employed in our algorithm are a family of codes constructed by concatenating Reed-Solomon codes with certain binary non-linear Gilbert's greedy codes meeting the Gilbert-Varshamov bound. It is well-known that concatenated codes such constructed can be made to meet the so-called *Zyablov bound*¹⁵

$$\kappa(C) \geq \max_{0 < \kappa(C_2) < 1 - H_2(\delta(C))} \kappa(C_2) \left(1 - \frac{\delta(C)}{H_2^{-1}(1 - \kappa(C_2))} \right) \quad (1)$$

Suppose we want a binary (m, K, d) -code for our algorithms, where m and d are fixed and our goal is to maximize K , conditioned on that the code is efficiently decodable. We pick a Reed-Solomon code $C_1 = \text{RS}_{q, m_1, k_1}$ and a Gilbert's greedy code $C_2 = \text{GV}_{m_2, d_2}$ with the following constraints: $m_1 m_2 \leq m$ ($m_1 m_2$ should be as close to m as possible), $d_1 d_2 \geq d$, $K_2 = 2^{m_2 \kappa(C_2)} \geq q > m_1$, and $2^{m_2} \leq \text{poly}(m_1)$. It is easy to check that there are large ranges of values for m_1 and m_2 , and optimizing the choice of d_2 (and therefore $\delta(C_2)$) makes our concatenated code $C = C_1 \diamond C_2$ both meets the Zyablov bound in Eqn. (1) and can be decoded in $\text{poly}(m)$ time.

We will denote the maximum rate as a function of the relative distance δ given by the Zyablov bound by $\kappa_Z(\delta)$, and similarly denote the maximum rate given by the Gilbert-Varshamov bound by $\kappa_{GV}(\delta)$ (i.e. $\kappa_{GV}(\delta) = 1 - H_2(\delta)$). Note that $\kappa_Z(\delta) \leq \kappa_{GV}(\delta)$ for all $0 \leq \delta \leq 1/2$, and the reason we use codes achieving only $\kappa_Z(\delta)$ is because such codes can be generated and decoded in $\text{poly}(m)$ time.

2.2 The Closest Pair Problem

Given n vectors x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$, the Closest Pair Problem is to find two vectors whose pairwise Hamming distance is minimum. For ease of exposition and without loss of generality, we will assume throughout the paper that there is a unique pair, namely x_0 and x_1 , that achieves the minimum pairwise distance d_{\min} . We will use d_2 to denote the second minimum pairwise distance, where $d_2 \geq d_{\min} + 1$. In the most general case, we do not make any assumption about m , d_{\min} or d_2 .

3 Main Algorithm for the Closest Pair Problem

We now present our Main Algorithm for the Closest Pair Problem. For ease of exposition, we make a somewhat unnatural assumption that the value of d_{\min} is given. However, as we show in Section 4, the algorithm can be modified to get rid of this assumption, with only a slight slowdown in running time.

Theorem 3.1 (Non-gapped version). *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} (and the second smallest distance can be as small as $d_{\min} + 1$). Suppose we are given the value of d_{\min} and let $\delta := d_{\min}/m$. Then there is a randomized algorithm running in $O(n \log^2 n \cdot 2^{(1-\kappa_Z(\delta)-\delta)m} \cdot \text{poly}(m))$ which finds the closest pair x_0 and x_1 with probability at least $1 - 1/n^2$.*

Proof. Our Main Algorithm for the Closest Pair problem is described in Algorithm 2, and the decoding subroutine $\text{Dec}(C, r, x)$ is illustrated in Algorithm 3. Note that we choose the minimum distance of C to be $d_{\min} + 1$, hence the decoding radius of C is $d_{\min}/2$ (without loss of generality, assume that d_{\min} is even).

¹⁵In fact, a stronger bound called *Blokh-Zyablov bound* can be achieved by applying multilevel concatenations (see e.g. [20] for a detailed discussion on multilevel concatenations of codes); however, as the improvement is minor, we only use single level concatenation in our code constructions to make the algorithms simpler.

For the correctness of the algorithm, first note that our algorithm will output the correct minimum distance if and only if x_0 is ever compared against x_1 for computing pairwise distance, and this happens if and only if x_0 and x_1 are adjacent in the sorted array after decoding. A sufficient condition for the latter is that the decoded vectors of x_0 and x_1 are identical and they are different from any other decoded vectors.

How many shift vectors $y \in \{0, 1\}^m$ in Algorithm 2 satisfy this condition? We will call such vectors *good* vectors. Denote the set of vectors lying at the “middle” between x_0 and x_1 by

$$\text{MID} = \{z \in \{0, 1\}^m : \text{dist}(x_0, z) = \text{dist}(z, x_1) = d_{\min}/2\}.$$

Note that any vector y that shifts a vector $z \in \text{MID}$ to a codeword $c \in C$ would be a *good* vector. To see this, first note that after such a shift, $y + z$ is a codeword in C , and both $y + x_0$ and $y + x_1$ lie within the decoding radius of $y + z$, and therefore will be decoded to $y + z$. Moreover, the shifted vector of any other input vector $y + x_i$, $2 \leq i \leq n - 1$, lies outside the decoding radius of $y + z$. This is because if it does, then by triangle inequality and the fact that the decoding radius of C is $d_{\min}/2$,

$$\begin{aligned} \text{dist}(x_0, x_i) &= \text{dist}(y + x_0, y + x_i) \\ &\leq \text{dist}(y + x_0, y + z) + \text{dist}(y + z, y + x_i) \\ &\leq d_{\min}/2 + d_{\min}/2 = d_{\min}, \end{aligned}$$

contradicting our assumption that x_0 and x_1 is the unique pair achieving the minimum distance.

How many such *good* vectors? There are in total $\binom{d_{\min}}{d_{\min}/2}$ vectors exist in MID, and all their pairwise distances are at most d_{\min} . Let c_1, c_2 be two distinct codewords in C . By our choice of the minimum distance of C , $\text{dist}(c_1, c_2) > d_{\min}$. Consider any two distinct vectors z_1 and z_2 in MID. Clearly applying these two shift vectors to the same codeword gives two distinct vectors, namely $c_1 + z_1$ and $c_1 + z_2$. Moreover, applying two distinct vectors in MID to two distinct codewords also results in two distinct shift vectors, because

$$\text{dist}(c_1 + z_1, c_2 + z_2) = \text{wt}(c_1 + c_2 + z_1 + z_2) > 0,$$

since $\text{wt}(c_1 + c_2) \geq d > d_{\min}$ but $\text{wt}(z_1 + z_2) = \text{dist}(z_1, z_2) \leq d_{\min}$.

Recall that C is a (m, K, d) -code and hence there are K codewords in C . It follows that there are in total $K \cdot \binom{d_{\min}}{d_{\min}/2}$ *good* vectors of this kind. Therefore

$$\Pr(\text{a random } y \text{ succeeds in finding the closest pair}) \geq \frac{K \cdot \binom{d_{\min}}{d_{\min}/2}}{2^m},$$

and hence repeatedly selecting

$$\begin{aligned} 2 \ln n \frac{2^m}{K \cdot \binom{d_{\min}}{d_{\min}/2}} &= O\left(\log n \frac{\sqrt{\delta m} 2^m}{2^{\kappa_Z(\delta) m} 2^{\delta m}}\right) \\ &= O(2^{(1-\kappa_Z(\delta)-\delta)m} m^{1/2} \log n) \end{aligned}$$

independent y 's will succeed with probability at least $1 - 1/n^2$, where in the last step we use the bound $\binom{n}{n/2} = O(\frac{2^n}{\sqrt{n}})$, a special case of Lemma 2.1.

Finally, note that each choice of shift vector y requires $n \cdot \text{poly}(m)$ time decoding as well as $O(n \log n \cdot m)$ sorting and comparing adjacent vectors, so the total running time of the algorithm is $O(n \log^2 n \cdot 2^{(1-\kappa_Z(\delta)-\delta)m} \cdot \text{poly}(m))$. \square

Algorithm 2: Main Algorithm for the Closest Pair Problem

input : A set of n vectors x_0, \dots, x_{n-1} in $\{0, 1\}^m$ and d_{\min}
output: Two vectors and their distance

- 1 generate a binary (m, K, d) -code C with $d = d_{\min} + 1$
- 2 **for** $i \leftarrow 1$ **to** $O(2^{(1-\kappa_Z(\delta)-\delta)m} m^{1/2} \log n)$ **do**
- 3 pick a random $y \in \{0, 1\}^m$
- 4 **for** $j \leftarrow 0$ **to** $n - 1$ **do**
- 5 $\tilde{x}_j \leftarrow \text{Dec}(C, \lfloor d_{\min}/2 \rfloor, y + x_j)$
- 6 **end**
- 7 sort $\tilde{x}_0, \dots, \tilde{x}_{n-1}$
- 8 (suppose the sorted sequence is $\tilde{x}_{s_0}, \dots, \tilde{x}_{s_{n-1}}$, where $\{s_0, \dots, s_{n-1}\}$ is a permutation of $\{0, 1, \dots, n-1\}$)
- 9 **for** $j \leftarrow 1$ **to** $n - 1$ **do**
- 10 compute $\text{dist}(x_{s_{j-1}}, x_{s_j})$
- 11 **end**
- 12 **end**
- 13 output the pair of vectors with minimum distance ever found and their distance

Algorithm 3: Dec(C, r, x)

input : A binary (m, K, d) -code C , a decoding radius $r < d/2$, and a vector $x \in \{0, 1\}^m$
output: A vector $\tilde{x} \in \{0, 1\}^m$

- 1 run the (efficient) decoding algorithm for C on input vector x , and let the output vector be \tilde{x}
- 2 **if** $\text{dist}(x, \tilde{x}) \leq r$ **then**
- 3 output \tilde{x}
- 4 **else**
- 5 output x
- 6 **end**

If we assume further that a decoding algorithm for some binary (m, K, d) -code C which meets the Gilbert-Varshamov bound is given as a black box, then the running time in Theorem 3.1 can be improved to $O(n \log^2 n \cdot 2^{(H_2(\delta)-\delta)m} \cdot \text{poly}(m))$. Note that this is not a totally unrealistic assumption, as for most interesting settings, $m = c \log n$ for some small constant c .¹⁶ Therefore, greedily searching for a binary code of block length m that meets the Gilbert-Varshamov bound is tantamount to running an $O(n^c)$ time preprocessing, which can be reused for any problem instance with the same vector length and minimum closest pair distance.

If there is a gap between d_2 and d_{\min} (this roughly corresponds to the *approximate closest pair* problem in [45]), then we can improve the running time of the Main Algorithm in Theorem 3.1 by exploiting an error correcting code with larger decoding radius.

Theorem 3.2 (Gapped version). *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} . Suppose we are given the values of d_{\min}*

¹⁶As in the settings of random vectors, e.g. the Light Bulb Problem, $m = c \log n$ is both necessary and sufficient to distinguish n stochastic bit sequences.

as well as the second minimum distance d_2 . Let $\delta := d_{\min}/m$ and $\delta' := d_2/m$. Then there is a randomized algorithm running in $O(n \log^2 n \cdot 2^{(1-\kappa_Z(\delta'))-\delta-(1-\delta)H_2(\frac{\delta'-\delta}{2(1-\delta)})m} \cdot \text{poly}(m))$ which finds the closest pair x_0 and x_1 with probability at least $1 - 1/n^2$. Moreover, the running time can be further improved to $O(n \log^2 n \cdot 2^{(H_2(\delta'))-\delta-(1-\delta)H_2(\frac{\delta'-\delta}{2(1-\delta)})m} \cdot \text{poly}(m))$, if we are given black box access to the decoding algorithm of an (m, K, d) -code which meets the Gilbert-Varshamov bound.

Proof. The proof follows a similar structure as the proof of Theorem 3.1. The main difference is now we pick a binary error correcting code of minimum distance $d_2 + 1$, thereby decoding radius $r = d_2/2 = \frac{1}{2}\delta'm$ (once again, for simplicity, we assume d_2 is even).

Accordingly, the “middle point” set is now defined as

$$\text{MID}_G = \{z \in \{0, 1\}^m : \text{dist}(x_0, z) \leq r \text{ and } \text{dist}(x_1, z) \leq r\}.$$

We now give a lower bound on the size of MID_G .

Without loss of generality, we assume $x_0 = 0^m$ and let $T = \{i \in [m] : (x_1)_i = 1\}$. Clearly $|T| = d_{\min}$. Let $i = |\{k \in T : (z)_k = 0\}|$ and $j = |\{k \in [m] \setminus T : (z)_k = 1\}|$. Then $\text{dist}(x_0, z) \leq r$ is equivalent to $d_{\min} - i + j \leq r$, and $\text{dist}(x_1, z) \leq r$ is equivalent to $i + j \leq r$. Therefore

$$\begin{aligned} |\text{MID}_G| &= \sum_{i+j \leq r} \sum_{d_{\min}-i+j \leq r} \binom{d_{\min}}{i} \binom{m-d_{\min}}{j} \\ &\geq \binom{d_{\min}}{d_{\min}/2} \binom{m-d_{\min}}{r-d_{\min}/2} \\ &= \Theta \left(\frac{2^{\delta m} 2^{(1-\delta)H_2(\frac{\delta'-\delta}{2(1-\delta)})m}}{\sqrt{\delta m} \sqrt{(1-\delta)m}} \right), \end{aligned}$$

where the last step follows from Lemma 2.1. The rest of the proof is identical to that of Theorem 3.1, and therefore is omitted. \square

3.1 A deterministic variant of the Main Algorithm

One can turn our randomized Main Algorithm into a deterministic one by exhaustively searching for all possible shift vectors $y \in \mathbb{F}_2^m$. A simple observation is that it suffices to check for all vectors in the Hamming ball of radius equals to the covering radius of the code C .

Theorem 3.3. *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} (and the second smallest distance can be as small as $d_{\min} + 1$). Suppose we are given the value of d_{\min} and let $\delta := d_{\min}/m$. Then there is a deterministic algorithm that finds the closest pair x_0 and x_1 with running time $O(n \log n \cdot 2^{H_2(1-\kappa_Z(\delta))m} \cdot \text{poly}(m))$. Moreover, if we are given as black box the decoding algorithm of a random Varshamov linear code with block length m and minimum distance $d_{\min} + 1$, then the running time is $O(n \log n \cdot 2^{H_2(H_2(\delta))m} \cdot \text{poly}(m))$.*

Proof. Let $\delta := d_{\min}/m$. It is well-known that for any linear $[m, k, d]_q$ -code C , the covering radius of C satisfies that $R(C) \leq m - k$. It follows that for Reed-Solomon code $\text{RS}_{q,m,k}$, $R(\text{RS}) \leq m - k < d$. We can either generate a random linear Varshamov code [47] similar to that described in Section 2.1.5 that meets the Gilbert-Varshamov bound and concatenate it with a Reed-Solomon code so that the resulting binary code is

a linear code. Then the covering radius of this concatenated code satisfies that $R(C) \leq (1 - \kappa_z(\delta))m$. Or, if preprocessing is allowed, we may simply generate a random linear Varshamov code of block length m , whose covering radius satisfies that $R(C) \leq (1 - \kappa_{GV}(\delta))m = H_2(\delta)m$.

Now the deterministic algorithm for finding the closest pair is similar to the Main Algorithm, except that instead of picking random shift vector y , the algorithm checks every $y \in B(0^m, R(C))$. It follows directly that the running time of the algorithm is $O(n \log n \cdot \text{poly}(m) \cdot \text{Vol}(B(0^m, R(C))))$. Here $\text{Vol}(B(0^m, R(C)))$ denotes the number of vectors within the Hamming ball $B(0^m, R(C))$, which is $2^{H_2(1 - \kappa_z(\delta))m}$ for the concatenated code, or $2^{H_2(H_2(\delta))m}$ for the random Varshamov linear code.

The correctness of the algorithm follows that, by the same argument of the correctness of Algorithm 2, any vector $z \in \text{MID}$ is at most $R(C)$ away from some codeword $c \in C$, namely $\text{dist}(z, c) = \text{wt}(z + c) \leq R(C)$. When vector $z + c$, which lies in $B(0^m, R(C))$, is chosen as the shift vector y , x_0 and x_1 will be the only two vectors decoded to c , therefore the algorithm successfully finds the closest pair. \square

We remark that our covering radius argument seems to be too rough, as there are many vectors in MID. Getting a more efficient deterministic algorithm, or derandomizing the Main Algorithm is an interesting open question of combinatorics in nature.

4 Searching for the Minimum Distance

In this section we show how to remove the assumption that the value of d_{\min} is given to the Main algorithm. Basically we show that one can use a binary-search like procedure to find d_{\min} without too much slowdown of the Main Algorithm. Our key observation is that, although the decoding radius is chosen to be $d_{\min}/2$ in the Main Algorithm, actually we can relax this requirement: indeed, any decoding radius between $d_{\min}/2$ and d_{\min} works.

Lemma 4.1. *The Main Algorithm works (with worse running time) as long as the binary error correcting code used has decoding radius $r = \lfloor \frac{d-1}{2} \rfloor$ satisfying $\frac{1}{2}d_{\min} \leq r \leq d_{\min}$.*

Proof. The proof is similar to the proof of Theorem 3.1, but we slightly generalize the original definition of MID as follows. Let

$$\text{MID}_1 = \{z \in \{0, 1\}^m : \text{dist}(x_0, z) = r \text{ and } \text{dist}(x_1, z) = d_{\min} - r\},$$

let

$$\text{MID}_2 = \{z \in \{0, 1\}^m : \text{dist}(x_1, z) = r \text{ and } \text{dist}(x_0, z) = d_{\min} - r\}.$$

and finally let $\text{MID}' = \text{MID}_1 \cup \text{MID}_2$.

Clearly the set MID' is non-empty. The key point is that any vector y that shifts some vector $z \in \text{MID}'$ to a codeword $c \in C$ must be a *good* vector, following a similar argument as in the proof of Theorem 3.1. The running time of the algorithm can also be calculated similarly. \square

Our algorithm for finding d_{\min} is illustrated in Algorithm 4. The correctness of Algorithm 4 follows from two simple facts: first, Algorithm 2 can never return a value less than d_{\min} ; second, when $d_{\min}/2 \leq r \leq d_{\min}$, by Lemma 4.1 and Theorem 3.1, Algorithm 2 returns the correct value of d_{\min} (with high probability).

In fact, to make our algorithm more efficient, for any $0 < \epsilon \leq 1$, we can search with decoding radius $r = 1, \lfloor (1 + \epsilon) \rfloor, \lfloor (1 + \epsilon)^2 \rfloor, \dots$. Note that by Lemma 4.1, the maximum value we will ever try is $(1 + \epsilon)d_{\min}/2$. As the running time of Algorithm 2 is monotone increasing with respect to the decoding radius r , so in order

Algorithm 4: Searching for d_{\min} **input :** A set of n vectors x_0, \dots, x_{n-1} in $\{0, 1\}^m$ **output:** The minimum pairwise distance d_{\min}

```

1  $r \leftarrow 1$ 
2 while true do
3   run Algorithm 2 with  $d = r$ ;
4   if the minimum distance  $d_{\min}$  returned in Algorithm 2 is at most  $2r$ 
5     then
6       | return  $d_{\min}$ ;
7     else
8       |  $r \leftarrow (1 + \epsilon)r$ ;
9     end
10 end

```

to bound the running time of searching for d_{\min} , it suffices to bound the running time of Algorithm 2 for $r = (1 + \epsilon)d_{\min}/2$. Following a similar analysis as in the proof of Theorem 3.1,

$$|\text{MID}'| \geq |\text{MID}_1| = \binom{d_{\min}}{(1 + \epsilon)d_{\min}/2}.$$

Therefore,

$$\begin{aligned} & \Pr(\text{a random } y \text{ succeeds in finding the closest pair}) \\ & \geq \frac{K \cdot \binom{d_{\min}}{(1 + \epsilon)d_{\min}/2}}{2^m} \\ & = \Theta\left(\frac{2^{\kappa_Z((1 + \epsilon)\delta)m} 2^{H_2(\frac{1 - \epsilon}{2})\delta m}}{\sqrt{\delta m} 2^m}\right), \end{aligned}$$

where in the last step we use bounds in Lemma 2.1. As the binary search calls at most $\log_{(1 + \epsilon)} d_{\min} < \log_{(1 + \epsilon)} m = O(\epsilon^{-1} \log m)$ times Algorithm 2, we therefore have the following theorem.

Theorem 4.2. *Let x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$ be n binary vectors such that x_0 and x_1 is the unique pair achieving the minimum pairwise distance d_{\min} . Then for any $\epsilon > 0$, there is a randomized algorithm running in :*

$$O(\epsilon^{-1} n \log^2 n \cdot 2^{(1 - \kappa_Z((1 + \epsilon)\delta) - \delta H_2(\frac{1 - \epsilon}{2}))m} \cdot \text{poly}(m)).$$

which finds the d_{\min} (as well as the closest pair x_0 and x_1) with probability at least $1 - 1/n^2$. The running time can be improved to :

$$O(\epsilon^{-1} n \log^2 n \cdot 2^{(H_2((1 + \epsilon)\delta) - \delta H_2(\frac{1 - \epsilon}{2}))m} \cdot \text{poly}(m)),$$

if we are given black-box decoding algorithms for an ensemble of $O(\epsilon^{-1} \log m)$ binary error correcting codes that meet the Gilbert-Varshamov bound.

5 The Light Bulb Problem

In this section, we apply our new algorithms for the Closest Pair Problem to a special case of it, namely the *Light Bulb Problem*.

In the Light Bulb Problem, we are given n sequences of bit strings X_0, X_1, \dots, X_{n-1} . All bits are generated independently, uniformly at random from $\{0, 1\}$, except that two strings, say X_0 and X_1 , are generated with non-zero linear correlation ρ ; that is, independently for each i , $\Pr((X_0)_i = (X_1)_i) = \frac{1+\rho}{2}$ and $\Pr((X_0)_i \neq (X_1)_i) = \frac{1-\rho}{2}$. The problem is to find this correlated pair of sequences.

First note that we may assume the Pearson correlation ρ is positive, as there is a simple randomized reduction from the negative ρ case to the positive ρ case: given an instance of the Light Bulb Problem with $\rho < 0$ randomly pick $n/2$ sequences and flip all the bits in these sequences. Then with probability $1/2$, the correlated pair become $-\rho$ correlated.

To apply our algorithms for the Closest Pair Problem to the Light Bulb Problem, the following standard result¹⁷ provides a randomized reduction from the latter to the former. We include a proof for completeness and to justify our choice of the dimension m .

Theorem 5.1. *If we pick $m = \frac{4 \ln 2 \cdot \log n}{\rho^2} (1 + o(1))$ bits at random from X_0, X_1, \dots, X_{n-1} to obtain n vectors x_0, x_1, \dots, x_{n-1} in $\{0, 1\}^m$, then with constant probability, x_0 and x_1 is the unique closest pair among these n vectors.*

Proof. For each pair of vectors x_i and x_j , $0 \leq i < j \leq n - 1$, define m indicator random variables $\{(I_{i,j})_k\}_{k \in [m]}$ such that $(I_{i,j})_k = 1$ if and only if $(x_i)_k \neq (x_j)_k$. Note that for any pair $i < j$, $\{(I_{i,j})_k\}_{k \in [m]}$ are m independent and identically distributed random variables, and $\text{dist}(x_i, x_j) = \sum_{k \in [m]} (I_{i,j})_k$. Specifically, $\Pr((I_{0,1})_k = 0) = \frac{1+\rho}{2}$ and $\Pr((I_{0,1})_k = 1) = \frac{1-\rho}{2}$; and $\Pr((I_{i,j})_k = 0) = \Pr((I_{i,j})_k = 1) = 1/2$ for all other $i < j$ pairs.

Note that each pairwise distance $\text{dist}(x_i, x_j)$ is a *binomial* random variable. In particular, $\text{dist}(x_0, x_1)$ is a $B(m, \frac{1-\rho}{2})$ random variable and all others are $B(m, 1/2)$ random variable. To argue about the distribution of distance between x_0 and x_1 , we need the following fact:

Fact 5.2 ([28]). *Binomial distribution $B(n, p)$ has median $\lfloor np \rfloor$ or $\lceil np \rceil$.*

Let $d_t := \mathbb{E}(\text{dist}(x_0, x_1)) = (1 - \rho)m/2$. Then by Fact 5.2, $\Pr(\text{dist}(x_0, x_1) \geq d_t) \leq 1/2$.

On the other hand, for any other pair x_i and x_j ,

$$\begin{aligned} \Pr(\text{dist}(x_i, x_j) < d_t) &= \Pr(\text{dist}(x_i, x_j) < \mathbb{E}(\text{dist}(x_i, x_j)) - \rho m/2) \\ &< e^{-(\rho m)^2/2m} = e^{-m\rho^2(1-o(1))/2} \\ &\leq \frac{1}{2n^2}, \end{aligned}$$

by a simple application of the Chernoff bound (e.g. Theorem A.1.1 in [6]). Now applying a union bound over all x_i and x_j pairs, we have that with probability at least $1/4$, $\text{dist}(x_0, x_1) < d_t$ and for all other pairs $\text{dist}(x_i, x_j) \geq d_t$, i.e., x_0 and x_1 is the unique closest pair among these n vectors. \square

Note that Theorem 5.1 implies that if we sample $m = \frac{4 \ln 2 \cdot \log n}{\rho^2} (1 + o(1))$ bits from the n random sequences, then with constant probability, we get an instance of the Closest Pair Problem with $d_{\min} <$

¹⁷This is a folklore bound. Indeed, similar analyses can be found in earlier work, although sometimes with slightly different tools (e.g., Hoeffding bound in place of Chernoff bound), but essentially they all aim to show that with a sufficiently high dimension, the planted correlation is unique, with high probability.

$(1 - \rho)m/2$. Now, by repeatedly running our randomized algorithm for Closest Pair Problem $O(\log n)$ times, each time taking independent samples from the input vectors, and then take a majority vote, then by combining a simple application of the Chernoff bound, Theorem 3.1 and Theorem 5.1, we obtain the following

Theorem 5.3. *There is a randomized algorithm for the Light Bulb Problem which runs in time*

$$O(n \cdot \text{poly}(\log n)) \cdot 2^{(1 - \kappa_Z(\frac{1-\rho}{2}) - \frac{1-\rho}{2}) \frac{4 \ln 2 \cdot \log n}{\rho^2} (1+o(1))}$$

and succeeds with probability at least $1 - 1/n^2$. The running time can be further improved to

$$O(n \cdot \text{poly}(\log n)) \cdot 2^{(H_2(\frac{1-\rho}{2}) - \frac{1-\rho}{2}) \frac{4 \ln 2 \cdot \log n}{\rho^2} (1+o(1))},$$

if we are allowed a one-time preprocessing time of $n^{2.773/\rho^2}$ to generate the decoding lookup table of a random Gilbert's $(m, K, (1 - \rho)m/2)$ -code.

Numerical calculations show that our new algorithm performs better than the improved Valiant's fast matrix multiplication algorithm [29] (which runs in $n^{1.582}$) when $\rho \geq 0.9967$ (equivalently when $\delta \leq 0.00165$). Moreover, if an $n^{2.773/\rho^2}$ -time preprocessing is allowed, then our algorithm runs faster for all $\rho \geq 0.9310$ (equivalently for all $\delta \leq 0.0345$).

Deterministic algorithm. Following [30], we say a deterministic algorithm solves the Light Bulb Problem if it is correct on *almost all instances*, i.e., if the algorithm fails on a randomly picked instance with probability at most $1/\text{poly}(n)$. Following a similar proof that of the randomized algorithm shown before, we have the following theorem on deterministic algorithm for the Light Bulb Problem

Theorem 5.4. *There is a deterministic algorithm for the Light Bulb Problem which runs in time*

$$O(n \cdot \text{poly}(\log n)) \cdot 2^{H_2(1 - \kappa_Z(\frac{1-\rho}{2})) \frac{4 \ln 2 \cdot \log n}{\rho^2} (1+o(1))}$$

and succeeds with probability at least $1 - 1/n^2$. The running time can be further improved to

$$O(n \cdot \text{poly}(\log n)) \cdot 2^{H_2(H_2(\frac{1-\rho}{2})) \frac{4 \ln 2 \cdot \log n}{\rho^2} (1+o(1))},$$

if we are allowed a one-time preprocessing time of $n^{2.773/\rho^2}$ to generate the decoding lookup table of a random Gilbert's $(m, K, (1 - \rho)m/2)$ -code.

Note that, like the randomized algorithm, our deterministic algorithm also needs to draw $O(m \log n)$ bits from each of the n sequences. However, the algorithm uses no random bits and the success probability is over the random instance the algorithm gets from the input.

As mentioned earlier, Alman [3] gave the currently best deterministic algorithm for the Light Bulb Problem, which runs in $O(n^{1.582})$ time. Unsurprisingly, the deterministic version of our algorithm outperforms the one in [3] when the Pearson correlation is very large. Specifically, by numerical calculation, our deterministic algorithm runs faster than Alman's when $\rho \geq 0.999948$. Moreover, if an $n^{2.773/\rho^2}$ -time preprocessing is allowed, then when $\rho \geq 0.9933$ (equivalently when $\delta \leq 0.0033$), we may take the vector length $m \leq 2.8101 \log n$ so that our deterministic algorithm runs in at most $O(n^{1.581})$.

6 Concluding Remarks and Open Problems

We propose a simple approach, namely a decoding-base method, to solve the classic Closest Pair Problem. Our results leave open several interesting questions. The way we derandomize our randomized algorithm is by a simple brute-force search. Is there a smarter and more efficient way to derandomize? Valiant’s fast matrix multiplication method [45] for the Light Bulb Problem is the only known algorithm that makes good use of the availability of larger amount of data. Is it possible to leverage the data size to improve the running time of our decoding approach? Another interesting open question is to study the Closest Pair Problem in the streaming model, as many real-life situations of the problem — such as in cyber security — are in fact in this setting.

Acknowledgements

We are most grateful to the anonymous referees for their detailed and invaluable comments and suggestions. We would like to thank Karthik C.S. for his comments.

References

- [1] A. Abboud, A. Rubinfeld, and R. Williams. Distributed PCP theorems for hardness of approximation in P. In *Proc. 58th Annual IEEE Symposium on Foundations of Computer Science*, pages 25–36, 2017.
- [2] A. Abboud, R. Williams, and H. Yu. More applications of the polynomial method to algorithm design. In *Proc. 26th ACM-SIAM Symposium on Discrete Algorithms*, pages 218–230, 2015.
- [3] J. Alman. An illuminating algorithm for the light bulb problem. In *Proc. 2nd Symposium on Simplicity in Algorithms*, pages 2:1–2:11, 2019.
- [4] J. Alman, Timothy M. Chan, and R. Williams. Polynomial representations of threshold functions and algorithmic applications. In *Proc. 57th Annual IEEE Symposium on Foundations of Computer Science*, pages 467–476, 2016.
- [5] J. Alman and R. Williams. Probabilistic polynomials and Hamming nearest neighbors. In *Proc. 56th Annual IEEE Symposium on Foundations of Computer Science*, pages 136–150, 2015.
- [6] N. Alon and J. Spencer. *The Probabilistic Method*. John Wiley and Sons, third edition, 2008.
- [7] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM*, 51:117–122, 2008.
- [8] A. Andoni and P. Indyk. Nearest neighbors in high-dimensional spaces. In J. Goodman, J. O’Rourke, and Csaba D. Toth, editors, *Handbook of Discrete and Computational Geometry*. Chapman and Hall/CRC, 3rd edition, 2017.
- [9] A. Andoni, T. Laarhoven, I. Razenshteyn, and E. Waingarten. Optimal hashing-based time-space trade-offs for approximate near neighbors. In *Proc. 28th ACM-SIAM Symposium on Discrete Algorithms*, pages 47–66, 2017.

- [10] S. Arya and D. Mount. Computational geometry: Proximity and location. In D. P. Mehta and S. Sahni, editors, *Handbook of Data Structures and Applications*. Chapman and Hall/CRC, 2005.
- [11] C. Aston, D. Ralph, D. Lalo, S. Manjeshwar, B. Gramling, D. DeFreese, A. West, D. Branam, L. Thompson, M. Craft, et al. Oligogenic combinations associated with breast cancer risk in women under 53 years of age. *Human genetics*, 116(3):208–221, 2005.
- [12] A. Becker, A. Joux, A. May, and A. Meurer. Decoding random binary linear codes in $2^{n/20}$: How $1 + 1 = 0$ improves information set decoding. In *Proc. 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 520–536, 2012.
- [13] A. Benson and G. Ballard. A framework for practical parallel fast matrix multiplication. In *Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 42–53, 2015.
- [14] J. Bentley. Multidimensional divide-and-conquer. *Communications of the ACM*, 23(4):214–229, 1980.
- [15] J. Cho, D. Nicolae, L. Gold, C. Fields, M. LaBuda, P. Rohal, M. Pickles, L. Qin, Y. Fu, J. Mann, et al. Identification of novel susceptibility loci for inflammatory bowel disease on chromosomes 1p, 3q, and 4q: evidence for epistasis between 1p and IBD1. *Proceedings of the National Academy of Sciences*, 95(13):7502–7507, 1998.
- [16] K. Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Computing*, 17(4):830–847, 1988.
- [17] H. Cordell. Detecting gene \times gene interactions that underlie human diseases. *Nature Reviews Genetics*, 10(6):392–404, 2009.
- [18] M. Datar, N. Immorlica, P. Indyk, and V. Mirrokni. Locality-sensitive hashing scheme based on p -stable distributions. In *Proc. 20th Symposium on Computational Geometry*, pages 253–262, 2004.
- [19] M. Dubiner. Bucketing coding and information theory for the statistical high-dimensional nearest-neighbor problem. *IEEE Transactions on Information Theory*, 56(8):4166–4179, 2008.
- [20] I. Dumer. Concatenated codes and their multilevel generalizations. In V. Pless and W. C. Huffman, editors, *Handbook of coding theory*, volume 2, pages 1911–1988. Amsterdam, The Netherlands: Elsevier Science, 1998.
- [21] D. Forney. *Concatenated Codes*. MIT Press, 1966.
- [22] K. Frazer, D. Ballinger, D. Cox, D. Hinds, L. Stuve, R. Gibbs, J. Belmont, A. Boudreau, P. Hardenbol, S. Leal, et al. A second generation human haplotype map of over 3.1 million SNPs. *Nature*, 449(7164):851–861, 2007.
- [23] E. Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31(3):504–522, 1952.
- [24] M. J. Golin, R. Raman, C. Schwarz, and M. Smid. Simple randomized algorithms for closest pair problems. *Nordic Journal of Computing*, 2(1):3–27, 1995.
- [25] V. Guruswami, A. Rudra, and M. Sudan. Essential coding theory. Book draft in preparation, available at <https://cse.buffalo.edu/faculty/atri/courses/coding-theory/book/web-coding-book.pdf>, 2018.

- [26] J. Huang, L. Rice, D. Matthews, and R. van de Geijn. Generating families of practical fast matrix multiplication algorithms. In *Parallel and Distributed Processing Symposium (IPDPS), 2017 IEEE International*, pages 656–667, 2017.
- [27] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annual ACM Symposium on the Theory of Computing*, pages 604–613, 1998.
- [28] R. Kaas and J. Buhrman. Mean, median and mode in binomial distributions. *Statistica Neerlandica*, 34(1):13–18, 1980.
- [29] M. Karppa, P. Kaski, and J. Kohonen. A faster subquadratic algorithm for finding outlier correlations. In *Proc. 27th ACM-SIAM Symposium on Discrete Algorithms*, pages 1288–1305, 2016.
- [30] M. Karppa, P. Kaski, J. Kohonen, and P. Ó Catháin. Explicit correlation amplifiers for finding outlier correlations in deterministic subquadratic time. In *Proc. 24th Annual European Symposia on Algorithms*, pages 52:1–52:17, 2016.
- [31] S. Khuller and Y. Matias. A simple randomized sieve algorithm for the closest-pair problem. *Information and Computation*, 118(1):34–37, 1995.
- [32] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. In *Proc. 30th Annual ACM Symposium on the Theory of Computing*, pages 614–623, 1998.
- [33] F. LeGall. Faster algorithms for rectangular matrix multiplication. In *Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 514–523, 2012.
- [34] F.J. MacWilliams and N. J. A. Sloane. *The Theory of Error-correction Codes*. North Holland, 1977.
- [35] A. May and I. Ozerov. On computing nearest neighbors with applications to decoding of binary linear codes. In *Proc. 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 203–228, 2015.
- [36] S. Meiser. Point location in arrangements of hyperplanes. *Information and Computation*, 106(2):286–303, 1993.
- [37] R. Motwani, A. Naor, and R. Panigrahi. Lower bounds on locality sensitive hashing. In *Proc. 22nd Symposium on Computational Geometry*, pages 154–157, 2006.
- [38] S. Musani, D. Shriner, N. Liu, R. Feng, C. Coffey, N. Yi, H. Tiwari, and D. Allison. Detection of gene \times gene interactions in genome-wide association studies of human population data. *Human heredity*, 63(2):67–84, 2007.
- [39] R. O’Donnell, Y. Wu, and Y. Zhou. Optimal lower bounds for locality-sensitive hashing (except when q is tiny). *ACM Transactions on Computation Theory*, 6(1), 2014.
- [40] V. Pan. Fast feasible and unfeasible matrix multiplication, April 2018. <https://arxiv.org/abs/1804.04102>.
- [41] R. Paturi, S. Rajasekaran, and J. Reif. The light bulb problem. *Information and Coputation*, 117:187–192, 1995.

- [42] M. Rabin. Probabilistic algorithms. *Algorithms and Complexity*, pages 21–30, 1976.
- [43] M. Smid. Closest-point problems in computational geometry. In J. Sack and J. Urrutia, editors, *Handbook of computational geometry*, pages 877–935. Elsevier Science Publishing, 1997.
- [44] A. Stothers. *On the complexity of matrix multiplication*. PhD thesis, The University of Edinburgh, 2010.
- [45] G. Valiant. Finding correlations in subquadratic time, with applications to learning parities and juntas. *Journal of the ACM*, 62(13), 2015. Earlier version in FOCS’12.
- [46] L. G. Valiant. Functionality in Neural Nets. In *First Workshop on Computational Learning Theory*, pages 28–39, 1988.
- [47] R. R. Varshamov. Estimate of the number of signals in error correcting codes. *Dokl. Akad. Nauk SSSR*, 117(5):739–741, 1957.
- [48] L. Welch and E. Berlekamp. Error correction for algebraic block codes, December 30 1986. US Patent 4,633,470.
- [49] R. Williams. Faster all-pairs shortest paths via circuit complexity. In *Proc. 46th Annual ACM Symposium on the Theory of Computing*, pages 664–673, 2014.
- [50] R. Williams. The polynomial method in circuit complexity applied to algorithm design (invited talk). In *Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 47–60, 2014.
- [51] R. Williams. On the difference between closest, furthest, and orthogonal pairs: Nearly-linear vs barely-subquadratic complexity. In *Proc. 29th ACM-SIAM Symposium on Discrete Algorithms*, pages 1207–1215, 2018.
- [52] V. Williams. Multiplying matrices faster than Coppersmith-Winograd. In *Proc. 44th Annual ACM Symposium on the Theory of Computing*, pages 887–898, 2012.