# A FASTER SUBQUADRATIC ALGORITHM FOR FINDING OUTLIER CORRELATIONS 

MATTI KARPPA, PETTERI KASKI, AND JUKKA KOHONEN<br>Aalto University, Department of Computer Science


#### Abstract

We study the problem of detecting outlier pairs of strongly correlated variables among a collection of $n$ variables with otherwise weak pairwise correlations. After normalization, this task amounts to the geometric task where we are given as input a set of $n$ vectors with unit Euclidean norm and dimension $d$, and for some constants $0<\tau<\rho<1$, we are asked to find all the outlier pairs of vectors whose inner product is at least $\rho$ in absolute value, subject to the promise that all but at most $q$ pairs of vectors have inner product at most $\tau$ in absolute value.

Improving on an algorithm of G. Valiant [FOCS 2012; J. ACM 2015], we present a randomized algorithm that for Boolean inputs ( $\{-1,1\}$-valued data normalized to unit Euclidean length) runs in time $$
\tilde{O}\left(n^{\max \{1-\gamma+M(\Delta \gamma, \gamma), M(1-\gamma, 2 \Delta \gamma)\}}+q d n^{2 \gamma}\right),
$$ where $0<\gamma<1$ is a constant tradeoff parameter and $M(\mu, \nu)$ is the exponent to multiply an $\left\lfloor n^{\mu}\right\rfloor \times\left\lfloor n^{\nu}\right\rfloor$ matrix with an $\left\lfloor n^{\nu}\right\rfloor \times\left\lfloor n^{\mu}\right\rfloor$ matrix and $\Delta=1 /\left(1-\log _{\tau} \rho\right)$. As corollaries we obtain randomized algorithms that run in time $$
\begin{gathered} \tilde{O}\left(n^{\frac{2 \omega}{3-\log _{\tau} \rho}}+q d n^{\frac{2\left(1-\log _{\tau} \rho\right)}{3-\log _{\tau} \rho}}\right) \\ \tilde{O}\left(n^{\frac{4}{2+\alpha\left(1-\log _{\tau} \rho\right)}}+q d n^{\frac{2 \alpha\left(1-\log _{\tau} \rho\right)}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right) \end{gathered}
$$ and in time where $2 \leq \omega<2.38$ is the exponent for square matrix multiplication and $0.3<\alpha \leq 1$ is the exponent for rectangular matrix multiplication. The notation $\tilde{O}(\cdot)$ hides polylogarithmic factors in $n$ and $d$ whose degree may depend on $\rho$ and $\tau$. We present further corollaries for the light bulb problem and for learning sparse Boolean functions.


## 1. Introduction

1.1. Scalable Correlation Analysis. The Pearson product-moment correlation coefficient 20, 48, 53, or briefly, correlation, is one of the most fundamental statistical quantities used to measure the strength of interaction between two sequences of $d$ numerical observations. Assuming we have $n$ such sequences, there are $\Theta\left(n^{2}\right)$ pairs of sequences, which immediately puts forth the algorithmic question of scalability of correlation-based analyses as our number of observables $n$ increases.

If we seek to explicitly compute all $\Theta\left(n^{2}\right)$ pairwise correlations, the size of our desired result forces us to quadratic $\Omega\left(n^{2}\right)$ scaling as a function of $n$. Thus, assuming

[^0]we seek subquadratif ${ }^{1}$ scaling in $n$, we must refine our objective towards aggregate analyses that only implicitly consider all the pairwise correlations. One natural goal in this setting is to identify interesting pairs of observables, subject to the assumption that there are few such pairs. In the context of correlation, one possible signal for an interesting interaction between a pair of observables is an abnormally large, or outlier, (absolute) correlation, as measured against a background parameter that bounds from above the absolute values of correlations between most pairs of observables.

Assuming our $n$ observables are each normalized to zero mean and unit standard deviation, the computational task of finding abnormally large correlations reduces to the following geometric setup:

Problem 1.1 (Outlier Correlations). Given as input a set of $n$ vectors with unit Euclidean norm and dimension $d$, find all the outlier pairs of vectors whose inner product is at least $\rho$ in absolute value, subject to the promise that all but at most $q$ pairs of vectors have inner product at most $\tau$ in absolute value, for some $0<\tau<\rho<1$. (The promise implies in particular that there are at most $q$ outlier pairs. Our interest is on inputs where $q$ is subquadratic in $n$.)

Remark 1.2. A useful variant of the problem is that the input vectors are each assigned one of two possible colors, and we want to find all the outlier pairs with distinct colors, subject to the same promise as above. Let us call this variant Bichromatic Outlier Correlations. In what follows we can work with the bichromatic variant since any instance of Outlier Correlations can be made bichromatic by a routine reduction with a multiplicative polylogarithmic cost in $n L^{2}$ The bichromatic variant in particular captures a batch-query database setting, where the vectors with one color constitute a database, the vectors with the other color constitute a batch of queries to the database, and the outlier pairs with distinct colors capture correlated matches in the database to the queries.

The question of scalability as $n$ increases is now parameterized by the dimension $d$ of the vectors and the thresholds $\rho, \tau$. Furthermore, we can impose restrictions on the structure of the vectors themselves to study restricted variants of the problem. From a discrete computational standpoint, a natural restriction is to assume the input vectors take values in $\{-1,1\}^{3}$ (Indeed, two vectors in $x, y \in\{-1,1\}^{d}$ have inner product $\langle x, y\rangle=\sum_{i=1}^{d} x_{i} y_{i}=\lambda d$ if and only if the corresponding normalized vectors have inner product $\langle x / \sqrt{d}, y / \sqrt{d}\rangle=\lambda$ with $-1 \leq \lambda \leq 1$.) Let us call $\{-1,1\}$-valued inputs Boolean inputs.

[^1]Remark 1.3. For Boolean vectors $x, y \in\{-1,1\}^{d}$, the inner product and the Hamming distance $D_{H}(x, y)=\left|\left\{i \in\{1,2, \ldots, d\}: x_{i} \neq y_{i}\right\}\right|$ are affinely equivalent via the identity

$$
\begin{equation*}
\langle x, y\rangle+2 D_{H}(x, y)=d \tag{1}
\end{equation*}
$$

Thus, for Boolean inputs the Outlier Correlations problem is a variant of near neighbor search under the Hamming metric- the larger the inner product, the smaller the Hamming distance, and vice versa. (Cf. [27, 40, 66].)
1.2. Earlier Lower Bounds and Upper Bounds. Already for Boolean inputs it is known that a quest for subquadratic scaling in $n$ is faced with fine-grained complexity-theoretic barriers. Indeed, unless the Strong Exponential Time Hypothesis [25] is false, any algorithm that for a bichromatic input detects the existence of an outlier pair with $n^{2-\epsilon}$ scaling in $n$ (and simultaneous $2^{o(d)}$ scaling in the dimension $d$ ) must have the property that $\epsilon$ depends on $\rho$ and $\tau$. More precisely, it follows by a simple local transformation (see from a result of Williams 61 that unless SETH is false, $\epsilon$ cannot be bounded from below by a positive constant when $|\rho-\tau| \rightarrow 0$.

On the positive side, algorithms for the $(1+\epsilon)$-approximate nearest neighbor problem can be used to identify outlier pairs $4^{4}$ Locality-sensitive hashing (LSH), introduced by Indyk and Motwani [28, is a well-known framework for finding $(1+$ $\epsilon$ )-approximate nearest neighbors. In their seminal paper, Indyk and Motwani describe an algorithm that, for $\epsilon>0$ and the Hamming metric, uses $\tilde{O}\left(n^{1+1 /(1+\epsilon)}+\right.$ $d n)$ space for supporting approximate nearest-neighbor queries to a data set of $n$ vectors of dimension $d$, and requires $\tilde{O}\left(d n^{1 /(1+\epsilon)}\right)$ query time $5^{5}$ Lower bounds are known for the hash function families 42, 44] which guarantee that algorithms based on data-independent LSH, such as Andoni and Indyk [5] with query time $d n^{1 /(1+\epsilon)^{2}+o(1)}$ for the Euclidean metric, are essentially optimal. Andoni, Indyk, Nguyen and Razenshteyn [7] present a framework that circumvents these lower bounds by introducing data dependence into the hash families, giving the improved query exponent $7 /\left(8(1+\epsilon)^{2}\right)+O\left(1 /(1+\epsilon)^{3}\right)+o(1)$ for the query time in the Euclidean metric. Andoni and Razenshteyn 8 optimize this exponent to $1 /\left(2(1+\epsilon)^{2}-1\right)+o(1)$ and Andoni et al. 6] present LSH optimized for angular distance.

Turning from approximate to exact solutions, Alman and Williams 4] show that the batch Hamming nearest neighbor problem with $n$ simultaneous queries to a database of $n$ vectors can be solved in randomized $n^{2-1 / O(c(n) \log c(n))}$ time if $d=c(n) \log n$ for $c: \mathbb{N} \rightarrow \mathbb{N}$. This scaling is subquadratic in $n$ if the dimension $d$ is bounded by $d=O(\log n)$.
1.3. The Curse of Weak Outliers and Valiant's Breakthrough. For Boolean inputs with unrestricted dimension, the scaling obtained via approximate nearest neighbors is subquadratic in $n$ for all constant $\rho>\tau$. However, such scaling suffers

[^2]from what could be called the curse of weak outliers. Namely, if $\rho$ is small (that is, the outlier correlations themselves are weak), the scaling in $n$ is essentially quadratic $n^{2-c \rho}$ for a positive constant $c$, even if the background $\tau$ decays to zero. Ideally, there should be a way to avoid such direct dependence on the value of $\rho$, as long as $\rho$ and $\tau$ do not converge.

In a breakthrough result, G. Valiant 55] showed that subquadratic scaling is possible, essentially independently of $\rho$, as long as $\log _{\tau} \rho=(\log \rho) /(\log \tau)$ remains bounded from above by a small positive constant. Here the $\operatorname{logarithmic~ratio~} \log _{\tau} \rho$ quantifies the gap between the background and the outliers; since $0<\tau<\rho<1$, both logarithms are negative, and the ratio provides an essentially inverted quantification of the gap so that the gap is wide when the ratio is close to 0 and narrow when the ratio is close to 1 (cf. Figure 1).

Let $2 \leq \omega<2.3728639$ be the exponent ${ }^{6}$ of square matrix multiplication 34. We say that a parameterized event $E_{n}$ happens with high probability (w.h.p.) in the parameter $n$ if $\operatorname{Pr}\left(E_{n}\right) \geq 1-f(n)$ for a function $f(n)$ with $\lim _{n \rightarrow \infty} n^{k} f(n)=0$ for all positive $k$.

Valiant's algorithm runs in two phases. The first one is the approximate detection ${ }^{7}$ phase where the input vectors are divided into blocks and it is decided which pairs of blocks contain one or more vector pairs with absolute inner product in the excess of $\rho d$, if any. The second phase is the listing phase where a brute-force search is performed on the pairs corresponding to the indicated blocks.
Theorem 1.4 (Valiant [55). For all constants $0<\tau<\rho<1$, the Outlier CorRELATIONS problem for Boolean inputs admits a randomized algorithm that runs in time $\tilde{O}\left(n^{\frac{5-\omega}{4-\omega}+\omega \log _{\tau} \rho}\right)$ for approximate detection and subsequent $\tilde{O}\left(q d n^{1 /(4-\omega)}\right)$ time for exact listing of all the outliers, w.h.p.
Remark 1.5. For any constant $\rho>0$, as $\tau \rightarrow 0$ we observe that the exponent for approximate detection in Theorem 1.4 approaches $1.50 \leq(5-\omega) /(4-\omega)<1.62$. Thus, Valiant's algorithm gives subquadratic scaling for all constant $\rho$, as long as $\tau$ is comparatively smaller.
Remark 1.6. A tacit property of Valiant's algorithm is that the running time for approximate detection is, up to polylogarithmic factors, independent of the dimension $d$ of the input. In particular, since the input is Boolean, the (normalized) input vectors are decoherent and hence it suffices to randomly sample the input to access the correlations. By decoherence we mean here the property that if we view the squares of the entries of a (normalized) input vector as a probability distribution over the $d$ dimensions, the probability mass is roughly uniformly distributed across the dimensions. With (normalized) Boolean input, the distribution is exactly uniform and thus eases concentration analyses.

Valiant's result opens up a quest to understand the extent of subquadratic scaling available for Outlier Correlations:

[^3]Assuming that the outliers are well-separated from the background correlations, that is, that $\log _{\tau} \rho$ is small, how close to linear scaling in $n$ can we get?
Our intent in this paper is to present further progress in such a quest.
1.4. Valiant's Algorithm in More Detail. Since our intent is to improve on Valiant's algorithm, it will be convenient to review the key ideas in its design.

Let us assume the input of our instance of Bichromatic Outlier CorreLations is given to us as two $d \times n$ Boolean matrices, $A \in\{-1,1\}^{d \times n}$ and $B \in\{-1,1\}^{d \times n}$. Our task is to find the outlier inner products (of absolute value at least $\rho d$ ) between columns of $A$ and $B$. Valiant's algorithm proceeds in the following four phases; the first and second phases together constitute a compression phase which enables the use of fast matrix multiplication in the third approximate detection phase, which is followed by an exact listing phase.

1. Expansion by uniform random sampling of a tensor power. It is possible to amplify the inner products between columns of $A$ and $B$ by individually tensoring the matrices $A$ and $B$, at the cost of increasing the dimension of the data from $d$ to $d^{p}$ for a positive integer $p$. Indeed, for any two $d$-dimensional real vectors, it holds that $\langle x, y\rangle^{p}=\left\langle x^{\otimes p}, y^{\otimes p}\right\rangle$, where $\otimes p$ indicates $p$-fold Kronecker product (tensoring) of the vector with itself. Accordingly, taking $p$-fold Kronecker products in the vertical dimension only (which we indicate with the notation $\downarrow \otimes p$ ) we obtain the $d^{p} \times n$ matrices $A^{\downarrow \otimes p}$ and $B^{\downarrow \otimes p}$. A key observation is that a uniform random sample of size $s$ (the value of $s$ will be fixed later) of the $d^{p}$ dimensions suffices because the input is decoherent (Boolean) and hence the sum of the sample is strongly concentrated, for example, via the Hoeffding bounds [23]. Accordingly, we assume that the matrices $A^{\downarrow \otimes p}$ and $B^{\downarrow \otimes p}$ in fact have dimensions $s \times n$.
2. Signed aggregation. Since comparatively few of the $n^{2}$ inner products between the columns of $A^{\downarrow \otimes p}$ and $B^{\downarrow \otimes p}$ are outliers, after amplification it is possible to aggregate the $n$ columns in $A^{\downarrow \otimes p}$ (respectively, $B^{\downarrow \otimes p}$ ) by randomly partitioning the columns into $n / t$ blocks of $t$ columns, and taking, with an independent uniform random sign for each column, the sum of the columns in each block. This produces an $s \times(n / t)$ matrix $\tilde{A}$ (respectively, $\tilde{B})$.
3. Approximate detection. Because the outer dimension has now decreased from $n$ to $n / t$, with careful selection of the parameters $s, t, p$ we can now afford to multiply the compressed matrices $\tilde{A}, \tilde{B}$ in subquadratic time to obtain the $(n / t) \times(n / t)$ product $\tilde{A}^{\top} \tilde{B}$. Because of careful amplification by sampling and signed aggregation, the pairs of blocks containing at least one outlier pair are with moderate probability signalled by entries in $\tilde{A}^{\top} \tilde{B}$ that have absolute value above a threshold value $8^{8}$
4. Exact listing. Finally, the outlier pairs can be computed with brute force by computing the $t^{2}$ pairwise inner products of columns of $A$ and $B$ within each signalled pair of blocks.

Theorem 1.4 follows by careful selection of the parameters $s, t, p$ and the signalling threshold.
1.5. Our Contribution. Our contribution in this paper amounts to the observation that the compression phase (expansion followed by signed aggregation) in Valiant's algorithm can be replaced by a faster compression subroutine. In essence, we rely on fast matrix multiplication as the algorithmic device to simultaneously

[^4]expand and aggregate; this requires that we replace the uniform random sampling of dimensions in the expansion phase of Valiant's algorithm with Cartesian product sampling to enable us to entangle expansion and aggregation. Despite the considerable decrease in entropy compared with a uniform random sample, we show that Cartesian product sampling on a tensor power remains roughly as sharply concentrated as a uniform random sample, enabling us to use roughly the same sample size $s$ as Valiant's algorithm for comparable compression, thus resulting in faster execution because of speedup given by fast matrix multiplication. The faster compression subroutine then enables a faster tradeoff that balances between compression and detection, as controlled by a tradeoff parameter $0<\gamma<1$.

Let us now state our main result for Outlier Correlations. For constants $\mu, \nu>0$, let $M(\mu, \nu)$ be the infimum of the values $\sigma>0$ such that there exists an algorithm that multiplies an $\left\lfloor n^{\mu}\right\rfloor \times\left\lfloor n^{\nu}\right\rfloor$ integer matrix with an $\left\lfloor n^{\nu}\right\rfloor \times\left\lfloor n^{\mu}\right\rfloor$ integer matrix in $O\left(n^{\sigma}\right)$ arithmetic operations. For example, $\omega=M(1,1)$.

Theorem 1.7 (Main). For all constants $0<\gamma<1$ and $\Delta \geq 1$, the Outlier CorRELATIONS problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\begin{equation*}
\tilde{O}\left(n^{\max \{1-\gamma+M(\Delta \gamma, \gamma), M(1-\gamma, 2 \Delta \gamma)\}} \tau^{-4}\right) \tag{2}
\end{equation*}
$$

for approximate detection and subsequent $\tilde{O}\left(q d n^{2 \gamma}\right)$ time for exact listing of all the outliers, w.h.p. The running time bounds hold uniformly for all $n^{-\Theta(1)} \leq \tau<\rho<1$ with $\log _{\tau} \rho \leq 1-\Delta^{-1}$.

For specific choices of $\gamma$ and $\Delta$ we obtain the following corollaries.
Corollary 1.8. For all constants $0<\tau<\rho<1$, the Outlier CorrelaTIONS problem for Boolean inputs admits a randomized algorithm that runs in time $\tilde{O}\left(n^{\frac{2 \omega}{3-\log _{\tau} \rho}}\right)$ for approximate detection and subsequent $\tilde{O}\left(q d n^{\frac{2\left(1-\log _{\tau} \rho\right)}{3-\log _{\tau} \rho}}\right)$ time for exact listing of all the outliers, w.h.p.

Let $0.30298<\alpha \leq 1$ be the exponent for rectangular matrix multiplication [33]. That is, $\alpha$ is the supremum of all values $\sigma \leq 1$ with $M(1, \sigma)=2$.

Corollary 1.9. For all constants $0<\tau<\rho<1$, the Outlier CorrelaTIONS problem for Boolean inputs admits a randomized algorithm that runs in time $\tilde{O}\left(n^{\frac{4}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right)$ for approximate detection and subsequent $\tilde{O}\left(q d n^{\frac{2 \alpha\left(1-\log _{\tau} \rho\right)}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right)$ time for exact listing of all the outliers, w.h.p.

Figure 1 displays the subquadratic scaling in $n$ for approximate detection obtained from Corollaries 1.8 and 1.9, contrasted with Theorem 1.4 of [55].

Remark 1.10. Corollary 1.9 implies that asymptotically it is possible to list extremely weak outliers with, say, $\rho=2^{-100}$ and $\tau=2^{-101}$, in time $\tilde{O}\left(n^{1.998}+\right.$ $\left.q d n^{0.003}\right)$.

Remark 1.11. In the setting of well-separated outliers with any constant $\rho$ and $\tau \rightarrow 0$, the exponent for approximate detection in Corollary 1.8 approaches $2 \omega / 3$, improving Valiant's exponent $(5-\omega) /(4-\omega)$ across the range $2 \leq \omega<2.38$.

Although our results show that exponents as low as $2 \omega / 3=4 / 3$ may be feasible for finding outlier correlations (if we assume $\omega=2$ and let $\tau \rightarrow 0$ ), the present


Figure 1. Illustration of subquadratic scaling in $n$ for approximate detection as a function of the parameter $\log _{\tau} \rho$. The matrix multiplication exponents $\omega$ and $\alpha$ are fixed to values $\omega=2.3728639$ and $\alpha=0.30298$ (left), and to values $\omega=2$ and $\alpha=1$ (right).
framework unfortunately appears not to be powerful enough to lower the exponent below $4 / 3$. It remains open whether the exponent can be lowered all the way to 1 (to linear scaling in $n$ ), and whether techniques other than fast matrix multiplication can be used to attain subquadratic scaling without the curse of weak outliers.

## 2. Related Work and Applications

The Outlier Correlations problem has many applications, since the inner product of two vectors can be used to measure the similarity of objects. For instance, [2, 9, 12, 13, 14, 15, 16, 31, 35, 36, 43, 49, 50, 51, 52, 54, 57, 58, 59, 63, 64, 65. Here we will be content with discussing a narrow set of applications of our results and related work. The proofs of all corollaries appear in $\$ 5$
2.1. The Light Bulb Problem. Theorem 1.7 gives as an almost immediate corollary a faster subquadratic algorithm (cf. [55, Corollary 2.2]) for solving the light bulb problem [56], which asks us to discover a hidden correlated pair of light bulbs among $n$ light bulbs blinking on and off independently and uniformly at random. Let us first state the problem in more precise terms and then give our improvement.
Problem 2.1 (Light Bulb). Suppose we are given as input a set of $n$ vectors in $\{-1,1\}^{d}$ consisting of (i) a planted pair with inner product at least $\rho d$ in absolute value for $0<\rho<1$, and (ii) $n-2$ independent uniform random vectors in $\{-1,1\}^{d}$. Our task is to find the planted pair among the $n$ vectors.$^{9}$

Corollary 2.2. For all constants $0<\epsilon<\omega / 3$, the Light Bulb problem admits a randomized algorithm that for all $d \geq 5 \rho^{-\frac{4 \omega}{9 \epsilon}-\frac{2}{3}} \log n$ runs in time

$$
\begin{equation*}
\tilde{O}\left(n^{\frac{2 \omega}{3}+\epsilon} \rho^{-\frac{8 \omega}{9 \epsilon}-\frac{4}{3}}\right) \tag{3}
\end{equation*}
$$

and finds the planted pair of vectors, with probability $1-o(1)$. The running time bound holds uniformly for all $n^{-\Theta(1)}<\rho<1$.

[^5]Prior to Valiant's algorithm, the first subquadratic algorithm for the light bulb problem was the randomized $\tilde{O}\left(n^{1+\left(\log \frac{1+\rho}{2}\right) /\left(\log \frac{1+\sigma}{2}\right)}\right)$ time algorithm of Paturi, Rajasekaran, and Reif [47, where $\sigma$ measures the absolute value of the inner product between the pair of vectors with the second largest inner product. Besides the algorithms based on locality-sensitive hashing within the context of approximate nearest neighbors (see $\$ 1.2$ ) that can also be used to solve light bulb problem, the bucketing codes approach of Dubiner [18] yields a randomized $\tilde{O}\left(n^{2 /(\rho+1)}\right)$ time algorithm for the light bulb problem, which was the fastest algorithm prior to Valiant's algorithm. May and Ozerov [38] present a recursive approach for solving the light bulb problem; with weak outliers also their algorithm converges to quadratic running time.
2.2. Learning Boolean Functions. The light bulb problem generalizes to the task of learning a parity function in the presence of noise.

Problem 2.3 (Parity with Noise). Let the support $S \subseteq\{1,2, \ldots, n\}$ of the parity function have size $k=|S|$ and let the noise rate be $0<\eta<1$. Our task is to find $S$, given access to independent examples of the form $(x, y)$, where (i) the input $x \in\{-1,1\}^{n}$ is chosen uniformly at random, and (ii) the label $y=z \cdot \prod_{j \in S} x_{j}$ is defined by independently choosing $z \in\{-1,1\}$ with $\operatorname{Pr}(z=-1)=\eta$.

The general case of unrestricted $k$ is studied by Blum, Kalai, and Wasserman [11] and Lyubashevsky [37. Studies of the sparse case with $k=O(1)$ include works by Grigorescu, Reyzin, and Vempala 22 ] and Valiant [55]. In the sparse case, we can use a split-and-list transformation presented by Valiant [55] p. 32] together with the algorithm underlying Corollary 2.2 to essentially match an algorithm of Valiant that relies on Fourier-analytic techniques (cf. [55, Theorems 2.4, 5.2, and 5.6]). Compared with Valiant's algorithm, our algorithm has a slightly worse tolerance for noise but better sample complexity:

Corollary 2.4. For all constants $0<\epsilon<\omega / 3$, the Parity with Noise problem admits a randomized algorithm that uses

$$
\begin{equation*}
d \geq(2 k+3) \cdot|1-2 \eta|^{-\frac{4 \omega}{9 \epsilon}-\frac{2}{3}} \log n \tag{4}
\end{equation*}
$$

examples, runs in time

$$
\begin{equation*}
\tilde{O}\left(n^{\frac{\omega+\epsilon}{3} k} \cdot|1-2 \eta|^{-\frac{8 \omega}{9 \epsilon}-\frac{4}{3}}\right), \tag{5}
\end{equation*}
$$

for all sufficiently large $k$, and finds the support of the parity function, with probability at least $1-o(1)$. The running time bound holds uniformly for all $n^{-\Theta(1)}<$ $|1-2 \eta|<1$.

Using algorithms of Feldman, Gopalan, Khot, and Ponnuswami 19, and Mossel, O'Donnell, and Servedio [41, from Corollary 2.4 one can obtain further corollaries in the context of learning sparse juntas and DNFs. We refer to 55 for a detailed exposition.
2.3. Subquadratic Algorithms in Small Dimension. If the dimension $d$ is very small, then subquadratic algorithms in $n$ for Outlier Correlations are available through dedicated space-partitioning data structures such as Voronoi diagrams, enabling near neighbor query times $\tilde{O}\left(d^{c}\right)$ for a constant $c$, but with exponential scaling in size as a function of $d$ [39, 60. Similar query-vs-size tradeoffs are available also in higher dimensions; for example, Kushilevitz, Ostrovsky, and Rabani [32]
presented a data structure with $\tilde{O}(d)$ query time and size $(d n)^{O(1)}$ for approximate nearest neighbors. Alman and Williams [4] obtain subquadratic scaling in $n$ for $d=O(\log n)($ see 1.2$)$.
2.4. Scaling in the Number of Pairs above the Background Correlation. Let us briefly study scaling in the case when the parameter $q$ is relatively large, such as the batch-quer $y^{10}$ case with $q=O(n)$, or the case $q=O\left(n^{2-\delta}\right)$ for a small constant $\delta>0$. In this situation we observe that in Theorem 1.7 the term $\tilde{O}\left(q d n^{2 \gamma}\right)$ can dominate the running time over (22). Indeed, assuming $q d$ is subquadratic in $n$, we can always obtain an overall running time that is subquadratic in $n$ by forcing a small enough $\gamma>0$ in Theorem 1.7. However, such a forced $\gamma$ may cause a suboptimal value of (2) over what would be available if $q$ was smaller. This bottleneck in scaling for large $q$ can be somewhat alleviated by pursuing a two-level recursive strategy, where we run the approximate detection algorithm recursively before proceeding to listing:

Theorem 2.5 (Main, Two-level). For all constants $0<\gamma, \kappa<1$ and $\Delta \geq 1$, the Outlier Correlations problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\tilde{O}\left(n^{\max \{1-\gamma+M(\Delta \gamma, \gamma), M(1-\gamma, 2 \Delta \gamma)\}} \tau^{-4}\right)
$$

for approximate detection and subsequent

$$
\tilde{O}\left(q n^{\gamma \max \{1-\kappa+M(\Delta \kappa, \kappa), M(1-\kappa, 2 \Delta \kappa)\}} \tau^{-4}+q d n^{2 \gamma \kappa}\right)
$$

time for exact listing of all the outliers, w.h.p. The running time bounds hold uniformly for all $n^{-\Theta(1)} \leq \tau<\rho<1$ with $\log _{\tau} \rho \leq 1-\Delta^{-1}$.

For specific choices $\gamma, \kappa, \Delta$ we obtain the analogs of Corollary 1.8 and 1.9 .
Corollary 2.6. For all constants $0<\tau<\rho<1$, the Outlier Correlations problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\tilde{O}\left(n^{\frac{2 \omega}{3-\log _{\tau} \rho}}\right)
$$

for approximate detection and subsequent

$$
\tilde{O}\left(q n^{\frac{2 \omega\left(1-\log _{\tau} \rho\right)}{\left(3-\log _{\tau} \rho\right)^{2}}}+q d n^{\frac{2\left(1-\log _{\tau} \rho\right)^{2}}{\left(3-\log _{\tau} \rho\right)^{2}}}\right)
$$

time for exact listing of all the outliers, w.h.p.

Corollary 2.7. For all constants $0<\tau<\rho<1$, the Outlier Correlations problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\tilde{O}\left(n^{\frac{4}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right)
$$

for approximate detection and subsequent

$$
\tilde{O}\left(q n^{\frac{4 \alpha\left(1-\log _{\tau} \rho\right)}{\left(2+\alpha\left(1-\log _{\tau} \rho\right)\right)^{2}}}+q d n^{\frac{2 \alpha^{2}\left(1-\log _{\tau} \rho\right)^{2}}{\left(2+\alpha\left(1-\log _{\tau} \rho\right)\right)^{2}}}\right)
$$

time for exact listing of all the outliers, w.h.p.

[^6]Remark 2.8. Let us recall from Remark 1.10 that Corollary 1.9 lists outliers with $\rho=2^{-100}$ and $\tau=2^{-101}$ in time $\tilde{O}\left(n^{1.998}+q d n^{0.003}\right)$. Corollary 2.7 improves this to time $\tilde{O}\left(n^{1.998}+q n^{0.003}+q d n^{0.0000045}\right)$.

Remark 2.9. In the setting of well-separated outliers with any constant $\rho$ and $\tau \rightarrow 0$, Corollary 1.8 lists outliers in time $\tilde{O}\left(n^{2 \omega / 3+\epsilon_{\tau}}+q d n^{2 / 3+\epsilon_{\tau}}\right)$. Corollary 2.6 improves this to time $O\left(n^{2 \omega / 3+\epsilon_{\tau}}+q n^{2 \omega / 9+\epsilon_{\tau}}+q d n^{2 / 9+\epsilon_{\tau}}\right)$. Here $\epsilon_{\tau} \rightarrow 0$ as $\tau \rightarrow 0$.
2.5. Discussion and Further Work. In addition to Valiant's algorithm, another breakthrough in the context of finding and approximating the outlier entries in a matrix product $A^{\top} B$ of two given $t \times s$ matrices $A, B$ was made by Pagh 45. Pagh's algorithm also computes a compressed version of the product $A^{\top} B$ by using 2 -wise independent hash families that are algebraically compatible with a cyclic group (of order $b$ ) to aggregate the operands $A, B$ with random signs into preimages of the hash function along the outer dimension $(s)$ and with explicit (uncompressed) summation along the inner dimension $(t)$. Because of the compatibility with the cyclic group, the compressed operands can be multiplied in essentially linear time via cyclic convolution (using the FFT for the cyclic group). The compressed product gives unbiased estimates of the product entries $\left(A^{\top} B\right)_{j_{1} j_{2}}$ with 2-wise independence controlling the variance via the Frobenius norm at $\left\|A^{\top} B\right\|_{F}^{2} / b$. Despite the essentially linear-time performance given by the use of FFT in evaluating the compressed products (thus enabling a large value of $b$ ), the combination of Frobenius control on variance and the linear scaling along the inner dimension $t$ appears not to yield improvements for the light bulb problem or for Outlier Correlations.

To our knowledge, there are no lower bounds that would preclude the use of multilinear algorithms other than fast matrix multiplication (such as Pagh's 45] algorithm discussed above) for Outlier Correlations, yet fast matrix multiplication remains the only known tool to obtain truly subquadratic scaling for weak outliers. As Valiant [55, §6] highlights, it would be of considerable interest to avoid fast matrix multiplication altogether to obtain practical subquadratic scaling for moderately-sized $n$. In fact, from this perspective our present results arguably proceed in the wrong direction by making heavier use of fast matrix multiplication to obtain asymptotically faster subquadratic scaling. For the state of the art in practical fast matrix multiplication algorithms, cf. [24] and [10].

Recently, Ahle, Pagh, Razenshteyn, and Silvestri [3] show that inner product similarity joins are hard to approximate in subquadratic time in $n$ unless the Orthogonal Vectors Conjecture (see §6) and subsequently SETH are false. Stated in terms of Outlier Correlations, they show that Boolean inputs with $\log _{\tau} \rho=1-o(1 / \sqrt{\log n})$ do not admit subquadratic scaling in $n$ unless OVC is false which would imply that SETH is false. Narrowing the gap between subquadratic approximability and inapproximability remains a topic for further work.

Let us conclude with a few questions for further work. Is $\log _{\tau} \rho$ the natural parameter for subquadratic scaling in the context of Outlier Correlations? Is it possible to improve the limiting exponent $2 \omega / 3$ in Corollary 1.8 and in Corollary 2.2. Curiously, the algorithm in our Corollary 2.4 and Valiant's algorithm [55, Algorithm 11] for Parity with Noise rely on somewhat different techniques, yet both arrive at the constant $\omega / 3$ for learning sparse parities-is this just a coincidence? Finally, is it possible to derandomize our algorithms or improve their
space usage? In subsequent work [30, a superset of present authors derandomize Valiant's algorithm, but restricted to ( $\pm 1$ )-valued inputs only and with more modest subquadratic running times. We refer to [46] and [29] for further recent work and motivation for derandomization and resource tradeoffs in the context of similarity search.

## 3. Preliminaries

This section collects terminology, notation, and background results used in the subsequent development.
3.1. The Hoeffding Bound. The Hoeffding bound establishes sharp concentration around the expectation of a sum of independent terms assuming we have control on the support of the terms.

Theorem 3.1 (Hoeffding [23, Theorem 2]). Let $Z_{1}, Z_{2}, \ldots, Z_{s}$ be independent random variables with $\ell_{i} \leq Z_{i} \leq u_{i}$ for all $i=1,2, \ldots, s$ and let $Z=Z_{1}+Z_{2}+\ldots+Z_{s}$. Then, for all $d>0$ it holds that

$$
\begin{equation*}
\operatorname{Pr}(Z-\mathrm{E}[Z] \geq d) \leq \exp \left(-\frac{2 d^{2}}{\sum_{i=1}^{s}\left(u_{i}-\ell_{i}\right)^{2}}\right) \tag{6}
\end{equation*}
$$

3.2. Inner Products, Restriction, Powering. Let us set up some notation regarding inner products and tensor powers. For convenience, let us write $[d]=$ $\{1,2, \ldots, d\}$. Let $x=\left(x_{1}, x_{2}, \ldots, x_{d}\right)$ and $y=\left(y_{1}, y_{2}, \ldots, y_{d}\right)$ be two vectors of indeterminates. The inner product

$$
\begin{equation*}
\langle x, y\rangle=x_{1} y_{1}+x_{2} y_{2}+\ldots+x_{d} y_{d} \tag{7}
\end{equation*}
$$

is a multilinear polynomial in the indeterminates $x, y$. For a (multi)set of indices $I \subseteq[d]$, we use the notation

$$
\langle x, y\rangle_{I}=\sum_{i \in I} x_{i} y_{i}
$$

for an inner product $\langle x, y\rangle$ taken over the indices in $I$.
For a positive integer $p$, we observe the $p$ th power of the inner product $\langle x, y\rangle$ is a polynomial in $x, y$ whose monomials can be indexed by $p$-tuples $\vec{i}=\left(i_{1}, i_{2}, \ldots, i_{p}\right) \in$ $[d]^{p}$ with $i_{1}, i_{2}, \ldots, i_{p} \in[d]$. Indeed, expanding the product of sums into a sum of products, we observe that

$$
\begin{equation*}
\langle x, y\rangle^{p}=\left(x_{1} y_{1}+x_{2} y_{2}+\ldots+x_{d} y_{d}\right)^{p}=\sum_{\vec{i} \in[d]^{p}} \prod_{\ell=1}^{p} x_{i_{\ell}} \prod_{\ell=1}^{p} y_{i_{\ell}} \tag{8}
\end{equation*}
$$

Let us write $x^{\otimes p}$ and $y^{\otimes p}$ for the $d^{p}$-dimensional vectors obtained by taking the $p$ th tensor power of $x$ and $y$ and whose coordinates are indexed by the $d^{p}$ possible $p$-tuples:

$$
x^{\otimes p}=\left(\prod_{\ell=1}^{p} x_{i_{\ell}}: \vec{i} \in[d]^{p}\right) \text { and } y^{\otimes p}=\left(\prod_{\ell=1}^{p} y_{i_{\ell}}: \vec{i} \in[d]^{p}\right)
$$

With this notation, we observe from (7) and (8) that

$$
\begin{equation*}
\langle x, y\rangle^{p}=\left\langle x^{\otimes p}, y^{\otimes p}\right\rangle \tag{9}
\end{equation*}
$$

3.3. Cartesian Sampling Lemma. The following lemma will be convenient when analysing the concentration of samples obtained in the compression algorithm. For an analogous analysis, cf. [17].

Lemma 3.2 (Cartesian Sampling). Let $s$ and $m$ be positive integer squares and let $x_{1}, x_{2}, \ldots, x_{m^{1 / 2}} \in\{-1,1\}$ and $y_{1}, y_{2}, \ldots, y_{m^{1 / 2}} \in\{-1,1\}$ with

$$
\sum_{u=1}^{m^{1 / 2}} x_{u}=\xi m^{1 / 2} \quad \text { and } \quad \sum_{v=1}^{m^{1 / 2}} y_{v}=\eta m^{1 / 2}
$$

for some $-1 \leq \xi, \eta \leq 1$. Suppose $S_{1}, S_{2}$ are two multisets of size $s^{1 / 2}$ selected by drawing two $s^{1 / 2}$-tuples consisting of elements of $\left[m^{1 / 2}\right.$ ] independently and uniformly at random (with repetition). Then, with high probability in $s$, we have

$$
\begin{equation*}
\left|\sum_{(u, v) \in S_{1} \times S_{2}} x_{u} y_{v}\right| \leq|\xi \eta| s+(|\xi|+|\eta|) s^{3 / 4} \log s+s^{1 / 2}(\log s)^{2} \tag{10}
\end{equation*}
$$

Furthermore, if

$$
\begin{equation*}
|\xi|,|\eta| \geq s^{-1 / 4} \log s \tag{11}
\end{equation*}
$$

then, with high probability in s, we have

$$
\begin{equation*}
\left|\sum_{(u, v) \in S_{1} \times S_{2}} x_{u} y_{v}\right| \geq|\xi \eta| s-(|\xi|+|\eta|) s^{3 / 4} \log s+s^{1 / 2}(\log s)^{2} \tag{12}
\end{equation*}
$$

Proof. For $j=1,2, \ldots, s^{1 / 2}$ let $X_{j} \in\{-1,1\}$ be a random variable that independently assumes each of the $m^{1 / 2}$ values $x_{u}$ with probability $m^{-1 / 2}$. For $j=1,2, \ldots, s^{1 / 2}$ let $Y_{j} \in\{-1,1\}$ be a random variable that independently assumes each of the $m^{1 / 2}$ values $y_{v}$ with probability $m^{-1 / 2}$. Consider the random variables $X=X_{1}+X_{2}+\ldots+X_{s^{1 / 2}}$ and $Y=Y_{1}+Y_{2}+\ldots+Y_{s^{1 / 2}}$. By linearity of expectation, we have $\mathrm{E}[X]=\xi s^{1 / 2}$ and $\mathrm{E}[Y]=\eta s^{1 / 2}$. Furthermore, from the Hoeffding bound (6) it follows that $|X-\mathrm{E}[X]| \leq s^{1 / 4} \log s$ and $|Y-\mathrm{E}[Y]| \leq s^{1 / 4} \log s$ with high probability in $s$. Thus, we conclude that with high probability in $s$ it holds that

$$
\begin{aligned}
|X Y| & \leq\left(|\xi| s^{1 / 2}+s^{1 / 4} \log s\right)\left(|\eta| s^{1 / 2}+s^{1 / 4} \log s\right) \\
& =|\xi \eta| s+s^{3 / 4}(|\xi|+|\eta|) \log s+s^{1 / 2}(\log s)^{2}
\end{aligned}
$$

Similarly, assuming that $|\xi|,|\eta| \geq s^{-1 / 4} \log s$, with high probability in $s$ we have the lower bound

$$
\begin{aligned}
|X Y| & \geq\left(|\xi| s^{1 / 2}-s^{1 / 4} \log s\right)\left(|\eta| s^{1 / 2}-s^{1 / 4} \log s\right) \\
& =|\xi \eta| s-s^{3 / 4}(|\xi|+|\eta|) \log s+s^{1 / 2}(\log s)^{2}
\end{aligned}
$$

3.4. Anti-concentration of Signed Aggregation. We recall the following anticoncentration lemma from the analysis of Valiant's algorithm.

Lemma 3.3 (Valiant [55, Lemma 3.2]). Let $C$ be a $t \times t$ matrix with entries $c_{i j}$. Suppose that $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{t} \in\{-1,1\}$ and $\beta_{1}, \beta_{2}, \ldots, \beta_{t} \in\{-1,1\}$ have been selected
independently and uniformly at random. Then,

$$
\operatorname{Pr}\left(\left|\sum_{i, j} \alpha_{i} \beta_{j} c_{i j}\right| \geq \frac{1}{4} \max _{i, j}\left|c_{i j}\right|\right) \geq \frac{1}{4} .
$$

## 4. The Algorithm

This section proves Theorem 1.7. We present a self-contained analysis and description of the entire algorithm underlying Theorem 1.7 even if our main technical contribution occurs in the compression subroutine that simultaneously expands and aggregates (cf. \$1.4 and \$1.5).
4.1. The Input and the Parameters. Let $0<\gamma<1$ and $\Delta \geq 1$ be fixed constants. Suppose we are given as input two $d \times n$ matrices $A, B \in\{-1,1\}^{d \times n}$ together with parameters $\rho, \tau$ that satisfy $n^{-\Theta(1)} \leq \tau<\rho<1$ and $\log _{\tau} \rho \leq 1-\Delta^{-1}$. For $i \in[d]$ and $j \in[n]$, let us write $a_{i j}$ and $b_{i j}$ for the entries of $A$ and $B$ at row $i$, column $j$. Furthermore, let us write $a_{j}$ and $b_{j}$ for the $j$ th column of $A$ and $B$, respectively.

The algorithm works with three positive integer parameters $s, t, p$ whose precise values we will fix in what follows. At this point, we can assume that $p$ is even, and that $s$ is a positive integer square. We will furthermore pad our input matrices with at most $t-1<n$ all-zero columns to ensure divisibility by $t$. Let us denote by $\bar{n}=\lceil n / t\rceil \cdot t$ the number of columns including padding.
4.2. The Algorithm. We describe the algorithm first and then proceed to analyse its correctness and running time. The algorithm executes a five-phase iteration $\left\lceil(\log n)^{2}\right\rceil$ times, and gives as its output the union of the outputs of all the iterations. The five phases are as follows:
0. Setup. Draw uniformly at random a partition $J_{1}, J_{2}, \ldots, J_{\bar{n} / t}$ of $[\bar{n}]$ into sets $J_{k}$ with $\left|J_{k}\right|=t$ for all $k \in[\bar{n} / t]$. For $k \in[\bar{n} / t]$, let us write $\underline{J}_{k}=J_{k} \cap[n]$ for the restriction of $J_{k}$ to the indices corresponding to non-zero columns. We observe that $1 \leq\left|\underline{J}_{k}\right| \leq t$. Draw independently and uniformly at random (with repetition) two $s^{1 / 2}$-tuples consisting of elements of $[d]^{p / 2}$ and let $I_{1}$ and $I_{2}$ be the resulting multisets of size $s^{1 / 2}$. Form the Cartesian product $I=I_{1} \times I_{2}$ of size $s$ and observe that all the elements of $I$ are $p$-tuples in $[d]^{p}$. Draw $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{\bar{n}} \in\{-1,1\}$ and $\beta_{1}, \beta_{2}, \ldots, \beta_{\bar{n}} \in\{-1,1\}$ independently and uniformly at random.

The compressed matrices (definition). Let us first define what we compute and only then give the algorithm that computes the matrices. For $\vec{i} \in I$ and $k \in[\bar{n} / t]$, define

$$
\begin{equation*}
\hat{a}_{\vec{i} k}=\sum_{j \in \underline{J}_{k}} \alpha_{j}\left(a_{j}^{\otimes p}\right)_{\vec{i}} \quad \text { and } \quad \hat{b}_{\overrightarrow{i k}}=\sum_{j \in \underline{J}_{k}} \beta_{j}\left(b_{j}^{\otimes p}\right)_{\vec{i}} . \tag{13}
\end{equation*}
$$

This defines the matrices $\hat{A}_{I}=\left(\hat{a}_{\vec{k} k}\right)$ and $\hat{B}_{I}=\left(\hat{a}_{\vec{i} k}\right)$, both of size $s \times(\bar{n} / t){ }^{11}$
$1+2$. Compression by simultaneous expansion and aggregation. We compute the compressed matrices $\hat{A}_{I}$ and $\hat{B}_{I}$ from the given input $A$ and $B$ as follows. We describe the computation only for $\hat{A}_{I}$, the computation for $\hat{B}_{I}$ is symmetric. Let us first establish that the task of computing $\hat{A}_{I}$ amounts to $\bar{n} / t$ matrix products

[^7]because our random sample $I$ decomposes into the Cartesian product $I=I_{1} \times I_{2}$. Indeed, we observe from (13) that if we write $\vec{i}=\left(\vec{i}_{1}, \vec{i}_{2}\right) \in I$ in terms of its parts $\vec{i}_{1} \in I_{1}$ and $\vec{i}_{2} \in I_{2}$, we have
\[

$$
\begin{equation*}
\hat{a}_{\vec{i} k}=\hat{a}_{\left(\vec{i}_{1}, \vec{i}_{2}\right), k}=\sum_{j \in \underline{J}_{k}} \alpha_{j}\left(a_{j}^{\otimes p / 2}\right)_{\vec{i}_{1}}\left(a_{j}^{\otimes p / 2}\right)_{\vec{i}_{2}} . \tag{14}
\end{equation*}
$$

\]

Thus, for each fixed $k \in[\bar{n} / t]$ we observe that 14 is in fact a matrix product. Indeed, define the $s^{1 / 2} \times t$ matrix $L_{k}=\left(\ell_{\vec{i}_{1} j}\right)$ to consist of the entries

$$
\begin{equation*}
\ell_{\vec{i}_{1} j}=\alpha_{j}\left(a_{j}^{\otimes p / 2}\right)_{\vec{i}_{1}} \tag{15}
\end{equation*}
$$

for $\vec{i}_{1} \in I_{1}$ and $j \in J_{k}$. Similarly, define the $s^{1 / 2} \times t$ matrix $R_{k}=\left(r_{\vec{i}_{2} j}\right)$ to consist of the entries

$$
\begin{equation*}
r_{\vec{i}_{2} j}=\left(a_{j}^{\otimes p / 2}\right)_{\vec{i}_{2}} \tag{16}
\end{equation*}
$$

for $\vec{i}_{2} \in I_{2}$ and $j \in J_{k}$. By (13) the $\left(\vec{i}_{1}, \vec{i}_{2}\right)$-entry of the product matrix $L_{k} R_{k}^{\top}$ is precisely the value $\hat{a}_{\left(\vec{i}_{1}, \vec{i}_{2}\right), k}$. Indeed,

$$
\left(L_{k} R_{k}^{\top}\right)_{\vec{i}_{1}, \vec{i}_{2}}=\sum_{j \in J_{k}} \ell_{\vec{i}_{1} j} r_{\vec{i}_{2} j}=\sum_{j \in \underline{J}_{k}} \ell_{\vec{i}_{1} j} r_{\vec{i}_{2} j}=\sum_{j \in \underline{J}_{k}} \alpha_{j}\left(a_{j}^{\otimes p / 2}\right)_{\vec{i}_{1}}\left(a_{j}^{\otimes p / 2}\right)_{\vec{i}_{2}}=\hat{a}_{\left(\vec{i}_{1}, i_{2}\right), k}
$$

To compute the matrix $\hat{A}_{I}$, we thus proceed as follows. For each $k \in[\bar{n} / t]$, we construct the matrices $L_{k}$ and $R_{k}$ entrywise using (15) and (16), respectively, and then we compute the $s^{1 / 2} \times s^{1 / 2}$ product $L_{k} R_{k}^{\top}$. This gives us the $(\bar{n} / t) \times s$ matrix $\hat{A}_{I}$, at the cost of $\bar{n} / t$ matrix multiplications with inner dimension $t$ and outer dimension $s^{1 / 2}$.
3. Approximate detection. We compute the matrix product $\hat{A}_{I}^{\top} \hat{B}_{I}$ and mark all entries $\left(k_{1}, k_{2}\right) \in[\bar{n} / t] \times[\bar{n} / t]$ that satisfy $\left|\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}}\right| \geq \frac{1}{8} \rho^{p} s$. (This requires one matrix product with inner dimension $s$ and outer dimension $\bar{n} / t$.)
4. Exact listing. Finally, for each marked entry $\left(k_{1}, k_{2}\right) \in[\bar{n} / t] \times[\bar{n} / t]$, we iterate over all $\left(j_{1}, j_{2}\right) \in \underline{J}_{k_{1}} \times \underline{J}_{k_{2}}$, compute the inner product $\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle$, and output $\left(j_{1}, j_{2}\right)$ as an outlier pair if $\left|\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle\right| \geq \rho d$.
4.3. Analysis (correctness). Proof outline. The correctness proof proceeds as follows. We start by establishing an upper bound on the absolute value of entries in the product matrix $\hat{A}_{I}^{\top} \hat{B}_{I}$, assuming only small partial inner products are aggregated in an entry. Dually, we then establish a lower bound for an entry of $\hat{A}_{I}^{\top} \hat{B}_{I}$ that contains at least one large partial inner product. We then establish control on the concentration of the partial inner products using the Cartesian concentration lemma (Lemma 3.2). With appropriate parameterization, we then verify that any single outlier is found (its entry in the product matrix is marked for listing) with good probability during one iteration, and amplify this probability by iteration to ensure that all outliers are found with high probability. Finally, we show that at most $q$ entries are marked for listing with high probability, completing the correctness proof.

Let us now proceed with the detailed proof. We start by fixing the value of the parameter $t$. With foresight, let $t$ be the unique integer that satisfies

$$
\begin{equation*}
n^{\gamma} \leq t<n^{\gamma}+1 \tag{17}
\end{equation*}
$$

Observe that the choice implies $n \leq \bar{n} \leq n+n^{\gamma}$.

We begin our analysis by studying the entries of the product matrix $\hat{A}_{I}^{\top} \hat{B}_{I}$. Let $k_{1}, k_{2} \in[\bar{n} / t]$ and consider the $\left(k_{1}, k_{2}\right)$-entry of $\hat{A}_{I}^{\top} \hat{B}_{I}$. From 13 we have

$$
\begin{align*}
\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}} & =\sum_{\vec{i} \in I}\left(\sum_{j_{1} \in \underline{J}_{k_{1}}} \alpha_{j_{1}}\left(a_{j_{1}}^{\otimes p}\right)_{\vec{i}}\right)\left(\sum_{j_{2} \in \underline{J}_{k_{2}}} \beta_{j_{2}}\left(b_{j_{2}}^{\otimes p}\right)_{\vec{i}}\right)  \tag{18}\\
& =\sum_{j_{1} \in \underline{J}_{k_{1}}} \alpha_{j_{1}} \sum_{j_{2} \in \underline{J}_{k_{2}}} \beta_{j_{2}}\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I} .
\end{align*}
$$

The value of an entry of $\hat{A}_{I}^{\top} \hat{B}_{I}$ that aggregates only small partial inner products. Let us first derive an upper bound for $\left|\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}}\right|$ subject to an upper bound for the absolute values of the partial inner products $\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}\right|$. That is, let us assume $U \geq 0$ is an upper bound $\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}\right| \leq U$ that holds for all $j_{1} \in J_{k_{1}}$ and $j_{2} \in J_{k_{2}}$ with high probability in $n$. Recall that $\left|J_{k_{1}}\right|=\left|J_{k_{2}}\right|=t$ and that the signs $\alpha_{j_{1}}, \beta_{j_{2}} \in\{-1,1\}$ have been selected independently and uniformly at random for $j_{1} \in J_{k_{1}}$ and $j_{2} \in J_{k_{2}}$. Let us fix $j_{1} \in J_{k_{1}}$ arbitrarily and analyse the concentration of the innermost sum $Z_{j_{1}}=\sum_{j_{2} \in J_{k_{2}}} \beta_{j_{2}}\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}$ in 18). Because $\beta_{j_{2}}$ and $\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}$ are independent, each summand $Z_{j_{1}, j_{2}}=\beta_{j_{2}}\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}$ is a random variable with zero expectation. Let us condition on our assumption $\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}\right| \leq U$ and thus conclude that $\left|Z_{j_{1}, j_{2}}\right| \leq U$ holds for each summand. Thus, since $t \leq n$, the Hoeffding bound (6) gives us that $\left|Z_{j_{1}}\right| \leq t^{1 / 2} U \log n$ with high probability in $n$. By the union bound we thus have that $\left|Z_{j_{1}}\right| \leq t^{1 / 2} U \log n$ holds for all $j_{1} \in J_{k_{1}}$ with high probability in $n$. Let us condition on this event. Observe that 18 equals $Z=\sum_{j_{1} \in J_{k_{1}}} \alpha_{j_{1}} Z_{j_{1}}$. Since $\alpha_{j_{1}}$ and $Z_{j_{1}}$ are independent, the product $\alpha_{j_{1}} Z_{j_{1}}$ has zero expectation. Thus, from $\left|\alpha_{j_{1}} Z_{j_{1}}\right| \leq t^{1 / 2} U \log n$ and the Hoeffding bound (6), we conclude that, with high probability in $n$, we have

$$
\begin{equation*}
\left|\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}}\right| \leq t U(\log n)^{2} \tag{19}
\end{equation*}
$$

The value of an entry of $\hat{A}_{I}^{\top} \hat{B}_{I}$ that aggregates at least one large partial inner product. Next, let us suppose that $L \geq 0$ is a lower bound $\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}\right| \geq L$ for some specific $j_{1}, j_{2} \in[n]$. Then, for the pair of blocks $\left(k_{1}, k_{2}\right) \in[\bar{n} / t] \times[\bar{n} / t]$ with $j_{1} \in \underline{J}_{k_{1}}$ and $j_{2} \in \underline{J}_{k_{2}}$ we conclude from 18 and Lemma 3.3 that, with probability at least $1 / 4$, we have

$$
\begin{equation*}
\left|\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}}\right| \geq \frac{L}{4} . \tag{20}
\end{equation*}
$$

When (20) holds, we say the pair $\left(j_{1}, j_{2}\right)$ is signalled by (the $\left(k_{1}, k_{2}\right)$-entry in) $\hat{A}_{I}^{\top} \hat{B}_{I}$.
Concentration of the partial inner products. Let $j_{1}, j_{2} \in[n]$. Let us now study the concentration of $\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}$ using the Cartesian concentration lemma (Lemma 3.2). Because $I=I_{1} \times I_{2}$, the partial inner product $\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I}$ decomposes into

$$
\begin{align*}
\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I_{1} \times I_{2}} & =\sum_{\vec{i} \in I_{1} \times I_{2}}\left(a_{j_{1}}^{\otimes p}\right)_{\vec{i}}\left(b_{j_{2}}^{\otimes p}\right)_{\vec{i}}  \tag{21}\\
& =\sum_{\left(\vec{i}_{1}, \vec{i}_{2}\right) \in I_{1} \times I_{2}}\left(a_{j_{1}}^{\otimes p / 2}\right)_{\vec{i}_{1}}\left(a_{j_{1}}^{\otimes p / 2}\right)_{\vec{i}_{2}}\left(b_{j_{2}}^{\otimes p / 2}\right)_{\vec{i}_{1}}\left(b_{j_{2}}^{\otimes p / 2}\right)_{\vec{i}_{2}}
\end{align*}
$$

Now suppose that $\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle=d \sigma$ for some $-1 \leq \sigma \leq 1$. From (9) we thus have $\left\langle a_{j_{1}}^{\otimes p / 2}, b_{j_{2}}^{\otimes p / 2}\right\rangle=d^{p / 2} \sigma^{p / 2}$. Recalling that $\left|I_{1}\right|=\left|I_{2}\right|=s^{1 / 2}$, take $\xi=\eta=\sigma^{p / 2}$,
$m=d^{p}, S_{1}=I_{1}, S_{2}=I_{2}, x_{\vec{u}}=\left(a_{j_{1}}^{\otimes p / 2}\right)_{\vec{u}}\left(b_{j_{2}}^{\otimes p / 2}\right)_{\vec{u}}$, and $y_{\vec{v}}=\left(a_{j_{1}}^{\otimes p / 2}\right)_{\vec{v}}\left(b_{j_{2}}^{\otimes p / 2}\right)_{\vec{v}}$ for all $\vec{u}, \vec{v} \in[d]^{p / 2}$ in Lemma 3.2 and thus observe from 21 and 10 that, with high probability in $s$, we have the upper bound

$$
\begin{align*}
\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I_{1} \times I_{2}}\right| & =\left|\sum_{\left(\vec{i}_{1}, \vec{i}_{2}\right) \in I_{1} \times I_{2}} x_{\vec{i}_{1}} y_{\vec{i}_{2}}\right|  \tag{22}\\
& \leq\left|\sigma^{p}\right| s+2\left|\sigma^{p / 2}\right| s^{3 / 4} \log s+s^{1 / 2}(\log s)^{2} .
\end{align*}
$$

Conversely, assuming that

$$
\begin{equation*}
\left|\sigma^{p}\right| \geq s^{-1 / 2}(\log s)^{2} \tag{23}
\end{equation*}
$$

so that $\sqrt[11]{ }$ holds, from 21 and $\sqrt{12}$ we have, with high probability in $s$, the lower bound

$$
\begin{align*}
\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I_{1} \times I_{2}}\right| & \geq\left|\sigma^{p}\right| s-2\left|\sigma^{p / 2}\right| s^{3 / 4} \log s+s^{1 / 2}(\log s)^{2}  \tag{24}\\
& \geq\left|\sigma^{p}\right| s-2\left|\sigma^{p / 2}\right| s^{3 / 4} \log s .
\end{align*}
$$

Our eventual choice of $s$ will grow at least as fast as a root function ${ }^{12}$ of $n$, so by the union bound we can assume that the upper bound $\sqrt[22]{ }$ ) and the lower bound (24) hold for all relevant $j_{1}, j_{2} \in[n]$ with high probability in $n$.

An upper bound for aggregated background correlations. Let us now fix $s$ to be the least integer square that satisfies ${ }^{13}$

$$
\begin{equation*}
\tau^{-2 p} \leq s \leq 2 \tau^{-2 p} \tag{25}
\end{equation*}
$$

Thus, from 22 with $|\sigma| \leq \tau$ we have

$$
\begin{equation*}
\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I_{1} \times I_{2}}\right| \leq 4 \tau^{p} s(\log s)^{2} \leq 4 \tau^{p} s(\log n)^{3} \tag{26}
\end{equation*}
$$

here we assume that our eventual choice for $s$ grows no faster than a polynomial function of $n$, that is $\log s=O(\log n)$ and hence $\log s \leq(\log n)^{3 / 2}$ for all large enough $n$. Let us now choose a value for the upper bound $U$ in 19. By 26) we can select $U=8 \tau^{p} s(\log n)^{3}$ to conclude that, with high probability $n$, for all $k_{1}, k_{2} \in[\bar{n} / t]$ we have

$$
\begin{equation*}
\left|\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}}\right| \leq 8 \tau^{p} \operatorname{st}(\log n)^{5} \tag{27}
\end{equation*}
$$

unless there is at least one pair $\left(j_{1}, j_{2}\right) \in \underline{J}_{k_{1}} \times \underline{J}_{k_{2}}$ with $\left|\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle\right|>\tau d$.
Any fixed outlier correlation will be listed with probability at least $1 / 4$ during any fixed iteration of the algorithm. For any outlier correlation we have $|\sigma| \geq \rho$. Let us assume that our eventual choice of $p$ will be such that for all large enough $p$ we have both

$$
\begin{equation*}
\rho^{p} \geq \tau^{p}(\log s)^{2} \geq s^{-1 / 2}(\log s)^{2} \tag{28}
\end{equation*}
$$

and

$$
\begin{equation*}
2\left|\sigma^{p / 2}\right| s^{3 / 4} \log s \leq 2\left|\sigma^{p / 2}\right| \tau^{p / 2} s \log s \leq \frac{1}{2}\left|\sigma^{p}\right| s \tag{29}
\end{equation*}
$$

[^8]Subject to these assumptions, we observe that 23 holds by 28) and from (24) and 29 we have

$$
\begin{equation*}
\left|\left\langle a_{j_{1}}^{\otimes p}, b_{j_{2}}^{\otimes p}\right\rangle_{I_{1} \times I_{2}}\right| \geq \frac{1}{2} \rho^{p} s \tag{30}
\end{equation*}
$$

Let us now choose a value for the lower bound $L$ in 20). By (30) we can select $L=\frac{1}{2} \rho^{p} s$ to conclude that any fixed $\left(j_{1}, j_{2}\right) \in[n] \times[n]$ with $\left|\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle\right| \geq \rho d$ is signalled by $\hat{A}_{I}^{\top} \hat{B}_{I}$ with probability at least $1 / 4$. Recall now that our threshold in the algorithm for marking an entry $\left(k_{1}, k_{2}\right) \in[\bar{n} / t] \times[\bar{n} / t]$ is $\left|\left(\hat{A}_{I}^{\top} \hat{B}_{I}\right)_{k_{1} k_{2}}\right| \geq$ $\frac{1}{8} \rho^{p} s=L / 4$. Thus, assuming $\left(j_{1}, j_{2}\right) \in[n] \times[n]$ is signalled by $\hat{A}_{I}^{\top} \hat{B}_{I}$, the pair of blocks $\left(k_{1}, k_{2}\right) \in[\bar{n} / t] \times[\bar{n} / t]$ with $\left(j_{1}, j_{2}\right) \in \underline{J}_{k_{1}} \times \underline{J}_{k_{2}}$ will be marked by the algorithm in the detection phase, and hence $\left(j_{1}, j_{2}\right)$ will be listed as an outlier in the listing phase.

With high probability in n, all outlier correlations will be listed. Since the algorithm has at least $(\log n)^{2}$ independent iterations, and each iteration outputs any fixed outlier correlation with probability at least $1 / 4$, it follows that this fixed outlier correlation is output with high probability in $n$. Taking the union bound over the at most $q \leq n^{2}$ outlier correlations, every outlier correlation is listed with high probability in $n$. Observe that this is unaffected by the padding of input data.

With high probability in n, at most $q$ pairs of blocks will be marked. It remains to complete the parameterization of the algorithm so that at most $q$ pairs of blocks $\left(k_{1}, k_{2}\right) \in[\bar{n} / t] \times[\bar{n} / t]$ will be marked with high probability in $n$. Recall that there are at most $q$ pairs $\left(j_{1}, j_{2}\right) \in[n] \times[n]$ with $\left|\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle\right|>\tau d$. Thus, recalling (27) and our marking threshold $\frac{1}{8} \rho^{p} s$, it suffices to set up the parameters so that $8 \tau^{p} s t(\log n)^{5}<\frac{1}{8} \rho^{p} s$. For example, we can select the unique even integer $p$ with

$$
\begin{equation*}
\frac{\log t+5 \log \log n+\log 128}{\log (\rho / \tau)} \leq p<\frac{\log t+5 \log \log n+\log 128}{\log (\rho / \tau)}+2 \tag{31}
\end{equation*}
$$

From (25) and (31) we thus have

$$
\begin{equation*}
t^{\frac{2 \log (1 / \tau)}{\log (\rho / \tau)}}(128 \log n)^{\frac{10 \log (1 / \tau)}{\log (\rho / \tau)}} \leq s<2 t^{\frac{2 \log (1 / \tau)}{\log (\rho / \tau)}}(128 \log n)^{\frac{10 \log (1 / \tau)}{\log (\rho / \tau)}} \tau^{-4} \tag{32}
\end{equation*}
$$

Now recall from (17) that $n^{\gamma} \leq t<2 n^{\gamma}$ and observe that

$$
1 \leq \frac{\log (1 / \tau)}{\log (\rho / \tau)}=\frac{1}{1-\log _{\tau} \rho} \leq \Delta
$$

Thus, from (32) we get

$$
\begin{equation*}
n^{2 \gamma}(128 \log n)^{10} \leq s<2^{1+2 \Delta} n^{2 \Delta \gamma}(128 \log n)^{10 \Delta} \tau^{-4} \tag{33}
\end{equation*}
$$

Since $\Delta \geq 1$ and $0<\gamma<1$ are constants, and recalling our assumption $\tau \geq n^{-\Theta(1)}$, both $s$ and $t$ grow at least as fast as a root function of $n$ and at most as fast as a polynomial function of $n$. Thus, high probabilities in $s$ and $t$ are high probabilities in $n$, and vice versa. It remains to justify our assumptions 28) and 29). To establish (28), observe that (31) and $t \geq n^{\gamma}$ imply $(\rho / \tau)^{p} \geq n^{\gamma}$. Furthermore, (33) and $\tau \geq n^{-\Theta(1)}$ imply $(\log s)^{2}=O\left((\log n)^{2}\right)$. Thus, 28) holds by 25). To establish $\sqrt{29}$, recall that $|\sigma| \geq \rho$ and hence it suffices to establish $(\rho / \tau)^{p / 2} \geq 4 \log s$. The reasoning is similar to what we used to establish 28). This completes the correctness proof.
4.4. Analysis (running time). We recall Theorem 1.7 ,

Theorem 1.7 (Main). For all constants $0<\gamma<1$ and $\Delta \geq 1$, the Outlier CorRELATIONS problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\begin{equation*}
\tilde{O}\left(n^{\max \{1-\gamma+M(\Delta \gamma, \gamma), M(1-\gamma, 2 \Delta \gamma)\}} \tau^{-4}\right) \tag{2}
\end{equation*}
$$

for approximate detection and subsequent $\tilde{O}\left(q d n^{2 \gamma}\right)$ time for exact listing of all the outliers, w.h.p. The running time bounds hold uniformly for all $n^{-\Theta(1)} \leq \tau<\rho<1$ with $\log _{\tau} \rho \leq 1-\Delta^{-1}$.

Proof. From (33) and (17) we have $s=\tilde{O}\left(n^{2 \Delta \gamma} \tau^{-4}\right)$ and $t=\Theta\left(n^{\gamma}\right)$. The matrices $\hat{A}_{I}$ and $\hat{B}_{I}$ are computed with $\bar{n} / t$ matrix multiplications, each with outer dimension $s^{1 / 2}=\tilde{O}\left(n^{\Delta \gamma} \tau^{-2}\right)$ and inner dimension $t=\Theta\left(n^{\gamma}\right)$. The exponent of running time for computing the matrices $\hat{A}_{I}$ and $\hat{B}_{I}$ from the input $A$ and $B$ is thus

$$
1-\gamma+M\left(\Delta \gamma-2 \log _{n} \tau, \gamma\right) \leq 1-\gamma+M(\Delta \gamma, \gamma)-4 \log _{n} \tau
$$

The matrix product $\hat{A}_{I}^{\top} \hat{B}_{I}$ has outer dimension $\bar{n} / t$ and inner dimension $s$. Thus, the exponent of running time for multiplying $\hat{A}_{I}^{\top}$ and $\hat{B}_{I}$ is

$$
M\left(1-\gamma, 2 \Delta \gamma-4 \log _{n} \tau\right) \leq M(1-\gamma, 2 \Delta \gamma)-4 \log _{n} \tau
$$

The running time (2) for approximate detection in Theorem 1.7 thus follows. (Indeed, the $\tilde{O}(\cdot)$-notation subsumes the polylogarithmic factors resulting from iterating the algorithm.)

To obtain the the running time for listing, by the analysis in $\$ 4.3$ at most $q$ entries of $\hat{A}_{I}^{\top} \hat{B}_{I}$ are marked with high probability in $n$. Each marked entry induces a computation of $t^{2}=\Theta\left(n^{2 \gamma}\right)$ inner products of dimension $d$, which results in the claimed running time $\tilde{O}\left(q d n^{2 \gamma}\right)$. This completes the proof of Theorem 1.7.

## 5. Corollaries

This section proves the corollaries in $\$ 1$ and $\$ 2$.
5.1. Proof of Corollary 1.8. We recall Corollary 1.8

Corollary 1.8. For all constants $0<\tau<\rho<1$, the Outlier CorrelaTIONS problem for Boolean inputs admits a randomized algorithm that runs in time $\tilde{O}\left(n^{\frac{2 \omega}{3-\log _{\tau} \rho}}\right)$ for approximate detection and subsequent $\tilde{O}\left(q d n^{\frac{2\left(1-\log _{\tau} \rho\right)}{3-\log _{\tau} \rho}}\right)$ time for exact listing of all the outliers, w.h.p.

Proof. Let us recall the following basic property of matrix multiplication exponents. We have

$$
M(\mu, \nu) \leq \begin{cases}(\omega-1) \mu+\nu & \text { if } \mu \leq \nu  \tag{34}\\ 2 \mu+(\omega-2) \nu & \text { if } \mu>\nu\end{cases}
$$

Let us now parameterize Theorem 1.7 to obtain Corollary 1.8. Let us take $\gamma=$ $\frac{1}{2 \Delta+1}$. In this case we have $1-\gamma=2 \Delta \gamma$. Recall the two terms in the maximum in (2). Since $\Delta \geq 1$, from (34) we have that the first term of the maximum in (2) is bounded by

$$
\begin{equation*}
1-\gamma+M(\Delta \gamma, \gamma) \leq \frac{4 \Delta-2+\omega}{2 \Delta+1} \tag{35}
\end{equation*}
$$

The second term of the maximum in $\sqrt{22}$ is

$$
\begin{equation*}
M(1-\gamma, 2 \Delta \gamma)=\frac{2 \Delta \omega}{2 \Delta+1} \tag{36}
\end{equation*}
$$

In particular, for $2 \leq \omega \leq 3$ and $\Delta \geq 1$ we observe that $2 \Delta \omega \geq 4 \Delta-2+\omega$, thus (36) dominates (35). Let us take $\Delta=\frac{1}{1-\log _{\tau} \rho}$ and observe that (36) simplifies to $\frac{2 \omega}{3-\log _{\tau} \rho}$. This establishes the running time for approximate detection in Corollary 1.8. The running time for listing is immediate by our choice of $\gamma$.
5.2. Proof of Corollary 1.9. We recall Corollary 1.9 .

Corollary 1.9. For all constants $0<\tau<\rho<1$, the Outlier CorrelaTIONS problem for Boolean inputs admits a randomized algorithm that runs in time $\tilde{O}\left(n^{\frac{4}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right)$ for approximate detection and subsequent $\tilde{O}\left(q d n^{\frac{2 \alpha\left(1-\log _{\tau} \rho\right)}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right)$ time for exact listing of all the outliers, w.h.p.

Proof. Let us take $\gamma=\frac{\alpha}{2 \Delta+\alpha}$, where $0.30298<\alpha \leq 1$ is the exponent for rectangular matrix multiplication [33], and recall that $\omega \leq 3-\alpha$. In particular, from (34) and $\Delta \geq 1$ we thus have

$$
\begin{align*}
1-\gamma+M(\Delta \gamma, \gamma) & \leq 1-\frac{\alpha}{2 \Delta+\alpha}+\frac{2 \alpha \Delta}{2 \Delta+\alpha}+\frac{(\omega-2) \alpha}{2 \Delta+\alpha} \\
& =\frac{2 \Delta+\alpha-\alpha+\alpha(2 \Delta+\omega-2)}{2 \Delta+\alpha} \\
& =\frac{2 \Delta(\alpha+1)+\alpha(\omega-2)}{2 \Delta+\alpha}  \tag{37}\\
& \leq \frac{2 \Delta(\alpha+1)+\alpha(1-\alpha)}{2 \Delta+\alpha} \\
& \leq \frac{4 \Delta}{2 \Delta+\alpha}
\end{align*}
$$

Furthermore, by definition of $\alpha$ and our choice of $\gamma$, we have

$$
\begin{equation*}
M(1-\gamma, 2 \Delta \gamma)=M\left(\frac{2 \Delta}{2 \Delta+\alpha}, \frac{2 \Delta \alpha}{2 \Delta+\alpha}\right)=\frac{4 \Delta}{2 \Delta+\alpha} \tag{38}
\end{equation*}
$$

Thus (38) dominates (37). Simplifying (38), we obtain $\frac{4}{2+\alpha\left(1-\log _{\tau} \rho\right)}$. This establishes the running time for approximate detection in Corollary 1.9. The running time for listing is immediate by our choice of $\gamma$.
5.3. Proof of Corollary 2.2. We recall Corollary 2.2

Corollary 2.2. For all constants $0<\epsilon<\omega / 3$, the Light Bulb problem admits a randomized algorithm that for all $d \geq 5 \rho^{-\frac{4 \omega}{9 \epsilon}-\frac{2}{3}} \log n$ runs in time

$$
\begin{equation*}
\tilde{O}\left(n^{\frac{2 \omega}{3}+\epsilon} \rho^{-\frac{8 \omega}{9 \epsilon}-\frac{4}{3}}\right) \tag{3}
\end{equation*}
$$

and finds the planted pair of vectors, with probability $1-o(1)$. The running time bound holds uniformly for all $n^{-\Theta(1)}<\rho<1$.

Proof. Fix a constant $0<\epsilon<\omega / 3$ and let $n^{-\Theta(1)}<\rho<1$ be given. With the objective of eventually applying Theorem 1.7, let us begin by selecting a suitable value of $\tau$. First, we want a small enough $0<\delta<1$ so that

$$
\frac{2 \omega}{3-\delta} \leq \frac{2 \omega}{3}+\epsilon
$$

or equivalently,

$$
\begin{equation*}
\delta \leq \frac{9 \epsilon}{2 \omega+3 \epsilon}=\frac{9}{2 \omega / \epsilon+3} . \tag{39}
\end{equation*}
$$

To obtain $\log _{\tau} \rho \leq \delta$, let us take $\tau=\rho^{1 / \delta}$. Because $0<\delta<1$ and $0<\rho<1$, we have $\tau<\rho$. Let us now check that for all large enough $d$, all the inner products between pairs of vectors in the input, except for the planted pair, are at most $\tau d$ in absolute value with probability $1-o(1)$. For $d \geq 5 \rho^{-2 / \delta} \log n$ it follows from the Hoeffding bound (6) and the union bound that with probability $1-o(1)$ all of the at most $n^{2}$ pairs of vectors that have at least one independent uniform random vector in the pair have inner product at most $\tau d$ in absolute value. Indeed, as $n \rightarrow \infty$ we have

$$
2 n^{2} \exp \left(-\frac{2(\tau d)^{2}}{4 d}\right)=2 n^{2} \exp \left(-\frac{\tau^{2} d}{2}\right) \leq 2 n^{2} \exp \left(-\frac{5}{2} \log n\right) \rightarrow 0
$$

It remains to find the planted pair with inner product at least $\rho d$ in absolute value. Let us apply the approximate detection algorithm in Theorem 1.7. To parameterize the algorithm, take $\Delta=1 /(1-\delta)$ and $\gamma=1 /(2 \Delta+1)$. Observe in particular that both $\Delta$ and $\gamma$ depend only on the constants $\delta$ and $\epsilon$, and do not depend on $\rho$ or $\tau$. Mimic the analysis in the proof of Corollary 1.8 to obtain from (2) the running time bound

$$
\begin{equation*}
\tilde{O}\left(n^{2 \omega / 3+\epsilon} \tau^{-4}\right) \tag{40}
\end{equation*}
$$

Observe furthermore that within the same time bound we can run the approximate detection algorithm recursively $O(\log \log n)$ times on the pair of blocks marked by the algorithm to find the planted pair, w.h.p. We obtain (3) by combining 40 with $\tau=\rho^{1 / \delta}$ and (39).
5.4. Proof of Corollary 2.4. We recall Corollary 2.4

Corollary 2.4. For all constants $0<\epsilon<\omega / 3$, the Parity with Noise problem admits a randomized algorithm that uses

$$
\begin{equation*}
d \geq(2 k+3) \cdot|1-2 \eta|^{-\frac{4 \omega}{9 \epsilon}-\frac{2}{3}} \log n \tag{4}
\end{equation*}
$$

examples, runs in time

$$
\begin{equation*}
\tilde{O}\left(n^{\frac{\omega+\epsilon}{3} k} \cdot|1-2 \eta|^{-\frac{8 \omega}{9 \epsilon}-\frac{4}{3}}\right), \tag{5}
\end{equation*}
$$

for all sufficiently large $k$, and finds the support of the parity function, with probability at least $1-o(1)$. The running time bound holds uniformly for all $n^{-\Theta(1)}<$ $|1-2 \eta|<1$.

Proof. This proof relies on a split-and-list idea outlined by [55, p. 32]; we present a proof here for completeness of exposition.

Let us start by setting up some notation. For a subset $A \subseteq[n]$ and a vector $\left(x_{1}, x_{2}, \ldots, x_{n}\right) \in\{-1,1\}^{n}$, let us write $x_{A}=\prod_{\ell \in A} x_{\ell}$. (For the empty set, define $x_{\emptyset}=1$.) Observe that for all $A, B \subseteq[n]$ it holds that $x_{A} x_{B}=x_{A \oplus B}$, where
$A \oplus B=(A \backslash B) \cup(B \backslash A)$ is the symmetric difference of $A$ and $B$. Let us write $\binom{[n]}{k}$ for the set of all $k$-subsets of $[n]$.

Fix a constant $0<\epsilon<\omega / 3$. Select a corresponding constant $0<\delta<1$ so that (39) holds. Let $0<\eta<1$ be the given noise rate with $n^{-\Theta(1)}<|1-2 \eta|<1$. Let $S \subseteq[n]$ be the support of the parity function, $|S|=k \geq 2$. Our task is to determine $S$ by drawing examples. Let us draw $d$ examples $(x, y) \in\{-1,1\}^{d} \times\{-1,1\}$. (We will fix a lower bound for $d$ in what follows.) Recall that each label $y$ in an example has the structure $y=z x_{S}$, where $z \in\{-1,1\}$ independently with $\operatorname{Pr}(z=-1)=\eta$.

Define a collection of $\binom{n}{\lfloor k / 2\rfloor}+\binom{n}{\lceil k / 2\rceil} \leq 2 n^{(k+1) / 2}$ vectors of dimension $d$ as follows. First, for each $J_{1} \in\binom{n}{\lfloor k / 2\rfloor}$, construct the vector $a_{J_{1}}$ whose entries are the values $x_{J_{1}}$, where $x$ ranges over the $d$ examples $(x, y)$ that we have drawn. Second, for each $J_{2} \in\binom{n}{[k / 2\rceil}$, construct the vector $b_{J_{2}}$ whose entries are the values $x_{J_{2}} y$, where $x$ and $y$ range over the $d$ examples $(x, y)$ that we have drawn.

Let us write $\operatorname{Bin}_{ \pm 1}(d, \beta)$ for the sum of $d$ random variables, each independently taking values in $\{-1,1\}$ such that $\beta$ is the probability of taking the value -1 .

Let us study the inner products between the vectors in our collection. We will use the notation " $\sum_{(x, y)}$ " to indicate a sum over the $d$ examples $(x, y)$ that we have drawn. Observe that for all $J_{1}, J_{1}^{\prime} \in\binom{n}{\lfloor k / 2\rfloor}$ we have

$$
\begin{equation*}
\left\langle a_{J_{1}}, a_{J_{1}^{\prime}}\right\rangle=\sum_{(x, y)} x_{J_{1}} x_{J_{1}^{\prime}}=\sum_{(x, y)} x_{J_{1} \oplus J_{1}^{\prime}} \tag{41}
\end{equation*}
$$

Similarly, for all $J_{2}, J_{2}^{\prime} \in\binom{n}{\lceil k / 2\rceil}$ we have

$$
\begin{equation*}
\left\langle b_{J_{2}}, b_{J_{2}^{\prime}}\right\rangle=\sum_{(x, y)} x_{J_{2}} y x_{J_{2}^{\prime}} y=\sum_{(x, y)} x_{J_{2} \oplus J_{2}^{\prime}} \tag{42}
\end{equation*}
$$

Finally, for all $J_{1} \in\binom{n}{\lfloor k / 2\rfloor}$ and $J_{2} \in\binom{n}{\lceil k / 2\rceil}$ we have

$$
\begin{equation*}
\left\langle a_{J_{1}}, b_{J_{2}}\right\rangle=\sum_{(x, y)} x_{J_{1}} x_{J_{2}} y=\sum_{(x, y)} x_{J_{1}} x_{J_{2}} x_{S} z=\sum_{(x, y)} x_{J_{1} \oplus J_{2} \oplus S} z \tag{43}
\end{equation*}
$$

Recalling that the vector $x$ in each example is an independent uniform random vector in $\{-1,1\}^{n}$, we observe ${ }^{14}$ from (41), 42), and (43) that each of the following inner products has distribution $\operatorname{Bin}_{ \pm 1}(d, 1 / 2)$ : (i) $\left\langle a_{J_{1}}, a_{J_{1}^{\prime}}\right\rangle$ for distinct $J_{1}$, $J_{1}^{\prime}$, (ii) $\left\langle b_{J_{2}}, b_{J_{2}^{\prime}}\right\rangle$ for distinct $J_{2}, J_{2}^{\prime}$, and (iii) $\left\langle a_{J_{1}}, b_{J_{2}}\right\rangle$ for all $J_{1}, J_{2}$ with $J_{1} \oplus J_{2} \neq S$. Furthermore, all the remaining inner products between distinct vectors in our collection, that is, $\left\langle a_{J_{1}}, b_{J_{2}}\right\rangle$ for all $J_{1}, J_{2}$ with $J_{1} \oplus J_{2}=S$, have distribution $\operatorname{Bin}_{ \pm 1}(d, \eta)$. This difference between distributions enables us to detect a pair $J_{1}, J_{2}$ with $J_{1} \oplus J_{2}=S$ and hence the set $S$.

Take $\rho=|1-2 \eta|$ and observe that $n^{-\Theta(1)}<\rho<1$. Let us take $\tau=\rho^{1 / \delta}$. Because $0<\delta<1$ and $0<\rho<1$, we have $\tau<\rho$. Furthermore, $\log _{\tau} \rho \leq \delta$. For

$$
\begin{equation*}
d \geq(2 k+3) \rho^{-2 / \delta} \log n \tag{44}
\end{equation*}
$$

it follows from the Hoeffding bound $\sqrt{6}$ and the union bound that with probability $1-o(1)$ all of the at most $2 n^{k+1}$ inner products with distribution $\operatorname{Bin}_{ \pm 1}(d, 1 / 2)$ are

[^9]at most $\tau d$ in absolute value. Indeed, as $n \rightarrow \infty$ we have
\[

$$
\begin{align*}
4 n^{k+1} \exp \left(-\frac{2(\tau d)^{2}}{4 d}\right) & =4 n^{k+1} \exp \left(-\frac{\tau^{2} d}{2}\right)  \tag{45}\\
& \leq 4 n^{k+1} \exp \left(-\frac{2 k+3}{2} \log n\right)=4 n^{-1 / 2} \rightarrow 0 .
\end{align*}
$$
\]

Observe that the expectation of $\operatorname{Bin}_{ \pm 1}(d, \eta)$ is $(1-2 \eta) d$ and that the absolute value of the expectation is $\rho d$. For $\eta \geq 1 / 2$ (respectively, for $\eta \leq 1 / 2$ ) the probability that $\operatorname{Bin}_{ \pm 1}(d, \eta)$ is at most (respectively, at least) its expectation is at least 1/4 [21, Theorem 1 and Corollary 3].

We want to consider two events over the $d$ examples drawn: (a) that the absolute value of at least one inner product with distribution $\operatorname{Bin}_{ \pm 1}(d, \eta)$ is at least $\rho d$, and (b) all of the at most $2 n^{k+1}$ inner products with distribution $\operatorname{Bin}_{ \pm 1}(d, 1 / 2)$ are at most $\tau d$. By above, (a) occurs with at least probability $1 / 4$, and, by (45), (b) occurs with probability $1-o(1)$.

Let us apply the approximate detection algorithm in Theorem 1.7. To parameterize the algorithm, take $\Delta=1 /(1-\delta)$ and $\gamma=1 /(2 \Delta+1)$. Observe in particular that both $\Delta$ and $\gamma$ depend only on the constants $\delta$ and $\epsilon$, and do not depend on $\rho$ or $\tau$. Mimic the analysis in the proof of Corollary 1.8 to obtain from (2) for all large enough $k$ the running time bound

$$
\begin{equation*}
\tilde{O}\left(n^{(\omega+\epsilon) k / 3} \tau^{-4}\right) \tag{46}
\end{equation*}
$$

Observe furthermore that within the same time bound we can run the approximate detection algorithm recursively $O(\log \log n)$ times on any one pair of blocks marked by the algorithm to find at least one inner product $\left\langle a_{J_{1}}, b_{J_{2}}\right\rangle$ with absolute value at least $\rho d$, and hence the set $S=J_{1} \oplus J_{2}$ with probability $1-o(1)$. Here we have of course conditioned on both events (a) and (b) above.

The time bound also allows us to increase the odds of success to $1-o(1)$ by performing $O(\log \log n)$ repetitions of the approximate detection algorithm, each time with new $d$ examples. Indeed, since the input is stochastic in nature, we have that (a) fails at most at probability $(3 / 4)^{\log \log n}$. Likewise, from $(45)$, we have the bound $O\left(n^{-1 / 2}\right)$ on the failure probability of (b). Hence, by the union bound, the total probability of input that causes the approximate detection algorithm to fail is bounded by $(3 / 4)^{\log \log n}+\frac{\log \log n}{n^{-1 / 2}} \rightarrow 0$, so the input enables us to correctly detect the parity with probability $1-o(1)$. It should also be noted that due to randomness in the approximate detection algorithm itself, the algorithm also succeeds with probability $1-o(1)$.

We obtain (5) by combining (46) with $\tau=\rho^{1 / \delta}$ and (39). Similarly, we obtain (4) by combining (44) with (39).
5.5. Proof of Theorem 2.5. We recall Theorem 2.5

Theorem 2.5 (Main, Two-level). For all constants $0<\gamma, \kappa<1$ and $\Delta \geq 1$, the Outlier Correlations problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\tilde{O}\left(n^{\max \{1-\gamma+M(\Delta \gamma, \gamma), M(1-\gamma, 2 \Delta \gamma)\}} \tau^{-4}\right)
$$

for approximate detection and subsequent

$$
\tilde{O}\left(q n^{\gamma \max \{1-\kappa+M(\Delta \kappa, \kappa), M(1-\kappa, 2 \Delta \kappa)\}} \tau^{-4}+q d n^{2 \gamma \kappa}\right)
$$

time for exact listing of all the outliers, w.h.p. The running time bounds hold uniformly for all $n^{-\Theta(1)} \leq \tau<\rho<1$ with $\log _{\tau} \rho \leq 1-\Delta^{-1}$.
Proof. The proof is otherwise identical to the proof of Theorem 1.7 in $\$ 4$ with the exception that we modify the listing phase as follows. For each of the at most $n^{2}$ pairs of blocks $k_{1}, k_{2} \in[\bar{n} / t]$ marked by the detection algorithm, we run the detection algorithm again, with parameter $\kappa$ replacing the parameter $\gamma$, and using the same value for the parameter $\Delta$. Indeed, each signalled pair of blocks can be viewed as an input of size $n^{\gamma}$ and dimension $d$. From these inputs we obtain as output at most $q$ pairs of blocks, each of size $n^{\gamma \kappa}$. Finally, we run the listing algorithm for these at most $q$ blocks. The running time and success probability follow immediately from Theorem 1.7 and the observation that the number of recursive invocations on inputs of size $n^{\gamma}$ is at most $q \leq n^{2}$, so by the union bound the conclusion holds w.h.p. in $n$.
5.6. Proof of Corollary 2.6. We recall Corollary 2.6

Corollary 2.6. For all constants $0<\tau<\rho<1$, the Outlier Correlations problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\tilde{O}\left(n^{\frac{2 \omega}{3-\log _{\tau} \rho}}\right)
$$

for approximate detection and subsequent

$$
\tilde{O}\left(q n^{\frac{2 \omega\left(1-\log _{\tau} \rho\right)}{\left(3-\log _{\tau} \rho\right)^{2}}}+q d n^{\frac{2\left(1-\log _{\tau} \rho\right)^{2}}{\left(3-\log _{\tau} \rho\right)^{2}}}\right)
$$

time for exact listing of all the outliers, w.h.p.
Proof. Analogous to Corollary 1.8, take $\gamma=\kappa=1 /(2 \Delta+1)$.
5.7. Proof of Corollary 2.7. We recall Corollary 2.7

Corollary 2.7. For all constants $0<\tau<\rho<1$, the Outlier Correlations problem for Boolean inputs admits a randomized algorithm that runs in time

$$
\tilde{O}\left(n^{\frac{4}{2+\alpha\left(1-\log _{\tau} \rho\right)}}\right)
$$

for approximate detection and subsequent

$$
\tilde{O}\left(q n^{\frac{4 \alpha\left(1-\log _{\tau} \rho\right)}{\left(2+\alpha\left(1-\log _{\tau} \rho\right)\right)^{2}}}+q d n^{\frac{2 \alpha^{2}\left(1-\log _{\tau} \rho\right)^{2}}{\left(2+\alpha\left(1-\log _{\tau} \rho\right)\right)^{2}}}\right)
$$

time for exact listing of all the outliers, w.h.p.
Proof. Analogous to Corollary 1.9, take $\gamma=\kappa=\alpha /(2 \Delta+\alpha)$.

## 6. A Lower Bound via Orthogonal Vectors

This section presents a local transformation from the orthogonal vectors problem (cf. [1, 62]) to Bichromatic Outlier Correlations. We present this transformation for completeness of exposition only and remark that it has been superseded by recent results of Ahle, Pagh, Razenshteyn, and Silvestri [3, Theorem 2]. First, let us recall the orthogonal vectors problem:
Problem 6.1 (Orthogonal Vectors). Given as input a set of $n$ vectors with dimension $d$ and entries in $\{0,1\}$, decide whether there is an orthogonal pair of vectors over the integers.

The sparsification lemma of Impagliazzo, Paturi, and Zane 26 and a lemma of Williams 61] yield the following conditional hardness result (cf. [62, Lemma A.1]) for Orthogonal Vectors:

Lemma 6.2 (Williams 61). Suppose there exists a constant $\delta>0$ and a randomized algorithm that for all constants $c \geq 1$ solves the Orthogonal Vectors problem with $d \leq c \log n$ in time $O\left(n^{2-\delta}\right)$, w.h.p. Then, the Strong Exponential Time Hypothesis is false.

The Orthogonal Vectors Conjecture (OVC) states that the algorithm assumed in Lemma 6.2 does not exist (cf. [1, 62]). Lemma 6.2 together with the following local transformation shows that Bichromatic Outlier Correlations cannot be solved in subquadratic time for inputs with $\log _{\tau} \rho$ arbitrarily close to 1 unless OVC and subsequently SETH are false:

Lemma 6.3. Suppose that there exists a constant $\delta>0$ and a randomized algorithm that for all constants $c \geq 1$ solves the Bichromatic Outlier Correlations problem for Boolean inputs with $d \leq c \log n$ and $|\rho-\tau| \leq 2 / d$ in time $O\left(n^{2-\delta}\right)$, w.h.p. Then, there exists a constant $\delta^{\prime}>0$ and a randomized algorithm that for all constants $c^{\prime} \geq 1$ solves the ORThogonal Vectors problem with $d^{\prime} \leq c^{\prime} \log n$ in time $O\left(n^{2-\delta^{\prime}}\right)$, w.h.p.
Proof. Without loss of generality we may work with a bichromatic version of Ortogonal Vectors where our input is a pair of matrices $S, T \in\{0,1\}^{d^{\prime} \times n}$. We transform the matrices $S, T$ into a pair of matrices $A, B \in\{-1,1\}^{\left(4 d^{\prime}+1\right) \times n}$ such that for all $j_{1}, j_{2} \in[n]$ it holds that $\left|\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle\right|$ is above a threshold value if and only if $\left\langle s_{j_{1}}, t_{j_{2}}\right\rangle=0$.

The local transformation is given by a function $h:\{0,1\}^{2} \rightarrow \mathbb{Z}$ that decomposes into two functions $u, v:\{0,1\} \rightarrow\{-1,1\}^{k}$ such that $h(x, y)=\langle u(x), v(y)\rangle$ with $h(0,0)=h(1,0)=h(0,1)=k_{1}$ for some integer $k_{1}>0$, and $h(1,1)=k_{2}$ for some integer $k_{2}<0$. One such pair of functions with $k=4, k_{1}=2$, and $k_{2}=-2$ is the following:

$$
\begin{aligned}
& u(x)=\left[\begin{array}{rrrr}
1 & 1-2 x & 1 & 1-2 x
\end{array}\right]^{\top}, \\
& v(y)=\left[\begin{array}{rrrr}
1 & 1 & 1-2 y & 2 y-1
\end{array}\right]^{\top} .
\end{aligned}
$$

Apply the function $u$ (respectively, the function $v$ ) to each element of $S$ (respectively, $T$ ) to produce a matrix $A$ (respectively, $B$ ) of size $4 d^{\prime} \times n$. Finally, append one full row of 1-entries to both matrices $A$ and $B$.

The matrices $A$ and $B$ yield a $\{-1,1\}$-valued input to OUTlier CorrelaTIONS with the following parameters: the number of vectors is $n$, the dimension is $d=4 d^{\prime}+1$, the outlier parameter is $\rho=\left(2 d^{\prime}+1\right) /\left(4 d^{\prime}+1\right)$, and the background parameter $\tau=\left(2 d^{\prime}-1\right) /\left(4 d^{\prime}+1\right)$. In particular, we observe that the inner products $\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle$ lie in the interval $\left[-2 d^{\prime}+1,2 d^{\prime}+1\right]$, and we have $\left|\left\langle a_{j_{1}}, b_{j_{2}}\right\rangle\right| \geq \rho\left(4 d^{\prime}+1\right)$ if and only if $\left\langle s_{j_{1}}, t_{j_{2}}\right\rangle=0$. Furthermore, without loss of generality we may assume that the instance $A, B$ has at most $q=O\left(n^{2-\eta}\right)$ outlier pairs for a constant $\eta>0$. Indeed, we may test uniformly at random $n^{1.1 \eta}$ pairs of vectors for outlier pairs. If we find an outlier pair, we are done; otherwise we can conclude that w.h.p. in $n$ the instance has at most $q$ outlier pairs. In the latter case we run the assumed algorithm for Outlier Correlations to find a pair if it exists. Observe that for our choice of $\rho, \tau$ we have $|\rho-\tau|=2 /\left(4 d^{\prime}+1\right)=2 / d$.

Remark 6.4. We can replace the assumption $|\rho-\tau| \leq 2 / d$ in Lemma 6.3 with

$$
\log _{\tau} \rho \geq \frac{\log \frac{2 d^{\prime}+1}{4 d^{\prime}+1}}{\log \frac{2 d^{\prime}-1}{4 d^{\prime}+1}}=1-\frac{\log \left(1+\frac{4}{d-3}\right)}{\log \left(2+\frac{6}{d-3}\right)}=1-\Theta(1 / d)
$$

## Acknowledgement

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement 338077 "Theory and Practice of Advanced Search and Enumeration."

A preliminary conference abstract of this paper appeared in Robert Krauthgamer (Ed.), Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA 2016, Arlington, VA, January 10-12, 2016), SIAM, Philadelphia, PA, 2016, pp. 1288-1305.

We thank the anonymous referees for their helpful comments and Rasmus Pagh and Thomas Dybdahl Ahle for useful discussions. Work done in part while the second author was visiting the Simons Institute for the Theory of Computing.

## References

[1] Abboud, A., Williams, R. R., and Yu, H. More applications of the polynomial method to algorithm design. In Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Philadelphia, PA, USA, 2015), Society for Industrial and Applied Mathematics (SIAM), pp. 218-230.
[2] Achlioptas, P., Schölkopf, B., and Borgwardt, K. M. Two-locus association mapping in subquadratic time. In Proceedings of the 17 th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (New York, NY, USA, 2011), ACM, pp. 726-734.
[3] Ahle, T. D., Pagh, R., Razenshteyn, I., and Silvestri, F. On the complexity of inner product similarity join. In Proceedings of the 35th ACM Symposium on Principles of Database Systems (New York, NY, USA, 2016), PODS, ACM, pp. 151-164.
[4] Alman, J., and Williams, R. Probabilistic polynomials and Hamming nearest neighbors. In 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS) (Los Alamitos, CA, USA, 2015), IEEE Computer Society, pp. 136-150.
[5] Andoni, A., And Indyk, P. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS) (Los Alamitos, CA, USA, 2006), IEEE Computer Society, pp. 459-468.
[6] Andoni, A., Indyk, P., Laarhoven, T., Razenshteyn, I., and Schmidt, L. Practical and optimal lsh for angular distance. In Advances in Neural Information Processing Systems 28 (NIPS) (Red Hook, NY, USA, 2015), Curran Associates, Inc., pp. 1225-1233.
[7] Andoni, A., Indyk, P., Nguyen, H. L., and Razenshteyn, I. Beyond locality-sensitive hashing. In Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Philadelphia, PA, USA, 2014), Society for Industrial and Applied Mathematics (SIAM), pp. 1018-1028.
[8] Andoni, A., and Razenshteyn, I. Optimal data-dependent hashing for approximate near neighbors. In Proceedings of the 47 th Annual ACM on Symposium on Theory of Computing (STOC) (New York, NY, USA, 2015), ACM, pp. 793-801.
[9] Bayardo, R. J., Ma, Y., and Srikant, R. Scaling up all pairs similarity search. In Proceedings of the 16th International Conference on World Wide Web (WWW) (New York, NY, USA, 2007), ACM, pp. 131-140.
[10] Benson, A. R., and Ballard, G. A framework for practical parallel fast matrix multiplication. In Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (New York, NY, USA, 2015), PPoPP 2015, ACM, pp. 42-53.
[11] Blum, A., Kalai, A., and Wasserman, H. Noise-tolerant learning, the parity problem, and the statistical query model. J. ACM 50, 4 (2003), 506-519.
[12] Charikar, M. Similarity estimation techniques from rounding algorithms. In Proceedings on 34 th Annual ACM Symposium on Theory of Computing (STOC) (New York, NY, USA, 2002), ACM, pp. 380-388.
[13] Chaudhuri, S., Ganti, V., and Kaushik, R. A primitive operator for similarity joins in data cleaning. In Proceedings of the 22nd International Conference on Data Engineering (ICDE) (Los Alamitos, CA, USA, 2006), IEEE Computer Society, p. 5.
[14] Cohen, E., Datar, M., Fujiwara, S., Gionis, A., Indyk, P., Motwani, R., Ullman, J. D., and Yang, C. Finding interesting associations without support pruning. IEEE Transactions on Knowledge and Data Engineering 13, 1 (2001), 64-78.
[15] Curtin, R. R., Gray, A. G., and Ram, P. Fast exact max-kernel search. In Proceedings of the 13th SIAM International Conference on Data Mining (SDM) (Philadelphia, PA, USA, 2013), Society for Industrial and Applied Mathematics (SIAM), pp. 1-9.
[16] Das, A., Datar, M., Garg, A., and Rajaram, S. Google news personalization: scalable online collaborative filtering. In Proceedings of the 16th International Conference on World Wide Web (WWW) (New York, NY, USA, 2007), ACM, pp. 271-280.
[17] Drineas, P., Kannan, R., and Mahoney, M. W. Fast Monte Carlo algorithms for matrices. I. Approximating matrix multiplication. SIAM J. Comput. 36, 1 (2006), 132-157.
[18] Dubiner, M. Bucketing coding and information theory for the statistical high-dimensional nearest-neighbor problem. IEEE Transactions on Information Theory 56, 8 (2010), 41664179.
[19] Feldman, V., Gopalan, P., Khot, S., and Ponnuswami, A. K. On agnostic learning of parities, monomials, and halfspaces. SIAM J. Comput. 39, 2 (2009), 606-645.
[20] Galton, F. Co-relations and their measurement, chiefly from anthropometric data. Proceedings of the Royal Society of London 45 (1888), 135-145.
[21] Greenberg, S., and Mohri, M. Tight lower bound on the probability of a binomial exceeding its expectation. Statistics and Probability Letters 86 (2014), 91-98.
[22] Grigorescu, E., Reyzin, L., and Vempala, S. On noise-tolerant learning of sparse parities and related problems. In Proceedings of the 22nd International Conference on Algorithmic Learning Theory (ALT) (Berlin, Germany, 2011), Springer, pp. 413-424.
[23] Hoeffding, W. Probability inequalities for sums of bounded random variables. J. Amer. Statist. Assoc. 58 (1963), 13-30.
[24] Huang, J., Rice, L., Matthews, D. A., and van de Geijn, R. A. Generating families of practical fast matrix multiplication algorithms. In 2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS) (Los Alamitos, CA, USA, 2017), IEEE Computer Society, pp. 656-667.
[25] Impagliazzo, R., and Paturi, R. On the complexity of $k$-SAT. J. Comput. Syst. Sci. 62, 2 (2001), 367-375.
[26] Impagliazzo, R., Paturi, R., and Zane, F. Which problems have strongly exponential complexity? J. Comput. Syst. Sci. 63, 4 (2001), 512-530.
[27] Indyk, P. Nearest neighbors in high-dimensional spaces. In Handbook of Discrete and Computational Geometry, second ed., Discrete Mathematics and its Applications (Boca Raton). Chapman \& Hall/CRC, Boca Raton, FL, USA, 2004, pp. 877-892.
[28] Indyk, P., and Motwani, R. Approximate nearest neighbors: Towards removing the curse of dimensionality. In Proceedings of the 30th Annual ACM Symposium on the Theory of Computing (STOC) (New York, NY, USA, 1998), ACM, pp. 604-613.
[29] Kapralov, M. Smooth tradeoffs between insert and query complexity in nearest neighbor search. In Proceedings of the 34th ACM Symposium on Principles of Database Systems (PODS) (New York, NY, USA, 2015), ACM, pp. 329-342.
[30] Karppa, M., Kaski, P., Kohonen, J., and Ó Catháin, P. Explicit correlation amplifiers for finding outlier correlations in deterministic subquadratic time. In 24 th Annual European Symposium on Algorithms, ESA 2016 (Dagstuhl, Germany, 2016), Schloss Dagstuhl-LeibnizZentrum für Informatik, pp. 52:1-52:17.
[31] Koenigstein, N., Ram, P., and Shavitt, Y. Efficient retrieval of recommendations in a matrix factorization framework. In 21st ACM International Conference on Information and Knowledge Management (CIKM) (New York, NY, USA, 2012), ACM, pp. 535-544.
[32] Kushilevitz, E., Ostrovsky, R., and Rabani, Y. Efficient search for approximate nearest neighbor in high dimensional spaces. In Proceedings of the 30th Annual ACM Symposium on the Theory of Computing (STOC) (New York, NY, USA, 1998), ACM, pp. 614-623.
[33] Le Gall, F. Faster algorithms for rectangular matrix multiplication. In 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS) (Los Alamitos, CA, USA, 2012), IEEE Computer Society, pp. 514-523.
[34] Le Gall, F. Powers of tensors and fast matrix multiplication. In International Symposium on Symbolic and Algebraic Computation (ISSAC) (New York, NY, USA, 2014), ACM, pp. 296303.
[35] Lee, D., Park, J., Shim, J., and Lee, S. An efficient similarity join algorithm with cosine similarity predicate. In Proceedings of the 21st International Conference on Database and Expert Systems Applications (DEXA) (Berlin, Germany, 2010), Springer, pp. 422-436.
[36] Low, Y., AND Zheng, A. X. Fast top-k similarity queries via matrix compression. In 21st ACM International Conference on Information and Knowledge Management (CIKM) (New York, NY, USA, 2012), ACM, pp. 2070-2074.
[37] Lyubashevsky, V. The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In Proceedings of the 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX) and the 9th International Workshop on Randomization and Computation (RANDOM) (Berlin, Germany, 2005), Springer, pp. 378-389.
[38] May, A., and Ozerov, I. On computing nearest neighbors with applications to decoding of binary linear codes. In Proceedings of Advances in Cryptology - EUROCRYPT 2015 34 th Annual International Conference on the Theory and Applications of Cryptographic Techniques (Berlin, Germany, 2015), Springer, pp. 203-228.
[39] Meiser, S. Point location in arrangements of hyperplanes. Inf. Comput. 106, 2 (1993), 286303.
[40] Minsky, M. L., and Papert, S. A. Perceptrons: An Introduction to Computational Geometry. MIT Press, Boston, MA, USA, 1969.
[41] Mossel, E., O'Donnell, R., and Servedio, R. A. Learning functions of k relevant variables. J. Comput. Syst. Sci. 69, 3 (2004), 421-434.
[42] Motwani, R., Naor, A., and Panigrahy, R. Lower bounds on locality sensitive hashing. SIAM J. Discrete Math. 21, 4 (2007), 930-935.
[43] Neyshabur, B., and Srebro, N. On symmetric and asymmetric LSHs for inner product search. In Proceedings of the 32nd International Conference on Machine Learning (ICML) (2015), Journal of Machine Learning Research, pp. 1926-1934.
[44] O’Donnell, R., Wu, Y., and Zhou, Y. Optimal lower bounds for locality-sensitive hashing (except when q is tiny). ACM TOCT 6, 1 (2014), Article 5.
[45] Pagh, R. Compressed matrix multiplication. ACM TOCT 5, 3 (2013), Article 9.
[46] Pagh, R. Locality-sensitive hashing without false negatives. In Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Philadelphia, PA, USA, 2016), Society for Industrial and Applied Mathematics (SIAM), pp. 1-9.
[47] Paturi, R., Rajasekaran, S., and Reif, J. H. The light bulb problem. In Proceedings of the Second Annual Workshop on Computational Learning Theory (COLT) (New York, NY, USA, 1989), ACM, pp. 261-268.
[48] Pearson, K. Notes on the history of correlation. Biometrika 13, 1 (1920), 25-45.
[49] Ram, P., and Gray, A. G. Maximum inner-product search using cone trees. In The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) (New York, NY, USA, 2012), ACM, pp. 931-939.
[50] Shrivastava, A., and Li, P. Asymmetric LSH (ALSH) for sublinear time maximum inner product search (MIPS). In Annual Conference on Neural Information Processing Systems (NIPS) (Red Hook, NY, USA, 2014), Curran Associates, Inc., pp. 2321-2329.
[51] Shrivastava, A., and Li, P. Asymmetric minwise hashing for indexing binary inner products and set containment. In Proceedings of the 24th International Conference on World Wide Web (WWW) (New York, NY, USA, 2015), ACM, pp. 981-991.
[52] Silva, Y. N., Aref, W. G., and Ali, M. H. The similarity join database operator. In Proceedings of the 26th International Conference on Data Engineering (ICDE) (Los Alamitos, CA, USA, 2010), IEEE Computer Society, pp. 892-903.
[53] Stigler, S. M. Francis Galton's account of the invention of correlation. Statistical Science 4, 2 (1989), 73-86.
[54] Teflioudi, C., Gemulla, R., and Mykytiuk, O. LEMP: fast retrieval of large entries in a matrix product. In Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (New York, NY, USA, 2015), ACM, pp. 107-122.
[55] Valiant, G. Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem. J. ACM 62, 2 (2015), Article 13.
[56] Valiant, L. G. Functionality in neural nets. In Proceedings of the First Annual Workshop on Computational Learning Theory (COLT) (New York, NY, USA, 1988), ACM, pp. 28-39.
[57] Wang, J., Li, G., and Feng, J. Fast-join: An efficient method for fuzzy token matching based string similarity join. In Proceedings of the 27th International Conference on Data Engineering (ICDE) (Los Alamitos, CA, USA, 2011), IEEE Computer Society, pp. 458-469.
[58] Wang, J., Li, G., and Feng, J. Can we beat the prefix filtering?: an adaptive framework for similarity join and search. In Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD) (New York, NY, USA, 2012), ACM, pp. 85-96.
[59] Wang, Y., Metwally, A., and Parthasarathy, S. Scalable all-pairs similarity search in metric spaces. In The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) (New York, NY, USA, 2013), ACM, pp. 829-837.
[60] Weber, R., Schek, H., and Blott, S. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. In Proceedings of 24 rd International Conference on Very Large Data Bases (VLDB) (Burlington, MA, USA, 1998), Morgan Kaufmann Publishers, pp. 194-205.
[61] Williams, R. A new algorithm for optimal 2-constraint satisfaction and its implications. Theor. Comput. Sci. 348, 2-3 (2005), 357-365.
[62] Williams, R., and Yu, H. Finding orthogonal vectors in discrete structures. In Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Philadelphia, PA, USA, 2014), Society for Industrial and Applied Mathematics (SIAM), pp. 1867-1877.
[63] Xia, C., Lu, H., Ooi, B. C., And Hu, J. Gorder: An efficient method for KNN join processing. In Proceedings of the 30th International Conference on Very Large Data Bases (VLDB) (Burlington, MA, USA, 2004), Morgan Kaufmann Publishers, pp. 756-767.
[64] Xiao, C., Wang, W., Lin, X., Yu, J. X., and Wang, G. Efficient similarity joins for nearduplicate detection. ACM Trans. Database Syst. 36, 3 (2011), Article 15.
[65] Zadeh, R. B., and Goel, A. Dimension independent similarity computation. Journal of Machine Learning Research 14, 1 (2013), 1605-1626.
[66] Zezula, P., Amato, G., Dohnal, V., and Batko, M. Similarity Search. The Metric Space Approach, vol. 32 of Advances in Database Systems. Springer, New York, NY, USA, 2006.


[^0]:    Key words and phrases. correlation, fast matrix multiplication, light bulb problem, rectangular matrix multiplication, similarity search.

[^1]:    ${ }^{1}$ For brevity, we will use the expression "subquadratic in $n$ " to mean $O\left(n^{2-\epsilon}\right)$ for some constant $\epsilon>0$ independent of $n$.
    ${ }^{2}$ Suppose the data matrix for Outlier Correlations is $D \in\{-1,1\}^{d \times n}$ and we are interested in discovering the pairs of distinct columns in $D$ with (absolute) inner product at least $\rho d$. Construct $\left\lceil\log _{2} n\right\rceil$ pairs of matrices $A_{k}, B_{k}$ for $k=0,1, \ldots,\left\lceil\log _{2} n\right\rceil-1$, where the matrix $A_{k}$ (respectively, $B_{k}$ ) contains the columns of $D$ whose index, viewed as a $\left\lceil\log _{2} n\right\rceil$-bit string, has a 0 (respectively, 1) in bit position $k$. Take the union of the solutions of the instances $A_{k}, B_{k}$ to obtain the solution for $D$. All distinct columns will be discovered because each pair of distinct columns differs in at least one bit position $k$.
    ${ }^{3}$ An arbitrary input on the unit Euclidean sphere can be quantized to a $\{-1,1\}$-valued input if one is willing to accept randomization and an expected loss $\lambda \mapsto 1-\frac{2}{\pi} \arccos \lambda$ in the parameters $\lambda \in\{\tau, \rho\}$ caused by rounding the arbitrary input into a Boolean input with random hyperplanes. See e.g. 55] Algorithm 4] or [12].

[^2]:    ${ }^{4}$ By (17, two Boolean vectors with inner product $\lambda d$ have Hamming distance $(1-\lambda) d / 2$. Thus, assuming there are no pairs whose inner product is strictly between $\tau$ and $\rho$, the vectors that belong to outlier pairs can be discovered by selecting a small enough $\epsilon$ so that an $(1+\epsilon)$-approximate nearest neighbor of a vector has to be a completion of such a vector to an outlier pair if such a completion exists. For example, by 1 we can take $1+\epsilon<(1-\tau) /(1-\rho)$ for the Hamming metric and $(1+\epsilon)^{2}<(1-\tau) /(1-\rho)$ for the Euclidean metric.
    ${ }^{5}$ The notation $\tilde{O}(\cdot)$ suppresses factors polylogarithmic in $n$ and $d$ whose degree may depend on $\rho$ and $\tau$.

[^3]:    ${ }^{6}$ We adopt the convention that all running time bounds that contain an exponent $f(\omega)$ of $n$ that depends on $\omega$ (or any other limiting exponent for matrix multiplication, such as $\alpha$ ) are tacitly stated as $f(\omega)+\epsilon$ for an arbitrarily small constant $\epsilon>0$.
    ${ }^{7}$ We may stop the algorithm after the approximate detection phase if we only want to decide whether the input contains at least one pair with absolute inner product at least $\rho d$, with the following approximation guarantees: (i) if all the inner products in the input have absolute value at most $\tau d$, the algorithm outputs false w.h.p.; and (ii) if the input contains at least one inner product with absolute value at least $\rho d$, the algorithm outputs true w.h.p.

[^4]:    ${ }^{8}$ The algorithm is iterated $\Theta(\log n)$ times to guarantee that every outlier pair is signalled in at least one pair of blocks during at least one iteration with high probability.

[^5]:    ${ }^{9}$ It follows from the Hoeffding bound (6) and the union bound that there is a constant $c>0$ such that for $d \geq c \rho^{-2} \log n$ with probability $1-o(1)$ the planted pair is the unique pair of vectors in the input with inner product at least $\rho d$ in absolute value.

[^6]:    ${ }^{10}$ Suppose we have $n$ observables in a database and $n$ observables that constitute queries to the database. Suppose furthermore that for each query there are $O(1)$ outlier-correlated observables in the database, and that our task is to find these outlier-correlated observables for each query.

[^7]:    ${ }^{11}$ We observe in particular these matrices are essentially what results if we execute phases and 1 and 2 of Valiant's algorithm (recall $\$ 1.4$, with the difference that our sample $I$ always has Cartesian product structure, whereas Valiant's algorithm uses a uniform random sample of $[d]^{p}$.

[^8]:    ${ }^{12}$ By a root function of $n$ we mean a function $n^{\beta}$ for a constant $0<\beta<1$.
    ${ }^{13}$ Observe that for all large enough positive real $x$ the interval $[x, 2 x]$ contains at least one integer square. Also observe that we have not yet fixed our value of $p$, but will do so later.

[^9]:    ${ }^{14}$ Indeed, for all nonempty $K \subseteq[n]$, the random variables $\sum_{(x, y)} x_{K}$ and $\sum_{(x, y)} x_{K} z$ have distribution $\operatorname{Bin}_{ \pm 1}(d, 1 / 2)$. For the symmetric difference it holds that $K_{1} \oplus K_{2}=\emptyset$ if and only if $K_{1}=K_{2}$.

