

---

# A Subquadratic Time Algorithm for Robust Sparse Mean Estimation

---

Ankit Pensia<sup>1</sup>

## Abstract

We study the algorithmic problem of sparse mean estimation in the presence of adversarial outliers. Specifically, the algorithm observes a *corrupted* set of samples from  $\mathcal{N}(\mu, \mathbf{I}_d)$ , where the unknown mean  $\mu \in \mathbb{R}^d$  is constrained to be  $k$ -sparse. A series of prior works has developed efficient algorithms for robust sparse mean estimation with sample complexity  $\text{poly}(k, \log d, 1/\epsilon)$  and runtime  $d^2 \text{poly}(k, \log d, 1/\epsilon)$ , where  $\epsilon$  is the fraction of contamination. In particular, the fastest runtime of existing algorithms is quadratic in the dimension, which can be prohibitive in high dimensions. This quadratic barrier in the runtime stems from the reliance of these algorithms on the sample covariance matrix, which is of size  $d^2$ .

Our main contribution is an algorithm for robust sparse mean estimation which runs in *subquadratic* time using  $\text{poly}(k, \log d, 1/\epsilon)$  samples, with similar results for robust sparse PCA. Our results build on algorithmic advances in detecting weak correlations, a generalized version of the light-bulb problem by Valiant (Valiant, 2015).

## 1. Introduction

Mean estimation, a fundamental unsupervised inference task studied in literature, may be described as follows: Given a family of distributions  $\mathcal{P}$  over  $\mathbb{R}^d$ , the algorithm observes a set of i.i.d. points from an unknown  $P \in \mathcal{P}$ , with the goal of outputting  $\hat{\mu}$  such that, with high probability,  $\|\hat{\mu} - \mu\|_2$  is small. Although this framework is well-studied in the literature, the data observed in practice may deviate from the i.i.d. assumption and additionally may contain outliers. Crucially, these outliers can easily break standard off-the-shelf estimators, for example, sample mean, geometric median, and coordinate-wise median. To address this challenge, the field of robust statistics was initiated in the 1960s, aiming to

develop algorithms that are robust to outliers (Huber, 1964; Andrews et al., 1972; Huber & Ronchetti, 2009). Before proceeding further, we formally define the contamination model we study in this paper.

**Definition 1.1** (Strong Contamination Model). Given a *corruption* parameter  $\epsilon \in (0, 1/2)$  and a distribution  $P$  on uncorrupted samples, an algorithm obtains samples from  $P$  with  $\epsilon$ -contamination as follows: (i) The algorithm specifies the number  $n$  of samples it requires. (ii) A set  $S$  of  $n$  i.i.d. samples from  $P$  is drawn but not yet shown to the algorithm. (iii) An arbitrarily powerful adversary then inspects  $S$ , before deciding to replace any subset of  $\lceil \epsilon n \rceil$  samples with arbitrarily corrupted points (“outliers”) to obtain the contaminated set  $T$ , which is then returned to the algorithm. We say  $T$  is an  $\epsilon$ -corrupted version of  $S$  and a set of  $\epsilon$ -corrupted samples from  $P$ .

Our focus will be on high-dimensional distributions, i.e., when the distribution  $P$  is over  $\mathbb{R}^d$  for large  $d$ . Dealing with outliers becomes harder in high dimensions because *classical* outlier screening procedures (which otherwise work well in low dimensions) rely on the norm of the data points and are too coarse to distinguish outliers from inliers. Nevertheless, a long line of research, spurred by advances in Diakonikolas et al. (2016); Lai et al. (2016), has developed a systematic theory of handling outliers in high-dimensional robust statistics (Diakonikolas & Kane, 2023). Notwithstanding this progress, major gaps persist in our *fine-grained* understanding of fast robust algorithms for data with additional structure.

Structured high-dimensional data distributions are ubiquitous in practice, e.g., natural images and sounds. Moreover, leveraging these underlying structures often dramatically improves algorithmic performance, e.g., in terms of error. A well-studied structure both in the theory and practice of high-dimensional statistics is *sparsity*, see, for example, the textbooks Eldar & Kutyniok (2012); Hastie et al. (2015); van de Geer (2016). Consequently, we concentrate our efforts on *structured* mean estimation, where we assume that the underlying mean is sparse, i.e., an overwhelming majority of its coordinates are zero.

In light of the challenges posed by outliers above and the prevalence and importance of sparsity, we study the problem of *robust* sparse mean estimation. We say a vector  $x \in \mathbb{R}^d$  is

---

<sup>1</sup>IBM Research. Correspondence to: Ankit Pensia <ankitp@ibm.com>.

$k$ -sparse if  $x$  has at most  $k$  non-zero entries. Our focus is on the practically relevant regime where  $k$  is much smaller than  $d$ , say, poly-logarithmic in  $d$ . We formally define robust sparse mean estimation below.

**Problem 1.2** (Gaussian Robust Sparse Mean Estimation). Let  $\epsilon_0 \in (0, 1/2)$  be a sufficiently small constant. Given  $\epsilon \in (0, \epsilon_0)$ , sparsity  $k \in \mathbb{N}$ , and a set of  $\epsilon$ -corrupted set of samples from  $\mathcal{N}(\mu, \mathbf{I}_d)$  with an unknown  $k$ -sparse mean  $\mu \in \mathbb{R}^d$ , the goal is to output an estimate  $\hat{\mu} \in \mathbb{R}^d$  such that  $\|\hat{\mu} - \mu\|_2$  is small with high probability.

Robust sparse mean estimation algorithms, efficient both in runtime and samples, were first developed in Balakrishnan et al. (2017), with sample complexity  $n = \text{poly}(k, \log d, 1/\epsilon)$ , runtime  $\text{poly}(d, n, 1/\epsilon)$ , and near-optimal error  $\|\hat{\mu} - \mu\|_2 = \tilde{O}(\epsilon)$ .<sup>1</sup> In particular, the sample complexity is only poly-logarithmic in the ambient dimension  $d$ , thereby permitting statistical inference with far fewer samples than the  $\Omega(d)$  samples required by unstructured mean estimation. Therefore, for our algorithm for Problem 1.2, we set as the first requirement this sample complexity of  $\text{poly}(k, \log d, 1/\epsilon)$ .

The focus of this work is to develop *fast* robust sparse mean estimation with the aforementioned sample complexity. Although the runtime of Balakrishnan et al. (2017) is polynomial in dimension, their algorithm uses the ellipsoid algorithm (which in turn solves a semidefinite program) and hence is not practical in high dimensions. Diakonikolas et al. (2019) then developed a *spectral* algorithm with similar error guarantees and sample complexity and an improved runtime of  $d^2 \text{poly}(k, \log d, 1/\epsilon)$ . Subsequent papers have proposed many algorithmic improvements and generalizations to a wider class of distributions (Zhu et al., 2022; Cheng et al., 2022; Diakonikolas et al., 2022a;c); see Section 1.3.

Despite this algorithmic progress, the fastest currently known algorithm for Problem 1.2 is that of Diakonikolas et al. (2019) with runtime scaling as  $d^2$ . This quadratic runtime of the algorithm can be prohibitive in high dimensions—the very setting that benefits most from sparsity (because of sample-efficiency). This quadratic dimension dependence is in stark contrast to the *non-robust* setting (i.e., the outlier-free regime), where there exists a simple (folklore) algorithm<sup>2</sup> with nearly-linear runtime, which is also minimax optimal. This motivates the following fundamental question:

<sup>1</sup>In the presence of outliers, vanishing error is usually not possible. In our setting, this is because it is impossible to distinguish two isotropic Gaussian distributions that are  $\Omega(\epsilon)$ -far apart in the presence of  $\epsilon$ -fraction of contamination.

<sup>2</sup>The algorithm computes the sample mean and thresholds entries to ensure sparsity, hence failing if there is even a single outlier. Moreover, natural attempts to make this algorithm robust, such as coordinate-wise median, incur an highly suboptimal error of  $\Omega(\epsilon\sqrt{k})$ .

**Question 1.3.** *Are there any nearly-linear time algorithms for robust sparse mean estimation?*

If we momentarily forgo sparsity (and the benefits that come along with it, e.g., the reduced sample complexity and interpretability) and focus on robust *dense* estimation, then positive answers are known to Question 1.3, see, e.g., Cheng et al. (2019a); Dong et al. (2019); Diakonikolas et al. (2022b;d). However, the sample complexities of the algorithms in these papers scale *linearly* with dimension,<sup>3</sup> which considerably exceeds our allowed budget of  $\text{poly}(k, \log d, 1/\epsilon)$  samples.

In fact, as alluded to earlier, existing attempts at answering Question 1.3 do not even break the *quadratic* runtime barrier. This is due to natural technical obstacles within current algorithms: to robustly estimate the mean, they crucially rely on the sample covariance matrix to detect outliers; but merely computing the sample covariance matrix costs  $\Omega(d^2)$  time! Sparsity also precludes common tricks such as the power iteration to bypass explicitly writing the covariance matrix. Indeed, in certain parameter regimes, even detecting atypical values of the covariance matrix from samples is conjectured to require  $\Omega(d^2)$  time (Dagan & Shamir, 2018). This begs the question whether this quadratic gap is inherent:

**Question 1.4.** *Is there an algorithm for robust sparse mean estimation that runs in  $d^{2-\Omega(1)}$  time and uses  $\text{poly}(k, \log d, 1/\epsilon)$  samples?*

The main result of our work is an affirmative answer to Question 1.4. We hope our answer paves the path for progress towards answering Question 1.3, which was highlighted as an important open problem in Diakonikolas (2019); Cheng (2021); Diakonikolas (2023). Our algorithm builds on advances in fast correlation detection algorithms by Valiant (Valiant, 2015).

## 1.1. Our Results

We establish the following result:

**Theorem 1.5** (Robust Sparse Mean Estimation in Subquadratic Time). *Let the contamination rate be  $\epsilon \in (0, \epsilon_0)$  for a small constant  $\epsilon_0 \in (0, 1/2)$  and  $k \in \mathbb{N}$  be the sparsity. Let  $T$  be an  $\epsilon$ -corrupted set of  $n$  samples from  $\mathcal{N}(\mu, \mathbf{I}_d)$  for an unknown  $k$ -sparse  $\mu \in \mathbb{R}^d$ . Then there is a randomized algorithm  $\mathcal{A}$  that takes as input the corrupted set  $T$ , contamination rate  $\epsilon$ , sparsity  $k \in \mathbb{N}$ , and a parameter  $q \in \mathbb{N}$ , and produces an estimate  $\hat{\mu}$  such that*

- (Sample Complexity and Error) *If  $n \gtrsim (k^{2q} \log d)/\epsilon^{2q}$ , then  $\|\hat{\mu} - \mu\|_2 \lesssim \epsilon\sqrt{\log(1/\epsilon)}$  with probability at least 0.9 over the randomness of*

<sup>3</sup>In fact, the overall runtime of these algorithms scales as  $\tilde{\Theta}(nd) = \tilde{\Theta}(d^2/\epsilon^2)$ , which is again quadratic in dimension.

the samples and the algorithm.<sup>4</sup>

- (Runtime) The algorithm runs in time at most  $d^{1.62+\frac{3}{4}} \text{poly}(\log(d), k^q, 1/\epsilon^q)$ .

Several remarks are in order. The error guarantee of Theorem 1.5,  $O(\epsilon\sqrt{\log(1/\epsilon)})$ , is nearly optimal even given infinite data and runtime.<sup>5</sup> The main contribution of Theorem 1.5 is the first algorithm for robust sparse mean estimation with runtime  $d^{2-\Omega(1)} \text{poly}(k/\epsilon)$  and sample complexity  $\text{poly}(k, \log d, 1/\epsilon)$  (by selecting  $q \in \mathbb{N}$  to be a constant bigger than 9), thereby affirmatively answering Question 1.4. As  $q$  increases, the dependence of the runtime on the dimension approaches  $d^{1.62}$ .<sup>6</sup> In particular, for a constant contamination rate  $\epsilon$ , we may set  $q$  as large as  $\gamma \left(\frac{\log d}{\log k}\right)$  for a small  $\gamma \in (0, 1)$ , and the algorithm retains sublinear (in  $d$ ) sample complexity  $d^{O(\gamma)}$  and subquadratic runtime  $d^{1.62+O(\gamma)} k^{O(1/\gamma)}$ . Finally, the sample complexity of Theorem 1.5 is (polynomially) larger than existing works; see Section 4 for further remarks.

We next focus on robust sparse principal component analysis (PCA), where given corrupted samples from  $\mathcal{N}(0, \mathbf{I}_d + \rho vv^\top)$  for a  $k$ -sparse unit vector  $v \in \mathbb{R}^d$ , the goal is to estimate the unknown  $v$ .

**Theorem 1.6 (Robust Sparse PCA in Subquadratic Time).** *Let  $T$  be an  $\epsilon$ -corrupted set of samples from  $\mathcal{N}(0, \mathbf{I}_d + \eta vv^\top)$  for  $\eta \in (0, 1)$  and a  $k$ -sparse unit vector. There is a randomized algorithm that takes as input the corrupted set  $T$ , contamination rate  $\epsilon$ , sparsity  $k \in \mathbb{N}$ , spike  $\eta$ , and a parameter  $q \in \mathbb{N}$ , and produces an estimate  $\hat{v}$  such that*

- (Sample Complexity and Error) If  $n \gtrsim \text{poly}((k^q \log d)/\epsilon^q)$ , then  $\|\hat{v}\hat{v}^\top - vv^\top\|_{\text{Fr}} \lesssim \sqrt{\epsilon \log(1/\epsilon)}/\eta$  with probability at least  $1 - \frac{1}{\text{poly}(d)}$  over the randomness of the samples and the algorithm.
- (Runtime) The algorithm runs in time at most  $d^{1.62+\frac{3}{4}} \text{poly}(n)$ .

This result gives the first subquadratic time algorithm for dimension-independent error, improving on the  $\Omega(d^2)$  runtime of Balakrishnan et al. (2017); Diakonikolas et al. (2019); Cheng et al. (2022). We note that the error guarantee of Theorem 1.6 is sub-optimal by a polynomial factor

<sup>4</sup>The success probability can be boosted to  $1 - \delta$  with a multiplicative increase of  $\log(1/\delta)$  in the sample complexity and the runtime by repeating the procedure.

<sup>5</sup>The information-theoretic optimal error is  $\Theta(\epsilon)$ . Moreover, it is computationally hard to beat  $\Theta(\epsilon\sqrt{\log(1/\epsilon)})$  in the statistical query lower model (Diakonikolas et al., 2017) and the low-degree polynomial tests (Brennan et al., 2021) under Definition 1.1.

<sup>6</sup>The constant 1.62 in the exponent comes from the subroutine for fast correlation detection (Valiant, 2015). See Section 1.3 for further details.

of  $\epsilon/\eta$  (like in Cheng et al. (2022)), since the information-theoretic optimal error is  $\epsilon/\eta$ . Despite this (polynomially) larger error, Theorem 1.6 is the first subquadratic time algorithm for robust sparse PCA with any non-trivial error, say, less than 0.01.

Our main technical ingredient in proving Theorems 1.5 and 1.6 is a result on detecting correlated coordinates of a high-dimensional distribution by Valiant (2015). We give an overview of Theorems 1.5 and 1.6 in Section 1.2, with details in Section 3 and Appendix D, respectively.

## 1.2. Overview of Techniques

We begin by presenting a brief overview of the landscape of current robust sparse mean estimation algorithms, followed by challenges in using these approaches to obtain an  $o(d^2)$  runtime, and then conclude by presenting our algorithm.

**(Dense) Robust Mean Estimation** Let  $\mu'$  and  $\Sigma'$  be the sample mean and the sample covariance of the current (corrupted) dataset. The general guiding principle in robust dense mean estimation is to use  $\Sigma'$  to check if there are harmful outliers and iteratively remove them. Recall that inliers are sampled from an isotropic covariance distribution  $\mathcal{N}(\mu, \mathbf{I}_d)$ . Thus, if we take  $\Theta(d/\epsilon^2)$  samples, then the variance of the inliers in any direction is  $(1 \pm \tilde{O}(\epsilon))$ . Moreover, the variance of any  $(1 - \epsilon)$  fraction of inliers is  $(1 \pm \tilde{O}(\epsilon))$ .

The following are the key insights in developing algorithms for robust dense mean estimation: (i) Outliers cannot change the sample mean  $\mu'$  of the data in any direction  $v$  without significantly increasing the covariance  $\Sigma'$  in the direction  $v$ , (ii) Given a direction of large variance  $v$  of the data (i.e., with variance larger than  $1 + \tilde{\Omega}(\epsilon)$ ), one can reliably remove outliers by projecting the data onto  $v$  and thresholding appropriately, and (iii) In the dense setting, the directions of large variance correspond to leading eigenvectors of the covariance matrix  $\Sigma'$ , and further they can be computed efficiently (in nearly-linear time) using power iteration. Thus, one can iteratively remove outliers as follows: compute (approximately) the leading eigenvalues and eigenvectors of the sample covariance matrix  $\Sigma'$  and remove the samples that have large projections along the computed direction.

### Adapting to Sparsity and Smaller Sample Complexity

For robust sparse mean estimation, one can adapt the above strategy by focusing only on the sparse directions  $v$ . Indeed, (i) and (ii) above are straight-forward and the resulting sample complexity of the algorithm is  $k \log(d)/\epsilon^2$  since we require concentration of the mean and the covariance only along  $k$ -sparse directions. However, the problem of computing the direction of the leading sparse eigenvalues of a matrix,  $\max_{v: \|v\|_2=1, k\text{-sparse}} v^\top \Sigma' v$ , is computationally-hard in the worst case. Inspired by the literature on sparse

PCA (d’Aspremont et al., 2007), Balakrishnan et al. (2017) proposed the following convex relaxation<sup>7</sup>:

$$\sup_{\{\mathbf{A}:\mathbf{A}\succeq 0, \text{tr}(\mathbf{A})=1, \|\mathbf{A}\|_1 \leq k\}} \langle \mathbf{A}, \Sigma' - \mathbf{I}_d \rangle. \quad (1)$$

Given such a feasible  $\mathbf{A}$  with value larger than  $\tilde{\Omega}(\epsilon)$ , one can remove outliers provided a larger sample complexity of  $(k^2 \log d)/\epsilon^2$ .<sup>8</sup> Although the resulting algorithm is polynomial-time and the desired sample complexity  $\text{poly}(k, \log d, 1/\epsilon)$ , the algorithm requires solving semidefinite programs (SDPs), for which the current algorithms require time superquadratic in dimension (Jiang et al., 2020).

**Spectral Algorithm of Diakonikolas et al. (2019)** To avoid solving the SDPs from the preceding paragraph, Diakonikolas et al. (2019) considers a different (and, in a sense, weaker) relaxation of sparse eigenvalues by relying on the distributional properties of Gaussians. Let  $\mathcal{B}_{k^2\text{-sparse}}$  be the set of all (sparse) matrices  $\mathbf{B}$  with Frobenius norm 1 and at most  $k^2$  non-zero entries. Importantly,  $\mathcal{B}_{k^2\text{-sparse}}$  contains all  $vv^\top$  for all  $k$ -sparse unit vectors  $v$ . Their starting point is the observation is that for any  $\mathbf{B} \in \mathcal{B}_{k^2\text{-sparse}}$ , we have  $\text{Var}_{x \sim \mathcal{N}(0, \mathbf{I}_d)}(x^\top \mathbf{B} x) = 2\|\mathbf{B}\|_{\text{Fr}}^2 = 2$ . Thus, for a fixed  $\mathbf{B} \in \mathcal{B}_{k^2\text{-sparse}}$ , the empirical mean of  $x^\top \mathbf{B} x$  over any  $1 - \epsilon$  fraction of inliers should be  $\text{tr}(\mathbf{B}) \pm \tilde{O}(\epsilon)$ . Moreover, standard uniform concentration arguments imply that this holds uniformly over  $\mathbf{B} \in \mathcal{B}_{k^2\text{-sparse}}$  given  $(k^2 \log d)/\epsilon^2$  samples. Their key observation is that the resulting (non-convex) optimization problem

$$\max_{\mathbf{B} \in \mathcal{B}_{k^2\text{-sparse}}} \langle \mathbf{B}, \Sigma' - \mathbf{I}_d \rangle \quad (2)$$

can be solved via standard matrix operations (despite being non-convex) without resorting to SDPs (as opposed to (1)); indeed, the optimal  $\mathbf{B}$  corresponds to the top- $k^2$  values of  $\Sigma' - \mathbf{I}_d$  in magnitude, computable in  $\tilde{O}(d^2)$  time given  $\Sigma'$ .

Given such a feasible  $\mathbf{B} \in \mathcal{B}_{k^2\text{-sparse}}$  that achieves  $\langle \mathbf{B}, \Sigma' - \mathbf{I}_d \rangle \geq \tilde{\Omega}(\epsilon)$ , Diakonikolas et al. (2019) also propose an efficient outlier-removal strategy. Overall, this yields an algorithm with runtime  $d^2 \text{poly}(k/\epsilon)$  and sample complexity  $(k^2 \log d)/\epsilon^2$ . While a significant improvement over Balakrishnan et al. (2017), this still unfortunately falls short of our target runtime of  $o(d^2)$ .

The main challenge in extending Diakonikolas et al. (2019)’s algorithm to get an  $o(d^2)$  runtime is that one needs to write down  $\Sigma'$  explicitly, which itself takes  $\Omega(d^2)$  time. Moreover, there is no known analog of power iteration for sparse

settings with provable guarantees (recall that in the dense setting, the power iteration can be implemented in nearly-linear runtime (Sachdeva & Vishnoi, 2014)). Our main technical insight is to use advances in fast algorithms for correlation detection initiated by Valiant (2015); see Section 2.3 for precise statements. Next, we explain why Valiant (2015) is potentially useful in our setting, explain the challenges in a direction application of their result, and our proposed fix.

### Fast Correlation Detection Algorithm To The Rescue

Denote the correlation detection algorithm in Valiant (2015) by  $\mathcal{A}_{\text{corr}}$ . In our setting, this algorithm guarantees that given a  $\rho \in (0, 1)$  and a large  $q \in \mathbb{N}$ , it can find (off-diagonal) coordinate pairs  $(i, j)$  such that  $|\Sigma'_{i,j}| \geq \rho$  in subquadratic time *as long as* there are at most  $\tilde{O}(d)$  off-diagonal coordinate pairs  $(i', j')$  such that  $|\Sigma'_{i',j'}| \geq \rho^q$ ; observe that  $\rho^q \ll \rho$ .

We now illustrate why such a subroutine may be useful. Suppose that  $\mathcal{A}_{\text{corr}}$  returns many correlated coordinate pairs. Then the optimal value in (2) must be large (if we optimize over  $\mathcal{B}_{k'\text{-sparse}}$  for some  $k'$  large enough<sup>9</sup>), and we can use those coordinates pairs to construct a  $k'$ -sparse  $\mathbf{B}$  that can be used to remove outliers as before (Diakonikolas et al., 2019). If, on the other hand,  $\mathcal{A}_{\text{corr}}$  returns only a few coordinate pairs, then we know that only these small set of coordinates are (potentially) corrupted, and we have reduced our problem to robust dense mean estimation on these coordinates; recall that in this setting, the sample mean is a good candidate for the remaining coordinates. Thus, a fast correlation detection algorithm leads to a subquadratic time algorithm to filter outliers (or declare victory) provided that there are only  $\tilde{O}(d)$  coordinate pairs with correlation larger than  $\rho^q$ .

### Challenges in Applying Fast Correlation Detection and A Proposed Fix

A priori, it is unclear why there must be only  $\tilde{O}(d)$  correlated coordinate pairs: indeed, the outliers are allowed to be dense (similar to inliers — recall that only the population mean is sparse), and, in the worst case, it is possible that they cause *all* coordinate pairs to be correlated (on the corrupted data). Thus, we need an alternative procedure to (i) detect if there are many  $\rho^q$ -correlated pairs and (ii) if so, find an alternative procedure to make progress.

Fortunately, it turns out that if there are  $\Omega(d)$  correlated pairs, then a random pair has  $\Omega(d/d^2) = \Omega(1/d)$  probability to be correlated. Hence, we can sample a relatively large — but subquadratic, say  $\Theta(d^{1.5})$  — number of random pairs to estimate the true count of correlated pairs. If random sampling does not find many such pairs, then the true count would anyway have been small with high probability, and we may safely invoke Valiant (2015)’s algorithm,

<sup>7</sup>For a matrix  $\mathbf{A} \in \mathbb{R}^{d \times d}$ ,  $\|\mathbf{A}\|_1$  denotes the  $\ell_1$ -norm of  $A$  when flattened as a  $d^2$ -dimensional vector

<sup>8</sup>This larger sample complexity,  $k^2$  versus  $k$ , is due to the stronger concentration required by the relaxation.

<sup>9</sup>Recall that we can interpret the maximum in (2) as Euclidean norm of the largest  $k^2$  entries, which would be at least  $\rho\sqrt{k^2} = \rho k$ .

solving the detection problem (i) above. On the other hand, if we do observe many (i.e., scaling polynomially with  $d$ )  $\rho^q$ -correlated pairs, then we know that the Frobenius norm of the largest  $k' = \text{poly}(k)$  entries of  $\Sigma' - \mathbf{I}_d$  must be large enough,  $\Omega(\rho^q \sqrt{k'})$ . In other words, we can find a relatively sparse  $\mathbf{B}' \in \mathcal{B}_{k'\text{-sparse}}$  such that  $\langle \Sigma' - \mathbf{I}_d, \mathbf{B}' \rangle \geq \tilde{\Omega}(\epsilon)$ . Thus, we can iteratively remove outliers (or declare victory when safe to do so) irrespective of the number of correlated pairs. Finally, the larger sample complexity of algorithm comes from invoking the filter on  $k'$ -sparse matrices  $\mathbf{B}$  for  $k' \gg k$ . We give a more detailed overview in Section 3.

**Robust Sparse PCA** We now give a very brief description of Theorem 1.6.

As a starting point, we remark that Theorem 1.6 does not follow directly from Theorem 1.5. In particular, a standard reduction from (robust sparse) PCA to (robust sparse) mean estimation relies on the fact that the mean of the random variable  $\mathbf{Y} - \mathbf{I}_d$  is exactly  $\rho v v^\top$  for  $\mathbf{Y} = x x^\top$ , where  $x \sim \mathcal{N}(0, \mathbf{I} + \rho v v^\top)$ , i.e., (robust) sparse PCA reduces to (robust) sparse mean estimation. However,  $\mathbf{Y}$  is a  $d^2$ -dimensional object and thus a naive application of Theorem 1.5 would yield a super-quadratic runtime, in fact,  $(d^2)^{1.62} = \Omega(d^3)$ . Moreover, the covariance of  $\mathbf{Y}$  is not isotropic and thus it is unclear if samples from  $\mathbf{Y}$  would satisfy the stability condition from Definition 2.3.

We show in Appendix D that a more white-box analysis of the ideas presented above can yield a subquadratic runtime: these ideas allow us to identify a small set of coordinates  $H$  (of size  $\text{poly}(k/\epsilon)$ ) such that the covariance matrix restricted to  $H$  is far from the identity matrix in Frobenius norm. Since the true covariance matrix of the inliers is  $\mathbf{I}_d + \rho v v^\top$  for a  $k$ -sparse vector  $v$ , the covariance matrix of the inliers restricted to  $H$  is far from identity if and only if  $v$  is approximately supported on  $H$ : If  $v$  is supported on  $H$ , we can simply estimate the top eigenvector of the covariance matrix (restricted to  $H$ ) and terminate the algorithm. On the other hand, if  $v$  is not supported on  $H$  (which can happen only if the outliers contribute disproportionately to the covariance matrix along  $H$ ), we show that an appropriate stability condition on inliers allows us to filter outliers.

### 1.3. Related Work

Our work is situated within the field of algorithmic robust statistics, and we refer the reader to Diakonikolas & Kane (2023) for an extensive exposition on the topic.

**Robust Sparse Estimation** Efficient algorithms for robust sparse mean estimation were first developed in Balakrishnan et al. (2017), giving an algorithm to compute  $\hat{\mu}$  with sample complexity  $\text{poly}(k, \log d, 1/\epsilon)$ , runtime  $\text{poly}(d, 1/\epsilon)$ , and near-optimal error  $\|\hat{\mu} - \mu\|_2 = \tilde{O}(\epsilon)$ . Their algorithm

used the ellipsoid algorithm with a separation oracle that requires solving an SDP. Invoking the ellipsoid algorithm can be avoided by using Zhu et al. (2022) or through iterative filtering, but the resulting algorithm still requires solving multiple SDPs, which, as noted earlier, is inherently slow. Bypassing the use of SDPs, Diakonikolas et al. (2019) developed the first spectral algorithm for robust sparse estimation with runtime  $d^2 \text{poly}(k, \log d, 1/\epsilon)$ . Another novel take on this problem was seen in Cheng et al. (2022), which gave an optimization-based algorithm showing that first-order stationary points of a natural non-convex objective suffice. Although the resulting algorithm relies on simple matrix operations, the derived runtime is super-quadratic in dimension. In a different direction, Diakonikolas et al. (2022a) and Diakonikolas et al. (2022c) developed algorithms for robust sparse mean estimation for a wider class of distributions: heavy-tailed distributions and light-tailed distributions with unknown covariance, respectively.

Robust sparse mean estimation is conjectured to have information-computation gaps (Diakonikolas et al., 2017; Brennan & Bresler, 2020; Diakonikolas et al., 2022a). Particularly, while there exist inefficient ( $d^k$ -time) algorithms for robust sparse mean estimation using  $(k \log(d))/\epsilon^2$  samples, all polynomial time algorithms are conjectured to require  $\Omega(k^2/\epsilon^2)$  samples (Brennan & Bresler, 2020).

**Fast Algorithms for Robust Estimation** Looking beyond polynomial runtime as the criterion of computational efficiency, a recent line of work has investigated *faster* algorithms for a variety of robust estimation tasks: mean estimation (Cheng et al., 2019a; Dong et al., 2019; Depersin & Lecué, 2022; Lei et al., 2020; Diakonikolas et al., 2022b;d; 2023a), covariance estimation (Cheng et al., 2019b), principal component analysis (Jambulapati et al., 2020; Diakonikolas et al., 2023b), list-decoding (Cherapanamjeri et al., 2020b; Diakonikolas et al., 2022b), and linear regression (Cherapanamjeri et al., 2020a; Diakonikolas et al., 2023a). The overarching goal in this line of work is to develop robust algorithms that have runtimes matching the corresponding non-robust off-the-shelf algorithms, thus reducing the computational overhead of robustness. However, none of these algorithms is tailored to sparsity and hence have sample complexity scaling (at least) linearly with the dimension.<sup>10</sup> *Our main contribution is the first subquadratic runtime algorithm for robust sparse mean estimation with sample complexity  $\text{poly}(k \log d, 1/\epsilon)$ .*

**Fast Correlation Detection** Given  $n$  vectors in  $\{\pm 1\}^d$  and two thresholds  $1 \geq \rho > \tau > 0$ , the correlation detection problem asks to find all coordinate pairs that have correlation at least  $\rho$  given that not too many pairs have

<sup>10</sup>As a result, the dependence on the runtime again becomes quadratic because  $nd = \Omega(d^2)$ .

correlation larger than  $\tau$ . This problem is a generalization of the light bulb problem (Valiant, 1988). The first subquadratic algorithm for both these problems was given by Valiant (2015), with further improvements and simplifications in Karppa et al. (2018; 2020); Alman (2019). Developing faster algorithms for these problems is an active area of research (Alman & Zhang, 2023), and further algorithmic improvements would likely also improve our runtime guarantees in Theorems 1.5 and 1.6.

## 2. Preliminaries

**Notation** For a random variable  $X$ ,  $\mathbb{E}[X]$  denotes its expectation. For a finite set  $S$  and a function  $g : S \rightarrow \mathbb{R}^d$ , we use  $\mathbb{E}_S[g(X)]$  to denote  $(\sum_{x \in S} g(X))/|S|$ . We use  $\text{poly}(\cdots)$  to denote an expression that is polynomial in its arguments. The notations  $\lesssim, \gtrsim, \asymp$  hide absolute constants.

For a vector  $x \in \mathbb{R}^d$ , we use  $\|x\|_0$  and  $\|x\|_2$  to denote the number of non-zero entries of  $x$  and the  $\ell_2$ -norm of  $x$ , respectively. For a vector  $x$  and  $k \in \mathbb{N}$ , we define the  $\|x\|_{2,k}$  norm as the maximum correlation between  $x$  and a unit  $k$ -sparse vector, i.e.,  $\|x\|_{2,k} := \sup_{v: \|v\|_0 \leq k} \langle v, x \rangle$ . Estimation in  $\|\cdot\|_{2,k}$  immediately yields an estimate that is close in  $\ell_2$  norm whenever  $\mu$  is  $k$ -sparse:

**Proposition 2.1** (Sparse estimation using  $\|\cdot\|_{2,k}$  norm (Cheng et al., 2022)). *Let  $x \in \mathbb{R}^d, y \in \mathbb{R}^d$ , where  $y$  is  $k$ -sparse. Let  $J \subset [d]$  be the top- $k$  coordinates of  $x$  in magnitude, breaking ties arbitrarily. Define  $x' \in \mathbb{R}^d$  to be  $x'_i = x_i$  if  $i \in J$  and 0 otherwise. Then  $\|x' - y\|_2 \leq 6\|x - y\|_{2,k}$ .*

Thus, in the sequel, we solve the harder problem of estimating an arbitrary mean  $\mu \in \mathbb{R}^d$  in the  $\|\cdot\|_{2,k}$  norm.

We denote matrices by bold capital letters, e.g.,  $\mathbf{A}, \mathbf{\Sigma}$ . We denote the  $d \times d$  identity matrix by  $\mathbf{I}_d$ , omitting the subscript when clear. For a matrix  $\mathbf{A}$ , we use  $\|\mathbf{A}\|_0$ , and  $\|\mathbf{A}\|_{\text{Fr}}$ , to denote the number of non-zero entries and the Frobenius norm, respectively. For matrices  $\mathbf{A}$  and  $\mathbf{B}$  of the same dimensions,  $\langle \mathbf{A}, \mathbf{B} \rangle$  denotes the trace inner product  $\text{tr}(\mathbf{A}^\top \mathbf{B})$ .

For a subset  $H \subset [d]$ , and a vector  $x \in \mathbb{R}^d$ , define  $(x)_H$  to be  $|H|$ -dimensional vector that corresponds to the coordinates in  $H$ . Similarly, for a matrix  $\mathbf{A}$ , we define  $(\mathbf{A})_H$  to be the  $|H| \times |H|$  matrix corresponding to coordinates in  $H$ . For a square matrix  $\mathbf{A}$ , we use  $\text{diag}(\mathbf{A})$  and  $\text{offdiag}(\mathbf{A})$  to denote its diagonal and offdiagonal, respectively.

For a finite set  $T \subset \mathbb{R}^d$ , we define  $\mu_T$  and  $\mathbf{\Sigma}_T$  to be the sample mean and the sample covariance of  $T$ , respectively.<sup>11</sup> When the set  $T$  is clear from context, for a coordinate pair  $(i, j) \in [d] \times [d]$  with  $i \neq j$ , we denote the correlation between these coordinates on  $T$  as  $\text{corr}(i, j) := \left| \frac{\mathbf{\Sigma}'_{i,j}}{\sqrt{\mathbf{\Sigma}'_{i,i} \mathbf{\Sigma}'_{j,j}}} \right|$  for  $\mathbf{\Sigma}' = \mathbf{\Sigma}_T$ . For a

$\rho \in (0, 1)$ , we say coordinates  $(i, j)$  are  $\rho$ -correlated if  $\text{corr}(i, j) \geq \rho$ .

Robust sparse estimation requires checking whether the current covariance matrix  $\mathbf{\Sigma}'$  has small quadratic forms,  $v^\top (\mathbf{\Sigma}' - \mathbf{I})v$ , for sparse unit vectors. For a matrix  $\mathbf{A}$  and  $k \in \mathbb{N}$ , we define the sparse operator norm,  $\|\mathbf{A}\|_{\text{op},k} := \sup_{v: \|v\|_2=1, \|v\|_0 \leq k} |v^\top \mathbf{A}v|$ . Since computing  $\|\cdot\|_{\text{op},k}$  is computationally hard, we look at the following relaxation from Diakonikolas et al. (2019): For a matrix  $\mathbf{A}$ , define  $\|\mathbf{A}\|_{\text{Fr},k^2} := \sup_{\mathbf{B}: \|\mathbf{B}\|_{\text{Fr}}=1, \|\mathbf{B}\|_0 \leq k^2} \langle \mathbf{A}, \mathbf{B} \rangle$ . It can be seen that  $\|\mathbf{A}\|_{\text{op},k} \leq \|\mathbf{A}\|_{\text{Fr},k^2}$  since  $\mathbf{B}$  could be all  $\pm v v^\top$  for  $k$ -sparse unit vectors  $v$ . Moreover,  $\|\mathbf{A}\|_{\text{Fr},k^2}$  is the Euclidean norm of the largest  $k^2$  entries (in magnitude) of  $\mathbf{A}$ .

Since we will routinely look at the projections of the points on a subset of coordinates, we formally define it below:

**Definition 2.2** (Projection of Pairs of Coordinates). Let  $H_{\text{pair}} \subset [d] \times [d]$  be a set of pair of coordinates such that  $(i, i) \notin H_{\text{pair}}$  for any  $i \in [d]$ . For an even  $m \in [d^2]$ , we define the operator  $\text{Proj}_m$  that takes any such  $H_{\text{pair}}$  and returns a set in  $[d]$  that has a cardinality at most  $m$  as follows:

- If  $|H_{\text{pair}}| \leq \frac{m}{2}$ , return  $\{i : (i, j) \in H_{\text{pair}} \text{ or } (j, i) \in H_{\text{pair}}\}$ .
- Otherwise, let any  $m/2$  distinct elements of  $H_{\text{pair}}$  be  $(i_1, i_2), \dots, (i_{m-1}, i_m)$ . Return  $\{i_j : j \in [m]\}$ .

When the subscript  $m$  is omitted, we take  $m$  to be  $d^2$ .

Informally, the operator returns a set  $H$  such that for any matrix  $\mathbf{A}$ , for small  $m$ ,  $\|(\mathbf{A})_H\|_{\text{Fr}}^2 \geq \sum_{(i,j) \in H_{\text{pair}}} \mathbf{A}_{i,j}^2$ , while for larger  $m$ ,  $\|(\mathbf{A})_H\|_{\text{Fr}}^2 \geq m \min_{(i,j) \in H_{\text{pair}}} \mathbf{A}_{i,j}^2$ .

### 2.1. Deterministic Condition on Inliers

A recurring notion in developing robust algorithms is that of *stability*, which stipulates that the first and second moment of the data not change much under removal of a small fraction of data points.

**Definition 2.3** (Stability). For an  $\epsilon \in (0, 1/2)$ ,  $\delta \geq \epsilon$ , and sparsity  $k \in \mathbb{N}$ , we say a set  $S \subset \mathbb{R}^d$  is  $(\epsilon, \delta, k)$ -stable with respect to  $\mu \in \mathbb{R}^d$  if the following holds for any subset  $S' \subseteq S$  with  $|S'| \geq (1 - \epsilon)|S|$ : (i)  $\|\mathbb{E}_{S'}[X - \mu]\|_{2,k} \leq \delta$ , and (ii)  $\|\mathbb{E}_{S'}[(X - \mu)(X - \mu)^\top] - \mathbf{I}_d\|_{\text{Fr},k^2} \leq \delta^2/\epsilon$ .

The following result gives a nearly-tight bound on the sample complexity required to ensure stability.

**Lemma 2.4** (Stability Sample Complexity (Cheng et al., 2022, Lemma 3.3)). *Let  $S$  be a set of  $n$  i.i.d. samples from a subgaussian distribution  $P$  over  $\mathbb{R}^d$  such that  $P$  has (i) mean  $\mu \in \mathbb{R}^d$ , (ii) identity covariance, and (iii) satisfies the Hanson-Wright inequality; In particular,  $\mathcal{N}(\mu, \mathbf{I}_d)$  satisfies all three properties. Then if  $n \gtrsim (k^2(\log d)/\epsilon^2)$ , then  $S$  is*

<sup>11</sup>Not to be confused with  $(\mathbf{\Sigma})_H$  when  $H \subset [d]$ .

$(\epsilon, \delta, k)$ -stable with high probability with respect to  $\mu$  for  $\delta = C\epsilon\sqrt{\log(1/\epsilon)}$  where  $C$  is a large absolute constant.

We note that the deterministic condition in Definition 2.3 is slightly stronger than Cheng et al. (2022)—the condition for the covariance—but their proof continues to work with the same sample complexity for Definition 2.3.

## 2.2. Certificate Lemma and Frobenius Norm Filtering

The following standard certificate lemma guides the algorithmic design: if the sample covariance matrix of the corrupted data is roughly isotropic, then the sample mean is a good estimate.

**Lemma 2.5** (Sparse Certificate Lemma, see, e.g., Balakrishnan et al. (2017)). *Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \delta, k)$ -stable with respect to  $\mu$  (cf. Definition 2.3). Then  $\|\mu_T - \mu\|_{2,k} \lesssim \delta + \sqrt{\epsilon\|\Sigma_T - \mathbf{I}_d\|_{\text{op},k}}$ .*

We now state the guarantee of filtering procedures, where the goal is to filter outliers from a corrupted set  $T$ . In dense mean estimation, the most common filters are based on the scores of the form  $(v^\top(x - \mu_T))^2$  for a direction of large variance  $v$ ; this filter is guaranteed to succeed as long as the covariance matrix  $\Sigma_T$  is far from the identity in operator norm. In our setting, we will need a stronger filter guarantee that is guaranteed to succeed under the weaker condition that the covariance (submatrix) matrix is far from the (submatrix) identity in the Frobenius norm; Observe that the operator norm of (the corresponding submatrix of)  $\Sigma_T - \mathbf{I}_d$  can be much smaller. The following lemma corresponds to the above situation and simplifies Diakonikolas et al. (2019, Steps 6-10 of Algorithm 1):

**Lemma 2.6** (Sparse Filtering Lemma). *Let  $\epsilon \in (0, \epsilon_0)$  for a small absolute constant  $\epsilon_0$ . Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \delta, k)$ -stable with respect to  $\mu$ . Let  $H \subset [d]$  be such  $\|(\Sigma_T - \mathbf{I})_H\|_{\text{Fr}} = \lambda$  for  $\lambda \gtrsim \delta^2/\epsilon$  and  $|H| \leq k$ . There exists an algorithm  $\mathcal{A}$  that takes  $T$ ,  $H$ ,  $\epsilon$ , and  $\delta$  and returns scores  $f : T \rightarrow \mathbb{R}_+$  so that  $\sum_{x \in S \cap T} f(x) \leq \sum_{x \in T \setminus S} f(x)$ , i.e., the sum of scores over inliers is less than that of outliers, and  $\max_{x \in T} f(x) > 0$ . Moreover, the algorithm runs in time  $d \cdot \text{poly}(|H||T|)$ .*

These scores can be used to filter points from  $T$  such that on expectation over the algorithm’s randomness, more outliers are removed than inliers (Diakonikolas & Kane, 2023). We give a proof of Lemma 2.6 in Appendix A.

Thus, if  $\Sigma_T - \mathbf{I}_d$  has large (sparse) Frobenius norm, then we can make progress by removing outliers. The contribution to this norm from the *diagonal* entries can be calculated efficiently in  $O(dn)$  time, and if large, then can also be used to remove outliers. Thus, we will assume that the corrupted set has already been pre-processed to satisfy the following:

**Condition 2.7** (Preprocessing). Let  $T$  be an  $\epsilon$ -corrupted

version of  $S$ , where  $S$  is  $(\epsilon, \delta, k)$ -stable. Suppose  $T$  satisfies  $\|\text{diag}(\Sigma_T - \mathbf{I}_d)\|_{\text{Fr},k^2} \leq \min(O(\delta^2/\epsilon), 0.5)$ .

For completeness, we give details in Appendix A.3.

## 2.3. Detecting Correlation in Subquadratic Time

We will use Valiant (2015, Theorem 1) that can detect  $\rho$ -correlated coordinates in subquadratic time if there are not too many  $\tau$ -correlated coordinates for  $\tau \ll \rho$ , say  $\rho^3$ . See Appendix A.4 for details:

**Theorem 2.8** (Fast Correlation Detection (Valiant, 2015)). *Let  $\rho \in (0, 1)$  be strong correlation threshold and  $\tau \in (0, 1)$  be margin threshold with  $\rho > 12\tau$ . Let  $T$  be a set of  $n$  vectors in  $\mathbb{R}^d$  such that there are at most  $s$   $\tau$ -correlated coordinate pairs. Then, there is an algorithm that takes  $\rho, \tau, T$  as input, and, with probability  $1 - o(1)$ , will output all  $\rho$ -coordinate pairs. Additionally, the runtime of the algorithm is  $(sd^{0.62} + d^{1.62+2.4\frac{\log(4/\rho)}{\log(1/3\tau)}}) \text{poly}(n, \log d, 1/\tau)$ .*

## 3. Sparse Certificates in Subquadratic Time

In this section, we explain our main technical contribution: a fast algorithm for robust sparse mean estimation under the stability condition.

**Theorem 3.1** (Robust Sparse Mean Estimation in Subquadratic Time (informal)). *Let the corruption rate be  $\epsilon \in (0, \epsilon_0)$ , where  $\epsilon_0$  is a sufficiently small absolute constant. Let  $k \in \mathbb{N}$  be the sparsity parameter and  $q \in \mathbb{N}$  the correlation decay parameter with  $q \geq 3$ . Let  $T$  be an  $\epsilon$ -corrupted version of a set  $S$ , where  $S$  satisfies  $(\epsilon, \delta, k')$ -stability with respect to  $\mu$  for  $k' = \Theta\left(\frac{(k)^q}{(\delta^2/\epsilon)^{q-1}}\right)$ . Then there is a randomized algorithm that takes as inputs  $T$ ,  $\epsilon$ ,  $\delta$ ,  $k$ , and  $q$  and produces an estimate  $\hat{\mu}$  such that with high probability:*

- ▶ (Error)  $\|\hat{\mu} - \mu\|_{2,k} \lesssim \delta$ .
- ▶ (Runtime) The algorithm runs in time at most  $d^{1.62+\frac{3}{q}} \text{poly}(|T|, \log(d), k^q, 1/\epsilon^q)$ .

The result above implies Theorem 1.5 using the sample complexity of stability from Lemma 2.4 and Proposition 2.1.

To establish Theorem 3.1, we start with the following blueprint for robust sparse mean estimation, with the aim of implementing it in  $o(d^2)$  time.

---

### Algorithm 1 Algorithmic Blueprint

---

- 1: Compute  $\|\Sigma_T - \mathbf{I}_d\|_{\text{Fr},k^2}$  approximately.
  - 2: **while**  $\|\Sigma_T - \mathbf{I}_d\|_{\text{Fr},k^2}$  is large **do**
  - 3:     Let  $H$  be the corresponding coordinates with large Frobenius norm.
  - 4:     Filter  $T$  using  $H$  in Lemma 2.6.
  - 5: **end while**
  - 6: Output the sample mean  $\mu_T$ .
-

The problem in implementing this blueprint naively is that Steps 1 and 3 in Algorithm 1 take  $d^2$  time. However, these are the only two bottlenecks: Each filtering step takes only  $d \text{poly}(k, n) = d \text{poly}(k/\epsilon)$  time, and there are at most  $n = \text{poly}(k/\epsilon)$  iterations. As we describe below, our goal in this section is to use Theorem 2.8 to speed up the Steps 1 and 3.

**Usefulness of Fast Correlation Detection** We will run Theorem 2.8 to identify the off-diagonal indices  $(i, j) \in [d] \times [d]$  that are correlated. Recall that Theorem 2.8 takes two arguments  $\rho$ —the threshold for strong correlation—and  $\tau$ —the margin threshold. Suppose we fix  $\rho$  to be small, roughly  $\frac{\delta^2}{\epsilon \text{poly}(k)}$ . Then Theorem 2.8 reports back all  $\rho$ -correlated coordinate pairs (The time taken is subquadratic in  $d$  only if the number of  $\tau$ -correlated pairs is small.) Let  $H_{\text{pair}} \subset [d] \times [d]$  be the set of coordinate pairs in the output, and define  $H$  to be the set of all coordinates that appear in  $H_{\text{pair}}$ ;  $\text{Proj}(H_{\text{pair}})$  from Definition 2.2 to be more formal. Then, one of the following two cases must be true:

- $|H|$  is small ( $\text{poly}(k/\epsilon)$ ): Since coordinates in  $H^{\text{G}}$  have correlations at most  $\rho$ , then we know that  $\|(\Sigma_T - \mathbf{I})_{H^{\text{G}}}\|_{\text{Fr}, k^2} \leq k\rho$ , which can be made less than  $\delta^2/\epsilon$  for  $\rho$  small enough ( $\rho \lesssim \delta^2/(k\epsilon)$ ). Thus, the sample mean on  $H^{\text{G}}$  is a good estimate in  $\ell_{2,k}$  norm (cf. Lemma 2.5). On  $H$ , we can use a dense mean estimation algorithm, which would be fast as  $|H| = \text{poly}(k/\epsilon)$ .
- $|H|$  is large: Since there are many coordinate pairs with correlation at least  $\rho$ , we can filter and iterate as follows: If we take any  $H' \subset H$  of size  $k'$ , then each row and column in  $(\Sigma_T - \mathbf{I})_{H'}$  has an entry larger than  $\rho$  (in absolute value), and thus  $\|(\Sigma_T - \mathbf{I})_{H'}\|_{\text{Fr}} \geq \rho\sqrt{k'}$ . By taking  $k'$  large enough (larger than  $\delta^4/(\rho^2\epsilon^2)$ ), the resulting quantity will be bigger than  $\delta^2/\epsilon$ , allowing us to filter if the inliers satisfy stability with parameter  $k'$  (cf. Lemma 2.6).

Thus, we can implement Steps 1 and 3 in Algorithm 1 fast using Theorem 2.8, so long as (i) inliers satisfy  $(\epsilon, \delta, k')$  stability, (ii) we choose  $\rho^2 \asymp \frac{\delta^4}{\epsilon^2 k^2}$  and  $k' \asymp \frac{\delta^4}{\epsilon^2 \rho^2} \asymp k^2$ , and (iii) there are not too many  $\tau$ -correlated pairs.

**Challenges in Using Fast Correlation Detection** Suppose we set  $\tau = \rho^q$  for some  $q \geq 3$ .<sup>12</sup> Looking at Theorem 2.8, we obtain a subquadratic time algorithm *only* if  $s$ , the number of  $\tau$ -correlated pairs for  $\tau := \rho^q$ , is much smaller than  $d^{1.38}$ ; In fact, we will impose  $s$  to be less than  $d$  so that it is not the dominant factor in the runtime. A priori, there is no reason for there to be at most  $d$  coordinate pairs (out of  $d^2$  pairs) that are  $\tau$ -correlated. Thus, we need a way to detect this situation and find an alternative way to make progress.

<sup>12</sup>We choose this parameterization because the runtime of Theorem 2.8 depends on  $\log(1/\rho)/\log(1/\tau) = 1/q$ .

**Proposed Solutions: Efficient Detection and Filtering** We begin with the detection procedure. If there are  $\Omega(d)$  many  $\tau$ -correlated pairs, then a pair sampled uniformly at random has a probability of  $\Omega(d^{-1})$  of being  $\tau$ -correlated. Thus, if we check many random coordinate pairs, superlinear but subquadratic, then we can accurately guess  $s$ .

In particular, let  $U$  be the number of  $\tau$ -correlated pairs that were observed out of  $m$  random pairs (sampled with replacement). Then,  $U$  is distributed roughly as  $\text{Ber}(m, s/d^2)$ . Binomial concentration implies that  $s \lesssim (d^2 U/m) + (\log d)(d^2/m)$  with probability at least  $(1 - 1/d^2)$ . Taking  $m$  to be  $d^{1.5}$ , we see that  $s \lesssim \sqrt{d}U + \sqrt{d} \log d$ . Thus, we obtain a fast (randomized) check to see if  $s$  is less than  $d$  than runs in  $mn = d^{1.5} \text{poly}(k/\epsilon)$  time: simply check if  $U \leq \sqrt{d}$ . Thus, it remains to ensure that we can make progress when  $U$  is large, in particular,  $\Omega(\sqrt{d})$ .

Let  $H'_{\text{pair}}$  be the  $\tau$ -correlated coordinates that were observed in the above procedure;  $U = |H'_{\text{pair}}|$ . Crucially, we are in the regime when  $|H'_{\text{pair}}| \geq \sqrt{d}$ . We want to use a small subset of the coordinates in  $H'_{\text{pair}}$  to filter outliers. Let  $H$  be  $k''$ -sized set of coordinates that appear in  $H'_{\text{pair}}$ ; formally,  $H := \text{Proj}_{k''}(H'_{\text{pair}})$ . Thus, each row and column in  $(\Sigma_T - \mathbf{I})_H$  has an entry larger than  $\tau$  in absolute value, implying that  $\|(\Sigma_T - \mathbf{I})_H\|_{\text{Fr}} \geq \tau\sqrt{k''}$ . Therefore, we can use this  $H$  to filter using Lemma 2.6 as long as the original set is also  $(\epsilon, \delta, k'')$ -stable and  $\sqrt{k''}\tau \gg \delta^2/\epsilon$ , i.e.,  $k'' \asymp \frac{\delta^4}{\epsilon^2 \tau^2} \asymp \frac{\delta^4}{\epsilon^2 \rho^{2q}} \asymp \frac{\delta^4}{\epsilon^2 (\delta^4/k\epsilon^2)^q} \asymp \frac{k^{2q}}{(\delta^4/\epsilon^2)^{q-1}}$ .

Observe that we require the inliers to satisfy  $(\epsilon, \delta, k'')$ -stability for  $k'' = \text{poly}(k^q/\epsilon^q)$ . By Lemma 2.4, a set of  $(k'')^2/\epsilon^2$  many i.i.d. points will satisfy this stability condition, giving us the sample complexity.

### 3.1. Fast Filters Using Correlation Detection

We now give formal guarantees of the key procedures outlined above, with the proofs deferred to Appendix B. First, consider the procedure that randomly samples the coordinates to estimate the number of weakly-correlated coordinates.

**Lemma 3.2.** *Let  $T \subset \mathbb{R}^d$  be a multiset with covariance matrix  $\Sigma'$ . Let  $m \in N$  be the sampling parameter. Let  $J_* \subset [d] \times [d]$  be the off-diagonal coordinate pairs such that  $|\Sigma'_{i,j}| \geq \tau$ . Then Algorithm 2 takes  $\tau$ ,  $m$ , and  $T$  as input and returns a set  $J \subset J_*$  in time  $O(m|T|)$  such that with probability  $1 - 1/d^2$ ,  $|J| \geq \frac{m|J_*|}{4d^2} - 16 \log d$ .*

---

#### Algorithm 2 RANDOMLYCHECKCOORDINATES

---

- 1: Let  $H_{\text{pair}} \subset [d] \times [d]$  of size  $m$ , with  $(i, j) \in H_{\text{pair}}$  sampled i.i.d. from off-diagonal elements.
  - 2: Let  $J \leftarrow \{(i, j) \in H_{\text{pair}} : |\Sigma'_{i,j}| \geq \tau\}$  for  $\Sigma' = \Sigma_T$
  - 3: **return**  $J$ .
-

Combining Algorithm 2 with Theorem 2.8, Algorithm 3 either returns a small subset of coordinates with large Frobenius norm or indicates when all tiny subsets have small Frobenius norm.

---

**Algorithm 3** Subroutine to Identify Corrupted Coordinates
 

---

**Require:** a finite set  $T \subset \mathbb{R}^d$  such that diagonal entries of  $\Sigma_T$  lie in  $[1/2, 2]$ , Frobenius threshold  $\kappa \in \mathbb{R}_+$ , correlation threshold  $\rho \in (0, 1)$ , margin threshold  $\tau$ , sampling parameter  $m \in \mathbb{N}$ , and correlation count  $s \in \mathbb{N}$ . We require the parameters to satisfy  $\frac{c_1 d^2}{m} (\frac{\kappa^2}{\tau^2} + \log d) \leq s$ .

- 1: Let  $H'_{\text{pair}}$  be the output of Algorithm 2 with  $\tau, m$ , and  $T$
- 2: **if**  $|H'_{\text{pair}}| \geq \kappa^2/\tau^2$  **then**
- 3:     **return**  $\text{Proj}_{\kappa^2/\tau^2}(H'_{\text{pair}})$
- 4: **end if**
- 5: Let  $H_{\text{pair}} \subseteq [d] \times [d]$  be the output of Theorem 2.8 with  $\tau, \rho, s$  on  $T$
- 6: **if**  $|H_{\text{pair}}| > \kappa^2/\rho^2$  **then**
- 7:     **return**  $\text{Proj}_{\kappa^2/\rho^2}(H_{\text{pair}})$ .
- 8: **end if**
- 9: Let  $H_2 \leftarrow \text{Proj}(H_{\text{pair}})$
- 10: **if**  $\|(\text{offdiag}(\Sigma_S))_{H_2}\|_{\text{Fr}} \geq \kappa$  **then**
- 11:     **return**  $H_2$
- 12: **end if**
- 13: **return** “ $\perp$ ”

---

**Proposition 3.3** (Subroutine to Identify Corrupted Coordinates Algorithm 3). *Let corrupted set  $T$ , Frobenius threshold  $\kappa \in \mathbb{R}$ , correlation threshold  $\rho \in (0, 1)$ , margin threshold  $\tau \in (0, \rho/12)$ , sampling parameter  $m \in \mathbb{N}$ , and correlation count  $s \in \mathbb{N}$ . Suppose that each diagonal entry of  $\Sigma_T$  lies in  $[1/2, 3/2]$  and  $\frac{c_1 d^2}{m} (\frac{\kappa^2}{\tau^2} + \log d) \leq s$ .*

*Then Algorithm 3 takes as input  $T, \kappa, \rho, \tau, m, s$  as input and satisfies the following with probability  $1 - 1/\text{poly}(d)$ :*

- *It either outputs a set  $H \subset [d]$  with  $|H| \leq \frac{\kappa^2}{\tau^2}$  and  $\|(\text{offdiag}(\Sigma_T))_H\|_{\text{Fr}} \geq \kappa/2$ .*
- *Else, it outputs “ $\perp$ ”. If it outputs “ $\perp$ ”, then  $\|\text{offdiag}(\Sigma_T)\|_{\text{Fr}, k^2}$  is at most  $2\kappa + 2\rho k$ .*

*Moreover, the algorithm runs in  $O\left(m + sd^{0.62} + d^{1.62+3\frac{\log(4/\rho)}{\log(1/3\tau)}}\right)\text{poly}(n, \log d, 1/\tau)$  time.*

We briefly sketch the proof of Theorem 3.1 below.

*Proof sketch of Theorem 3.1.* We formalize the blueprint from Algorithm 1 with the filtering set  $H$  given by Proposition 3.3. Combined with Lemma D.5, a standard martingale argument yields Theorem C.1; see Appendix C.  $\square$

## 4. Discussion

In this article, we presented the first subquadratic time algorithm for robust sparse mean estimation. We now discuss some related open problems and avenues for improvement. First, the sample complexity of Theorem 1.5 is polynomially larger than  $k^2(\log d)/\epsilon^2$ —the sample complexity of existing (quadratic runtime) algorithms.<sup>13</sup> Bridging this gap is an important problem. Second, Theorem 1.5 is specific to isotropic structured distributions such as Gaussians (more broadly, isotropic distributions  $P$  with mean  $\mu$  that satisfy  $\text{Var}_{x \sim P}((x - \mu)^\top \mathbf{A}(x - \mu)) \lesssim \|\mathbf{A}\|_{\text{Fr}}^2$  for symmetric matrices  $\mathbf{A}$ ). Indeed, both Diakonikolas et al. (2019); Cheng et al. (2022) rely on the isotropy and the aforementioned variance structure of quadratic polynomials to avoid solving SDPs that appear in Balakrishnan et al. (2017). Developing a similarly fast algorithm for unstructured distributions (with sparse means) is still open to our knowledge. Third, because Theorem 1.5 relies on Valiant (2015), which in turn relies on fast matrix multiplication, the resulting algorithm may not offer practical benefits for moderate dimensions; see the discussion in Valiant (2015). We believe overcoming these limitations is an important practically-motivated question. Finally, Question 1.3 remains open.

## Acknowledgements

We are grateful to Sushrut Karmalkar, Jasper Lee, Thanasis Pittas, and Kevin Tian for discussions on robust sparse mean estimation. We also thank Ilias Diakonikolas and Daniel Kane for their many insights on robustness.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## References

- Alman, J. An illuminating algorithm for the light bulb problem. In *Proc. 2nd Symposium on Simplicity in Algorithms (SOSA)*, 2019.
- Alman, J. and Zhang, H. Generalizations of matrix multiplication can solve the light bulb problem. In *Proc. 64th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2023.
- Andrews, D. F., Bickel, P. J., Hampel, F. R., Huber, P. J., Rogers, W. H., and Tukey, J. W. *Robust Estimates of*

<sup>13</sup>The sample complexity of  $\tilde{\Theta}(k^2/\epsilon^2)$  is also conjectured to be near-optimal among computationally-efficient algorithms (Brennan & Bresler, 2020).

- Location: Survey and Advances*. Princeton University Press, Princeton, NJ, USA, 1972.
- Balakrishnan, S., Du, S. S., Li, J., and Singh, A. Computationally Efficient Robust Sparse Estimation in High Dimensions. In *Proc. 30th Annual Conference on Learning Theory (COLT)*, 2017.
- Brennan, M. and Bresler, G. Reducibility and statistical-computational gaps from secret leakage. In *Proc. 33rd Annual Conference on Learning Theory (COLT)*, 2020.
- Brennan, M., Bresler, G., Hopkins, S. B., Li, J., and Schramm, T. Statistical query algorithms and low-degree tests are almost equivalent. In *Proc. 34th Annual Conference on Learning Theory (COLT)*, 2021.
- Cheng, Y. High-dimensional robust statistics: Faster algorithms and optimization landscape. **Robustness in High-dimensional Statistics and Machine Learning at IDEAL Institute**, 2021. See timestamp 22:00 in the talk.
- Cheng, Y., Diakonikolas, I., and Ge, R. High-dimensional robust mean estimation in nearly-linear time. In *Proc. 30th Annual Symposium on Discrete Algorithms (SODA)*, 2019a. doi: 10.1137/1.9781611975482.171.
- Cheng, Y., Diakonikolas, I., Ge, R., and Woodruff, D. P. Faster algorithms for high-dimensional robust covariance estimation. In *Proc. 32nd Annual Conference on Learning Theory (COLT)*, 2019b.
- Cheng, Y., Diakonikolas, I., Kane, D. M., Ge, R., Gupta, S., and Soltanolkotabi, M. Outlier-robust sparse estimation via non-convex optimization. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022.
- Cherapanamjeri, Y., Aras, E., Tripuraneni, N., Jordan, M. I., Flammarion, N., and Bartlett, P. L. Optimal Robust Linear Regression in Nearly Linear Time. *arXiv*, abs/2007.08137, 2020a.
- Cherapanamjeri, Y., Mohanty, S., and Yau, M. List decodable mean estimation in nearly linear time. In *Proc. 61st IEEE Symposium on Foundations of Computer Science (FOCS)*, 2020b.
- Dagan, Y. and Shamir, O. Detecting correlations with little memory and communication. In *Proc. 31st Annual Conference on Learning Theory (COLT)*, 2018.
- d’Aspremont, A., Ghaoui, L. E., Jordan, M. I., and Laffont, G. R. G. A direct formulation for sparse pca using semidefinite programming. *SIAM Review*, 49(3):434–448, 2007.
- Depersin, J. and Lecué, G. Robust subgaussian estimation of a mean vector in nearly linear time. *The Annals of Statistics*, 50(1):511–536, 2022.
- Diakonikolas, I. Computational-Statistical Tradeoffs and Open Problems. STOC 2019 Tutorial: Recent Advances in High-Dimensional Robust Statistics, 2019. See page number 36 in <http://www.iliasdiakonikolas.org/stoc19-tutorial/Tradeoffs-and-Open-Problems.pdf>.
- Diakonikolas, I. Algorithmic Robust Statistics. Statistical thinking in the age of AI : robustness, fairness and privacy (Meeting in Mathematical Statistics), 2023. Available online at <https://youtu.be/HKm0L2Cy69Y?t=3527>.
- Diakonikolas, I. and Kane, D. M. *Algorithmic High-Dimensional Robust Statistics*. Cambridge University Press, 2023.
- Diakonikolas, I., Kamath, G., Kane, D. M., Li, J., Moitra, A., and Stewart, A. Robust Estimators in High Dimensions without the Computational Intractability. In *Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016. doi: 10.1109/FOCS.2016.85.
- Diakonikolas, I., Kane, D. M., and Stewart, A. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *Proc. 58th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2017. doi: 10.1109/FOCS.2017.16.
- Diakonikolas, I., Kane, D. M., Karmalkar, S., Price, E., and Stewart, A. Outlier-robust high-dimensional sparse estimation via iterative filtering. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019.
- Diakonikolas, I., Kane, D. M., Karmalkar, S., Pensia, A., and Pittas, T. Robust sparse mean estimation via sum of squares. In *Proc. 35th Annual Conference on Learning Theory (COLT)*, 2022a.
- Diakonikolas, I., Kane, D. M., Kongsgaard, D., Li, J., and Tian, K. Clustering mixture models in almost-linear time via list-decodable mean estimation. In *Proc. 54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022b.
- Diakonikolas, I., Kane, D. M., Lee, J. C. H., and Pensia, A. Outlier-Robust Sparse Mean Estimation for Heavy-Tailed Distributions. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, 2022c.
- Diakonikolas, I., Kane, D. M., Pensia, A., and Pittas, T. Streaming Algorithms for High-Dimensional Robust Statistics. In *Proc. 39th International Conference on Machine Learning (ICML)*, 2022d.
- Diakonikolas, I., Kane, D. M., Pensia, A., and Pittas, T. Near-optimal algorithms for gaussians with huber contamination: Mean estimation and linear regression. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, 2023a.

- Diakonikolas, I., Kane, D. M., Pensia, A., and Pittas, T. Nearly-linear time and streaming algorithms for outlier-robust PCA. In *Proc. 40th International Conference on Machine Learning (ICML)*, 2023b.
- Dong, Y., Hopkins, S. B., and Li, J. Quantum Entropy Scoring for Fast Robust Mean Estimation and Improved Outlier Detection. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019.
- Eldar, Y. C. and Kutyniok, G. *Compressed sensing: theory and applications*. Cambridge University Press, 2012.
- Hastie, T., Tibshirani, R., and Wainwright, M. *Statistical Learning with Sparsity: The Lasso and Generalizations*. 2015.
- Huber, P. J. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, March 1964. ISSN 0003-4851. doi: 10.1214/aoms/1177703732.
- Huber, P. J. and Ronchetti, E. M. *Robust Statistics*. John Wiley & Sons, 2009.
- Jambulapati, A., Li, J., and Tian, K. Robust sub-gaussian principal component analysis and width-independent Schatten packing. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.
- Jiang, H., Kathuria, T., Lee, Y. T., Padmanabhan, S., and Song, Z. A faster interior point method for semidefinite programming. In *Proc. 61st IEEE Symposium on Foundations of Computer Science (FOCS)*, 2020.
- Karppa, M., Kaski, P., and Kohonen, J. A faster subquadratic algorithm for finding outlier correlations. *ACM Trans. Algorithms*, 14(3), 2018.
- Karppa, M., Kaski, P., Kohonen, J., and Catháin, P. Ó. Explicit correlation amplifiers for finding outlier correlations in deterministic subquadratic time. *Algorithmica*, 82(11), 2020.
- Lai, K. A., Rao, A. B., and Vempala, S. Agnostic Estimation of Mean and Covariance. In *Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 665–674, 2016. doi: 10.1109/FOCS.2016.76.
- Lei, Z., Luh, K., Venkat, P., and Zhang, F. A fast spectral algorithm for mean estimation with sub-gaussian rates. In *Proc. 33rd Annual Conference on Learning Theory (COLT)*, 2020.
- Polyanskiy, Y. and Wu, Y. *Information Theory: From Coding to Learning*. Cambridge University Press, 2023.
- Sachdeva, S. and Vishnoi, N. K. Faster algorithms via approximation theory. *Foundations and Trends® in Theoretical Computer Science*, 9:125–210, 2014.
- Valiant, G. Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem. *Journal of the ACM*, 62(2), 2015. ISSN 0004-5411. doi: 10.1145/2728167.
- Valiant, L. G. Functionality in neural nets. In *Proc. of the Seventh AAAI National Conference on Artificial Intelligence*, pp. 629–634. AAAI Press, 1988.
- van de Geer, S. *Estimation and Testing Under Sparsity*. École d’Été de Probabilités de Saint-Flour. Springer, 2016. ISBN 978-3-319-32774-7.
- Zhu, B., Jiao, J., and Steinhardt, J. Robust estimation via generalized quasi-gradients. *Information and Inference: A Journal of the IMA*, pp. 581–636, 2022. doi: 10.1093/imaia/iaab018.

## A. Preliminaries

In this section, we include details that were omitted from Section 2. Appendix A.1 provides the proof of Lemma 2.6. Appendix A.2 contains details about the standard template of filtering based algorithms, and Appendix A.3 applies this template to ensure the preprocessing condition. Finally, Appendix A.4 gives further details about the correlation detection algorithm.

### A.1. Defining Sparse Scores

In this section, we give the proof of Lemma 2.6.

**Lemma 2.6** (Sparse Filtering Lemma). *Let  $\epsilon \in (0, \epsilon_0)$  for a small absolute constant  $\epsilon_0$ . Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \delta, k)$ -stable with respect to  $\mu$ . Let  $H \subset [d]$  be such  $\|(\Sigma_T - \mathbf{I})_H\|_{\text{Fr}} = \lambda$  for  $\lambda \gtrsim \delta^2/\epsilon$  and  $|H| \leq k$ . There exists an algorithm  $\mathcal{A}$  that takes  $T, H, \epsilon$ , and  $\delta$  and returns scores  $f : T \rightarrow \mathbb{R}_+$  so that  $\sum_{x \in S \cap T} f(x) \leq \sum_{x \in T \setminus S} f(x)$ , i.e., the sum of scores over inliers is less than that of outliers, and  $\max_{x \in T} f(x) > 0$ . Moreover, the algorithm runs in time  $d \cdot \text{poly}(|H||T|)$ .*

*Proof.* As a first step, we simply project the data points along the coordinates in  $H$ , and by abusing the notation, call the projected set  $T$ . In the remainder of this proof,  $\mathbf{I}$  refers to  $\mathbf{I}_{|H|}$ . Computing this projected set takes at most  $d|T||H|$  time.

Let  $\lambda = \|\Sigma_T - \mathbf{I}\|_{\text{Fr}}$ . Define  $\mathbf{A} = (\Sigma_T - \mathbf{I})/\|\Sigma_T - \mathbf{I}\|_{\text{Fr}}$  so that  $\mathbf{A}$  maximizes the trace inner product with  $\Sigma_T - \mathbf{I}$  over unit Frobenius norm matrices. Define the function  $g(x) := (x - \mu_T)^\top \mathbf{A}(x - \mu_T) - \text{tr}(\mathbf{A})$ . We first compute the average of  $g(x)$  over the  $T$  below:

$$\mathbb{E}_T[g(x)] = \mathbb{E}_T[\langle (x - \mu_T)(x - \mu_T)^\top - \mathbf{I}, \mathbf{A} \rangle] = \langle \Sigma_T - \mathbf{I}, \mathbf{A} \rangle = \lambda.$$

By abusing notation again, we use  $S$  to denote the projection of the inliers  $S$  along the coordinates in  $H$ ; Note that the projected set also inherits  $(\epsilon, \delta, k)$ -stability and the  $\|\cdot\|_{\text{Fr}, k^2}$  reduces to the standard Frobenius norm. Thus, for any large subset of  $S' \subset S$  with  $|S'| \geq (1 - \epsilon)|S|$ , we use Lemma 2.5 to obtain the following:

$$\begin{aligned} |\mathbb{E}_{S'}[g(x)]| &= |\mathbb{E}_{S'}[\langle (x - \mu_T)(x - \mu_T)^\top - \mathbf{I}, \mathbf{A} \rangle]| \\ &= |\langle \Sigma_{S'} - \mathbf{I}, \mathbf{A} \rangle + 2(\mu - \mu_T)^\top \mathbf{A}(\mu - \mu_{S'}) + (\mu - \mu_T)^\top \mathbf{A}(\mu - \mu_T)| \\ &\leq \|\Sigma_{S'} - \mathbf{I}\|_{\text{Fr}} + 2\|\mu - \mu_T\|_2 \|\mathbf{A}\|_{\text{Fr}} \|\mu - \mu_{S'}\|_2 + \|\mu - \mu_T\|_2^2 \|\mathbf{A}\|_{\text{Fr}}^2 \\ &\lesssim \delta^2/\epsilon + 2(\delta + \sqrt{\epsilon\lambda})\delta + (\delta + \sqrt{\epsilon\lambda})^2 && \text{(using stability and Lemma 2.5)} \\ &\lesssim \delta^2/\epsilon + 3\delta^2 + 4\delta\sqrt{\epsilon\lambda} + \epsilon\lambda \\ &\lesssim \delta^2/\epsilon + 7\delta^2 + 2\epsilon\lambda && \text{(using } 2ab \leq a^2 + b^2\text{)} \\ &\lesssim \delta^2/\epsilon + \epsilon\lambda, \end{aligned} \tag{3}$$

where we used that  $\epsilon \leq 1$ . The following helper result, which is a slight generalization of Diakonikolas & Kane (2023, Proposition 2.19) from non-negative  $g$ 's to real-valued  $g$ 's will be useful.

**Claim A.1.** *Let  $h : S \rightarrow \mathbb{R}$  be a real-valued function on a finite set  $S$ . Further suppose that  $|\mathbb{E}_{S'}[h(X)]| \leq \tau$  for all sets  $S' \subset S$  with  $|S'| \geq (1 - \epsilon)|S|$ . Then for  $f'(x) = h(x)\mathbf{1}_{h(x) \geq 3\tau/\epsilon}$ , we have that  $\mathbb{E}_S[f'(x)] \leq 3\tau$ .*

*Proof.* We shall do it in two steps. First, we show the following: for all subsets  $S'' \subset S$  with  $|S''| \leq \epsilon|S|$ .

$$\mathbb{E}_{S''}[\max(h(X), 0)] \leq \tau + 2\tau/\epsilon. \tag{4}$$

To that end, for all sets  $S''$  with  $|S''| = \epsilon|S|$ , the triangle inequality implies

$$|\mathbb{E}_{S''}[h(X)]| = \left| \frac{\mathbb{E}_S[h(X)] - (1 - \epsilon)\mathbb{E}_{S \setminus S''}[h(X)]}{\epsilon} \right| \leq 2\tau/\epsilon. \tag{5}$$

Let  $S_*$  be the top  $\epsilon|S|$  entries of  $S$  in the increasing order of  $h(\cdot)$ , not just in the absolute value. Then establishing (4) is equivalent to establish an upper upper bound on  $\frac{1}{|S_*|} \sum_{x \in S_*} \max(h(X), 0)$ . Thus, if all the entries of  $S_*$  are bigger than

0, then (4) follows by (5). We shall show that  $h(x)$  on  $S_*$  is lower bounded by  $-\tau$ . Under this condition, we see that the desired result follows similarly by (5):  $\frac{1}{|S_*|} \sum_{x \in S_*} \max(h(X), 0) \leq \frac{1}{|S_*|} \sum_{x \in S_*} (h(X) + \tau) \leq \tau + 2\tau/\epsilon$ . We now establish that  $\min_{x \in S_*} h(x) \geq -\tau$ . If there exists an  $x \in S_*$  with  $h(X) \leq -\tau$ , then the average of  $S \setminus S_*$  must be less than  $-\tau$ , contradicting the assumption that  $|\mathbb{E}_{S \setminus S_*}[h(X)]| \leq \tau$ . Thus, we have established (4).

Given (4), we see that the fraction of points with  $h(x) \geq 3\tau/\epsilon$  must be less than  $\epsilon|S|$ . Otherwise, the conditional average over those points would be at least  $3\tau/\epsilon \geq 2\tau/\epsilon + \tau$ , contradicting (4). Therefore, the function  $f$  is non-zero only on at most  $\epsilon$ -fraction of  $S$ . Therefore, the non-negativity of  $f$  implies that

$$\frac{1}{|S|} \sum_{x \in S} f(X) \leq \frac{1}{|S|} \max_{S'' \subset S: |S''| \leq \epsilon|S|} \sum_{x \in S''} f(X) \leq \frac{1}{|S|} \max_{S'' \subset S: |S''| \leq \epsilon|S|} \sum_{x \in S''} \max(h(X), 0) \leq \epsilon(\tau + 2\tau/\epsilon), \quad (6)$$

where we use non-negativity of  $f$  and (4). □

Let  $R := C'(\delta^2/\epsilon + \epsilon\lambda)$  for a large enough constant  $C'$  be the bound from (3). Combining this with the claim above, we see that defining  $f(x)$  to be  $g(x)\mathbf{1}_{x \geq 3R/\epsilon}$ , the sum of scores over inliers is small:

$$\sum_{x \in S \cap T} f(x) \leq \sum_{x \in S} f(x) \leq 3R|S| = 3C'(\delta^2/\epsilon + \epsilon\lambda)|S| \leq 0.25\lambda|S|, \quad (7)$$

where we use that  $3C'\epsilon \leq 1/8$  and  $3C'\delta^2/\epsilon \leq \lambda/8$ .

On the other hand,  $\mathbb{E}_{x \in T \setminus S}[f(x)]$  must be large argued as argued below. We observe that  $f(x) \geq g(x) - 3R/\epsilon$ , and thus applying (3) to  $T \cap S$ , which is of size at least  $(1 - \epsilon)|T| = (1 - \epsilon)|S|$ , we obtain

$$\begin{aligned} \sum_{x \in T \setminus S} f(x) &\geq \sum_{x \in T \setminus S} (g(x) - 3R/\epsilon) = \left( \left( \sum_{x \in T} g(x) \right) - \left( \sum_{x \in T \cap S} g(x) \right) \right) - (|T \setminus S|3R/\epsilon) \\ &\geq (\lambda|T|) - (|T \cap S|R) - \epsilon|T|(3R/\epsilon) \\ &\geq \lambda|T| - 4|T|R \geq \lambda|T|/2, \end{aligned} \quad (8)$$

where the last inequality follows if we show that  $R \leq \lambda/8$ . Indeed, this follows if  $C'\delta^2/\epsilon \leq \lambda/16$  and  $C'\epsilon \leq 1/16$ . Combining (7) and (8), we get the desired result; the claim on  $\max_{x \in T} f(x)$  follows from (8). The complete algorithm is given below:

---

**Algorithm 4** Quadratic Scores

---

- 1:  $T \leftarrow \{(x)_H : x \in T\}$  ▷ Projection onto  $H$
  - 2: Let  $\mathbf{A}$  be the matrix  $(\Sigma_T - \mathbf{I})/\|\Sigma_T - \mathbf{I}\|_{\text{Fr}}$  ▷ The matrix such that  $\langle \mathbf{A}, \Sigma_T - \mathbf{I} \rangle = \|\Sigma_T - \mathbf{I}\|_{\text{Fr}}$
  - 3: Define  $g(x) := (x - \mu_T)^\top \mathbf{A}(x - \mu_T) - \text{tr}(\mathbf{A})$
  - 4: Define  $f(x)$  to be  $g(x)$  if  $g(x) \geq 3C' \left( \frac{\delta^2}{\epsilon^2} + \lambda \right)$  otherwise 0
  - 5: Return  $f$
- 

□

## A.2. Randomized Filtering

We will use the following template of filtering algorithm from [Diakonikolas & Kane \(2023, Section 2.4.2\)](#), after a slight change in parameters.

**Algorithm 5** Randomized Filtering

---

```

1: Let  $T_1 \leftarrow T$ 
2:  $i \leftarrow 1$ 
3: while  $T_i \neq \emptyset$  and  $T_i$  does not satisfy the stopping condition  $\mathcal{S}$  do
4:   Get the scores  $f : T_i \rightarrow \mathbb{R}_+$  satisfying  $\sum_{x \in T_i \cap S} f(x) \leq \sum_{x \in T_i \setminus S} f(x)$  and  $\max_{x \in T_i} f(x) > 0$ 
5:    $T_{i+1} \leftarrow T_i$ 
6:   for each  $x \in T_i$  do
7:     Remove the point  $x$  from  $T_{i+1}$  with probability  $\frac{f(x)}{\max_{x \in T_i} f(x)}$ 
8:   end for
9:    $i \leftarrow i + 1$ 
10: end while
11: return  $T_i$ 

```

---

These filtering algorithms have become a standard template in algorithmic robust statistics. Here, the stopping condition  $\mathcal{S}$  can be a generic condition that can be evaluated in  $\mathcal{T}_{\text{stopping}}$  time using some algorithm  $\mathcal{A}_s$ —it can be a randomized algorithm (using independent randomness from Line 6) that succeeds with probability  $1 - \tau$ . We also require that whenever the stopping condition is not satisfied and the set  $T_i$  is an  $10\epsilon$ -corruption of  $S$ , then the scores  $f : T_i \rightarrow \mathbb{R}_+$  satisfying the guarantees of Line 4 can be computed in time  $\mathcal{T}_{\text{score}}$ .

**Theorem A.2** (Guarantee of Algorithm 5; Diakonikolas & Kane (2023, Theorem 2.17)). *If the above stopping conditions and filter conditions are met, then Algorithm 5 returns a set  $T' \subseteq T$  such that, with probability at least  $8/9 - \tau|T|$ ,*

1. Algorithm 5 runs in time  $O(|T|(\mathcal{T}_{\text{stopping}} + \mathcal{T}_{\text{score}} + |T|))$ .
2. Each set  $T_i \subseteq T$  observed throughout the run of the algorithm (which includes  $T'$ ) is an  $10\epsilon$ -corruption of  $S$ .
3.  $T'$  satisfies the stopping condition  $\mathcal{S}$ .

### A.3. Preprocessing

In this section, we outline how to ensure Condition 2.7 quickly using Algorithm 5 and Lemma 2.6.

**Condition 2.7** (Preprocessing). Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \delta, k)$ -stable. Suppose  $T$  satisfies  $\|\text{diag}(\Sigma_T - \mathbf{I}_d)\|_{\text{Fr}, k^2} \leq \min(O(\delta^2/\epsilon), 0.5)$ .

Given the sparse filtering lemma (Lemma 2.6), we can simply filter along the diagonals to ensure Condition 2.7 as shown below.

**Claim A.3.** *Let  $\epsilon \in (0, \epsilon_0)$  for a small absolute constant  $\epsilon_0$ . Let  $c$  be a small enough absolute constant and  $C$  be a large enough constant. Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \delta, k^2)$ -stable with respect to  $\mu$  such that  $\delta^2/\epsilon \leq c$ . Then there is a randomized algorithm  $\mathcal{A}$  that takes as inputs  $\epsilon, \delta$ , and  $k$  such that it outputs a set  $T' \subseteq T$  such that with probability at least  $8/9$ : (i)  $T'$  is at most  $10\epsilon$ -corruption of  $S$  and (ii)  $\|\text{diag}(\Sigma_{T'} - \mathbf{I}_d)\|_{\text{Fr}, k^2} \leq C\delta^2/\epsilon \leq 0.1$ , and (iii) the algorithm runs in time  $\tilde{O}(|T|dk^2 + |T|^2d)$ .*

*Proof.* We can simply invoke Algorithm 5 with the stopping condition on the set  $T_i$  set to  $\|\text{diag}(\Sigma_{T_i} - \mathbf{I}_d)\|_{\text{Fr}, k} \leq C\delta^2/\epsilon$  for a constant  $C$  large enough. To evaluate this stopping condition, we can simply compute the matrix  $\text{diag}(\Sigma_{T'} - \mathbf{I}_d)$  in  $O(d|T|)$  time. The associated  $\|\cdot\|_{\text{Fr}, k^2}$  can be easily calculated by computing the Euclidean norm of its largest  $k^2$  entries, again computable in  $\tilde{O}(dk^2)$  time.

If the stopping condition is not satisfied, Lemma 2.6 returns the required scores. Thus, we get the desired guarantees on the set  $T$  from Theorem A.2.  $\square$

The next result shows that any further small modifications of the preprocessed sets retains small sparse operator norm.

**Lemma A.4.** *Let  $C$  be a large enough constant  $C > 0$ . Let  $T'' \subset T'$  be two  $O(\epsilon)$ -contamination of  $S$  such that  $S$  is an  $(C\epsilon, \delta, k)$ -stable with respect to  $\mu$ . Suppose that  $\|\text{diag}(\Sigma_{T'} - \mathbf{I}_d)\|_{\text{op}, k} \lesssim \delta^2/\epsilon$ . Then  $\|\text{diag}(\Sigma_{T''} - \mathbf{I}_d)\|_{\text{op}, k} \lesssim \delta^2/\epsilon$*

*Proof.* First we note that the lower bound on the sparse eigenvalues follow rather directly as shown below. We make use of the equality  $\Sigma_{T''} := \frac{1}{|T''|^2} \sum_{x,y \in T''} (x-y)(x-y)^\top$ . For any sparse unit vector  $v$ , we use the stability condition applied to  $S \cap T''$  to obtain the following:

$$\begin{aligned}
 v^\top \Sigma_{T''} v &= \frac{1}{|T''|^2} \sum_{x,y \in T''} (v^\top (x-y))^2 \geq \frac{1}{|T''|^2} \sum_{x,y \in S \cap T''} (v^\top (x-y))^2 && \text{(using non-negativity)} \\
 &= \frac{|S \cap T''|^2}{|T''|^2} v^\top \Sigma_{S \cap T''} v \geq (1 - O(\epsilon)) v^\top \Sigma_{S \cap T''} v \\
 &= (1 - O(\epsilon)) \left( v^\top \left( \mathbb{E}_{S \cap T''} [(X - \mu)(X - \mu)^\top] \right) v - (v^\top (\mu - \mu_{S \cap T''}))^2 \right) \\
 &\geq (1 - O(\epsilon)) (1 - O(\delta^2/\epsilon) - O(\delta^2)) && \text{(using stability)} \\
 &\geq 1 - O(\delta^2/\epsilon), && (9)
 \end{aligned}$$

where the last inequality uses  $\epsilon \leq \delta$ . For the upper bound, we observe that for any matrix  $\mathbf{A}$ ,  $\|\text{diag}(\mathbf{A})\|_{\text{op},k}$  is attained by 1-sparse unit vectors  $v$ , i.e.,  $\|\mathbf{A}\|_{\text{op},1} = \|\text{diag}(\mathbf{A})\|_{\text{op},k}$ . Thus, for any  $T'' \subset T'$  and that  $|T'| \leq |T''|(1 + O(\epsilon))$ : for any 1-sparse unit vector  $v$ ,

$$\begin{aligned}
 v^\top \Sigma_{T''} v &= \frac{1}{|T''|^2} \sum_{x,y \in T''} (v^\top (x-y))^2 \leq \frac{|T'|^2}{|T''|^2} \frac{1}{|T'|^2} \sum_{x,y \in T'} (v^\top (x-y))^2 && \text{(using nonnegativity)} \\
 &= \frac{|T'|^2}{|T''|^2} v^\top \Sigma_{T'} v \leq (1 + O(\epsilon)) v^\top \Sigma_{T'} v \leq (1 + O(\epsilon))(1 + O(\delta^2/\epsilon)) = 1 + O(\delta^2/\epsilon), && (10)
 \end{aligned}$$

where we use that for 1-sparse unit vectors  $v$ ,  $v^\top \Sigma_{T'} v = v^\top \text{diag}(\Sigma_{T'}) v$ . Combining Equations (9) and (10) for all 1-sparse unit vectors  $v$ , we obtain that  $\|\text{diag}(\Sigma_{T''} - \mathbf{I})\|_{\text{op},k} = O(\delta^2/\epsilon)$ . □

#### A.4. Fast Correlation Detection

In this section, we show how to obtain Theorem 2.8 from Valiant (2015, Theorem 1).

**Theorem A.5** (Robust Correlation Detection For Boolean Vectors in Subquadratic Time Valiant (2015, Theorem 1)). *Consider a set of  $n'$  vectors in  $\{-1, 1\}^{d'}$  and constants  $\rho, \tau \in [0, 1]$  with  $\rho > \tau$  such that for all but at most  $s$  pairs  $u, v$  of distinct vectors,  $|u^\top v| / \|u\|_2 \|v\|_2 \leq \tau$ . There is an algorithm that, with probability  $1 - o(1)$ , will output all pairs of vectors whose normalized inner product is least  $\rho$ . Additionally, the runtime of the algorithm is*

$$\left( sd' n'^{0.62} + n'^{1.62 + 2.4 \frac{\log(1/\rho)}{\log(1/\tau)}} \right) \text{poly}(\log n, 1/\tau).$$

An improved algorithm with better runtime was then provided in Karppa et al. (2018, Corollary 1.8), but we choose the version above for its simplicity. We provide the proof of Theorem 2.8, the version we used in this work, from Theorem A.5 using standard arguments below:

*Proof of Theorem 2.8 from Valiant (2015, Theorem 1).* Let  $X \in \mathbb{R}^{d \times n}$  denote the matrix with the columns of  $X$  denoting the centered vectors of  $T$ . That is, if  $T = \{z_1, \dots, z_n\}$ , then the  $i$ -th column of  $X$  is equal to  $z_i - \mu_T$ . Let  $X_i$  denote the  $i$ -th row of the matrix  $X$ . For  $i, j \in [d] \times [d]$ , the correlation between  $i$ -th and  $j$ -th coordinate on  $T$ ,  $\text{corr}(i, j)$ , is equal to  $\left| \frac{X_i^\top X_j}{\|X_i\|_2 \|X_j\|_2} \right|$ . Thus, we would like to apply Theorem A.5 with the rows of  $X$  (thus  $n' = d$  and  $d' = n$ ). However,  $X'$  is not a binary matrix.

A standard reduction allows us to compute a binary matrix that preserves the correlation between the rows of  $X$ . Let  $G \in \mathbb{R}^{n \times m}$  be a matrix with independent  $\mathcal{M}(0, 1)$  entries. Let  $Y = XG$  be in  $\mathbb{R}^{d \times m}$  and let  $Y' = \text{sgn}(Y) \in \mathbb{R}^{d \times m}$ , where  $\text{sgn}$  is applied elementwise. Thus,  $Y'$  is a boolean matrix as required in Theorem A.5. The following arguments show that the correlation between the rows of  $Y$  is preserved for  $m$  large enough. Let  $Y_i, Y'_i$  denote the  $i$ -th row of the matrices  $Y, Y'$ .

**Lemma A.6** (Valiant (2015, Lemma 21)). *If  $m \geq 10 \log(n)/\gamma^2$ , then with probability  $1 - o(1)$ , we have that for all  $i \neq j \in [n]$ , we have that*

$$\left| \frac{\langle Y'_i, Y'_j \rangle}{\|Y'_i\|_2 \|Y'_j\|_2} - \frac{2}{\pi} \arcsin \left( \frac{\langle X_i, X_j \rangle}{\|X_i\|_2 \|X_j\|_2} \right) \right| \leq \gamma.$$

In particular, if the original correlation is less than  $\tau$  in the absolute value, then the corresponding correlation in  $Y'$  in the absolute value is at most  $(2/\pi) \arcsin(\tau) + \gamma \leq (4/\pi)\tau + \gamma \leq 2\tau + \gamma$ , where we use  $|\arcsin(x)| \leq 2|x|$ . Similarly, if the original correlation is at least  $\rho$  in absolute value, then the new correlation is at least  $(2/\pi) \arcsin(\rho) - \gamma \geq (2/\pi)\rho - \gamma \geq \rho/2 - \gamma$ . Choosing  $\gamma = \tau$ , the new matrix  $Y'$  satisfies the guarantees of Theorem A.5 with  $n' = d, d' = m = 10 \log(n)/\tau^2, \tau' = 3\tau, \rho' = \rho/4$  (using  $\tau \leq \rho/4$ ). The time taken to compute  $Y$  and  $Y'$  is at most  $ndm$ . Thus, the total runtime is at most

$$\begin{aligned} & \left( nd \frac{\log n}{\gamma^2} + s \frac{\log n}{\tau^2} d^{0.62} + d^{1.62+2.4 \frac{\log(4/\rho)}{\log(1/3\tau)}} \text{poly}(\log n, 1/\tau) \right) \\ & \leq \left( sd^{0.62} + d^{1.62+2.4 \frac{\log(4/\rho)}{\log(1/3\tau)}} \right) \cdot \text{poly}(n(\log d)/\gamma^2). \end{aligned}$$

The probability of success can be boosted using repetition, if needed. □

## B. Key Subroutines and Their Proofs

In this section, we provide the details omitted from Section 3. Appendix B.1 and Appendix B.2 provide the proofs of Lemma 3.2 and Proposition 3.3, respectively.

### B.1. Proof of Lemma 3.2

**Lemma 3.2.** *Let  $T \subset \mathbb{R}^d$  be a multiset with covariance matrix  $\Sigma'$ . Let  $m \in \mathbb{N}$  be the sampling parameter. Let  $J_* \subset [d] \times [d]$  be the off-diagonal coordinate pairs such that  $|\Sigma'_{i,j}| \geq \tau$ . Then Algorithm 2 takes  $\tau, m$ , and  $T$  as input and returns a set  $J \subset J_*$  in time  $O(m|T|)$  such that with probability  $1 - 1/d^2$ ,  $|J| \geq \frac{m|J_*|}{4d^2} - 16 \log d$ .*

*Proof.* We repeat the algorithm below:

---

**Algorithm 3.1** RANDOMLYCHECKCOORDINATES

---

- 1: Let  $H_{\text{pair}} \subset [d] \times [d]$  of size  $m$ , with  $(i, j) \in H_{\text{pair}}$  sampled i.i.d. from off-diagonal elements.
  - 2: Let  $J \leftarrow \{(i, j) \in H_{\text{pair}} : |\Sigma'_{i,j}| \geq \tau\}$  for  $\Sigma' = \Sigma_T$
  - 3: **return**  $J$ .
- 

To show correctness, we shall use the following concentration inequality for Binomials: If  $X \sim \text{Ber}(n, p)$ , then with probability  $1 - \delta$ ,  $|\sqrt{X} - \sqrt{np}| \leq 2\sqrt{\log(1/\delta)}$ . See, for example, Polyanskiy & Wu (2023, Equations (15.21) and (15.22)). In particular, with probability  $1 - 1/d^4$ ,  $\sqrt{X} \geq \sqrt{np} - 4\sqrt{\log d}$ , which implies  $X \geq 0.25np - 16 \log d$ .<sup>14</sup>

The probability that a single pair in  $H_{\text{pair}}$  has correlation of magnitude at least  $\tau$  is exactly  $(|J_*|/d(d-1))$ , and thus  $|J| \sim \text{Ber}(m, |J_*|/(d(d-1)))$ . Therefore, applying the Binomial concentration, with probability  $1 - 1/d^2$ , it holds that  $|J| \geq \frac{m|J_*|}{4d^2} - 16 \log d$ . The claim about the runtime is immediate. □

### B.2. Proof of Proposition 3.3

**Proposition 3.3** (Subroutine to Identify Corrupted Coordinates Algorithm 3). *Let corrupted set  $T$ , Frobenius threshold  $\kappa \in \mathbb{R}$ , correlation threshold  $\rho \in (0, 1)$ , margin threshold  $\tau \in (0, \rho/12)$ , sampling parameter  $m \in \mathbb{N}$ , and correlation*

---

<sup>14</sup>We use that if  $a, b, c \in \mathbb{R}_+$  then  $a \geq b - c$  implies that  $a^2 \geq b^2/4 - c^2$ . The proof is as follows: if  $b \geq 2c$  then  $a \geq b/2$  and thus  $a \geq b^2/4 \geq b^2/4 - c^2$ ; otherwise  $b^2/4 \leq c^2$  and thus  $a^2 \geq 0 \geq b^2/4 - c^2$  holds trivially.

count  $s \in \mathbb{N}$ . Suppose that each diagonal entry of  $\Sigma_T$  lies in  $[1/2, 3/2]$  and  $\frac{c_1 d^2}{m} \left( \frac{\kappa^2}{\tau^2} + \log d \right) \leq s$ .

Then Algorithm 3 takes as input  $T, \kappa, \rho, \tau, m, s$  as input and satisfies the following with probability  $1 - 1/\text{poly}(d)$ :

- It either outputs a set  $H \subset [d]$  with  $|H| \leq \frac{\kappa^2}{\tau^2}$  and  $\|(\text{offdiag}(\Sigma_T))_H\|_{\text{Fr}} \geq \kappa/2$ .
- Else, it outputs “ $\perp$ ”. If it outputs “ $\perp$ ”, then  $\|\text{offdiag}(\Sigma_T)\|_{\text{Fr}, k^2}$  is at most  $2\kappa + 2\rho k$ .

Moreover, the algorithm runs in  $O\left(m + sd^{0.62} + d^{1.62+3\frac{\log(4/\rho)}{\log(1/3\tau)}}\right)\text{poly}(n, \log d, 1/\tau)$  time.

We restate the algorithm with extended notation in Algorithm 6.

---

**Algorithm 6** Main Subroutine (Expanded version of Algorithm 3)

---

**Require:** Frobenius threshold  $\kappa \in \mathbb{R}_+$ , a finite set  $T \subset \mathbb{R}^d$  such that diagonal entries of  $\Sigma_T$  lie in  $[1/2, 2]$ , correlation threshold  $\rho \in (0, 1)$ , weak correlation threshold  $\tau$ , sampling parameter  $m \in \mathbb{N}$ . We require the parameters to satisfy

$$\frac{c_1 d^2}{m} \left( \frac{\kappa^2}{\tau^2} + \log d \right) \leq s. \quad (11)$$

**Ensure:** With high probability, output either (i) a set  $H \subset [d]$  with  $|H| \leq k''$  and  $\|(\text{offdiag}(\Sigma_T))_H\|_{\text{Fr}} \geq 0.5\kappa$  or (ii) “ $\perp$ ”; If it outputs “ $\perp$ ”, then  $\|(\text{offdiag}(\Sigma_T))_H\|_{\text{Fr}} \leq 2\kappa + 2k\rho$ .

- 1: Let  $H'_{\text{pair}}$  be the output of Random Subsampling Algorithm (Algorithm 2) with  $m$  and  $\tau$
  - 2: **if**  $|H'_{\text{pair}}| \geq \kappa^2/\tau^2$  **then**
  - 3:     Set  $H_1 \leftarrow \text{Proj}_{\kappa^2/\tau^2}(H_{\text{proj}})$
  - 4:     **return**  $H_1$
  - 5: **else**
  - 6:     Let  $H_{\text{pair}} \subseteq [d] \times [d]$  be the output of Theorem 2.8 with  $\tau, \rho, s$  on  $T$
  - 7:     **if**  $|H_{\text{pair}}| > \kappa^2/\rho^2$  **then**
  - 8:          $H_3 \leftarrow \text{Proj}_{\kappa^2/\rho^2}(H_{\text{proj}})$
  - 9:         **return**  $H_3$
  - 10:     **else**
  - 11:         Let  $H_2 \leftarrow \text{Proj}(H_{\text{pair}})$
  - 12:         **if**  $\|(\text{offdiag}(\Sigma_S))_{H_2}\|_{\text{Fr}} \geq \kappa$  **then**
  - 13:             **return**  $H_2$
  - 14:         **else**
  - 15:             **return** “ $\perp$ ”
  - 16:         **end if**
  - 17:     **end if**
  - 18: **end if**
- 

*Proof.* Let  $\mathbf{A}$  be the matrix with diagonal entries zero and off-diagonal entry  $(i, j)$  equal to  $\frac{\Sigma'_{i,j}}{\sqrt{\Sigma'_{i,i}\Sigma'_{j,j}}}$ . Since diagonal entries of  $\Sigma_T$  lie in  $[0.5, 1.5]$ , the entries of  $\mathbf{A}$  and  $\text{offdiag}(\Sigma_T) = \text{offdiag}(\Sigma_T - \mathbf{I}_d)$  have the same signs and have the magnitude up to a factor of 2.

Thus, at the cost of constant factor, we will upper bound  $\|\mathbf{A}\|_{\text{Fr}, k^2}$  and lower bound  $\|(\mathbf{A})_H\|_{\text{Fr}}$  instead of dealing with  $\text{offdiag}(\Sigma_T - \mathbf{I}_d)$ .

We will use the notation from Algorithm 3. Let  $s_*$  be the number of off-diagonal coordinates of  $\mathbf{A}_*$  that have absolute value larger than  $\tau$ . Our algorithm will be correct on the event when both Theorem 2.8 and Lemma 3.2 succeed, which happens with high probability. For the rest of this proof, we condition on both of these algorithms succeeding.

We first consider the case when  $s_* > s$ . Then, the lower bound on  $s$  in the statement (cf. (11)), coupled with Lemma 3.2, implies that

$$|H'_{\text{pair}}| > \frac{ms_*}{4d^2} - 16 \log d > \frac{ms}{4d^2} - 16 \log d$$

$$> \frac{m}{4d^2} \left( \frac{c_1 d^2}{m} \left( \frac{\kappa^2}{\tau^2} + \log d \right) \right) - 16 \log d > \frac{c_1 \kappa^2}{2\tau^2},$$

where we use  $c_1$  is large enough, say  $c_1 \geq 100$ . By definition of  $H_1 = \text{Proj}_{\kappa^2/\tau^2}(H_{\text{proj}})$ , there are at least  $|\kappa^2/\tau^2|$  entries in  $(\mathbf{A})_{H_1}$  with absolute value at least  $\tau$ , and thus  $(\mathbf{A})_{H_1}$  has Frobenius norm at least  $\sqrt{\kappa^2/\tau^2}\tau = \kappa$ .

Consider the alternate case when  $s_* < s$ . If  $|H'_{\text{pair}}|$  happens to be large, then the same argument as above implies the correctness and the runtime. Thus, in the rest of this proof, we consider the case when we enter Line 5. Since  $s_* \leq s$ , Theorem 2.8 runs in the desired time and finds all off-diagonal coordinate pairs of  $A$ , collected in  $H_{\text{pair}}$ , with entries bigger than  $\rho$  in the promised time. If  $H_{\text{pair}}$  has more than  $\kappa^2/\rho^2$  entries, then by definition of  $H_3 = \text{Proj}_{\kappa^2/\rho^2}(H_{\text{pair}})$ , the same argument as above implies that the Frobenius norm of  $(\mathbf{A})_{H_3}$  is at least  $\kappa$ .

If  $|H_{\text{pair}}| < \kappa^2/\rho^2$ , then  $H_2$  is defined to contain all coordinates that appear in  $H_{\text{pair}}$ . If Line 12 succeeds, then the algorithm returns a subset of coordinates satisfying the desired conditions (small size and Frobenius norm larger than  $\kappa$ ). Suppose the if condition is not satisfied (and we return “ $\perp$ ”). We argue that  $\|\mathbf{A}\|_{\text{Fr},k^2}$  is small enough. Write  $\mathbf{A} = \mathbf{B} + \mathbf{B}'$ , where  $\mathbf{B}_{i,j}$  is non-zero only when  $i \in H_2, j \in H_2$  with values equal to  $\mathbf{A}_{i,j}$  and zero otherwise ( $\mathbf{B}'$  is then defined to be  $\mathbf{A} - \mathbf{B}$ ). By definition of  $H_2$ , we know each entry of  $\mathbf{B}'$  is at most  $\rho$  in absolute value. By triangle inequality,  $\|\mathbf{A}\|_{\text{Fr},k^2} \leq \|\mathbf{B}\|_{\text{Fr},k^2} + \|\mathbf{B}'\|_{\text{Fr},k^2} = \|(\mathbf{A})_{H_2}\|_{\text{Fr},k^2} + \|\mathbf{B}'\|_{\text{Fr},k^2} \leq \|(\mathbf{A})_{H_2}\|_{\text{Fr}} + k\rho \leq \kappa + k\rho$ . The accuracy guarantee follows by noting that the entries of  $\mathbf{A}$  and  $\Sigma - \mathbf{I}$  are within a factor of 2. Finally, the runtime guarantees are immediate by Theorem 2.8 and Lemma 3.2.  $\square$

## C. Complete Algorithm and Proof of Correctness

We now present the main technical result of our paper, the formal version of Theorem 3.1, below.

**Theorem C.1** (Robust Sparse Mean Estimation in Subquadratic Time; Formal version of Theorem 3.1). *Let  $c$  be a small enough absolute constant and  $C$  be a large enough absolute constant. Consider the corruption rate  $\epsilon \in (0, \epsilon_0)$ , where  $\epsilon_0$  is a small enough absolute constant. Let  $k \in \mathbb{N}$  be the sparsity parameter and  $q \in \mathbb{N}$  the correlation decay parameter with  $q \geq 3$ . Let  $T$  be an  $\epsilon$ -corrupted version of a set  $S$ , where  $S$  satisfies  $(C\epsilon, \delta, k')$ -stability with respect to  $\mu$  for  $k' := \frac{(Ck)^q}{(\delta^2/\epsilon)^{q-1}}$  and  $\delta^2/\epsilon \leq c$ . Then there is a randomized algorithm (Algorithm 7) that takes as inputs  $T, \epsilon, \delta, k$ , and  $q$  and produces an estimate  $\hat{\mu}$  such that, with a probability at least  $1 - 1/d^2$  over the randomness of the algorithm, we have the following guarantees:*

- (Error)  $\|\hat{\mu} - \mu\|_{2,k} \lesssim \delta$ .
- (Runtime) The algorithm runs in time at most  $d^{1.62 + \frac{3}{q}} \text{poly}(|T|, \log d, k^q, 1/\epsilon^q)$ .

We present the proof in Appendix C.1 below. We then show how Theorem C.1 implies Theorem 1.5 in Appendix C.2.

### C.1. Proof of Theorem C.1

The complete algorithm is given below: We now present the proof of its correctness:

*Proof of Theorem C.1.* The first step of Algorithm 7 is the fast preprocessing step from Claim A.3, which takes at most  $\tilde{O}(dk^2|T|^2)$  time and removes not too many inliers with high probability. In particular, the returned set  $T'$  has diagonal values in  $[1/2, 2]$ . In fact, Lemma A.4 implies that the subsequent sets  $T_i$ 's will satisfy this property, at least until we remove more than  $\Omega(\epsilon)$  fraction of points, which is the regime of interest anyway (cf. Theorem A.2). Thus, Proposition 3.3 will be applicable.

We will follow the standard proof template of randomized filtering algorithm from Theorem A.2, with the stopping condition provided by Proposition 3.3. The choice of parameters in Line 4 are such that whenever Proposition 3.3 does not return “ $\perp$ ” (and  $T_i$  is  $10\epsilon$ -corruption of  $S$ ), it returns a set  $H$  satisfying the guarantees of Lemma 2.6. To see this, observe that whenever Algorithm 3 outputs a subset  $H \subset [d]$ , then  $|H| \leq \kappa^2/\tau^2$  and  $\|(\Sigma_{T_i} - \mathbf{I}_d)_H\|_{\text{Fr}} > \|(\text{offdiag}(\Sigma_{T_i}))_H\|_{\text{Fr}} > \kappa/4 \gtrsim \delta^2/\epsilon$ . Moreover, whenever Proposition 3.3 returns “ $\perp$ ”, then  $\|(\Sigma_{T_i} - \mathbf{I}_d)_H\|_{\text{Fr},k^2} \lesssim \kappa + k\rho \lesssim \delta^2/\epsilon$  since  $\kappa \asymp \delta^2/\epsilon$  and  $\rho \asymp \delta^2/\epsilon$ . The choice of sparsity parameter in the stability of inliers is  $k'$ , which is larger than  $\kappa^2/\tau^2$ , because the choice of  $\kappa$  above and  $\tau = (\rho/12)^q$  imply that  $\kappa^2/\tau^2$  is of the order  $(\delta^2/\epsilon)/(c'\rho^q) \asymp (\delta^2/\epsilon)/((c'\delta^2/\epsilon)/k)^q \asymp \frac{k^q}{C^q(\delta^2/\epsilon)^{q-1}}$ . Thus, the scores generated by  $H$  using Lemma 2.6 give more weights to outliers than inliers. Theorem A.2 now implies that the final set  $T''$ ,

**Algorithm 7** Main Algorithm

**Require:** corruption rate  $\epsilon \in (0, 1)$ , stability parameter  $\delta \in (0, 1)$ , corrupted set  $T \subset \mathbb{R}^d$ , correlation-threshold  $\rho \in (0, 1)$ , correlation-decay  $q \in \mathbb{N}$ , sparsity  $k \in \mathbb{N}$ , sampling parameter  $m \in \mathbb{N}$ . We require  $T$  to be an  $\epsilon$ -corrupted version of an  $(C\epsilon, \delta, k')$  stable set with respect to  $\mu$  for  $k' := \frac{(Ck)^q}{(\delta^2/\epsilon)^{q-1}}$  and  $\delta^2/\epsilon \leq c$  for a small absolute constant  $c > 0$ .

**Ensure:**  $\hat{\mu} \in \mathbb{R}^d$  such that, with high probability,  $\|\hat{\mu} - \mu\|_{2,k} \lesssim \delta$ .

- 1:  $T' \leftarrow$  Filter  $T$  using Claim A.3 ▷ Preprocessing along the diagonals to ensure Condition 2.7
- 2:  $i \leftarrow 1$
- 3:  $T_i \leftarrow T'$
- 4:  $H \leftarrow$  output of Proposition 3.3 with inputs: corrupted set  $T_i$ , Frobenius threshold  $\kappa = C'\delta^2/\epsilon$  for a large constant  $C'$ , correlation threshold  $\rho = (\delta^2/\epsilon)/k$ , margin threshold  $\tau = (\rho/12)^q$ , sampling parameter  $m \asymp d(\kappa^2/\tau^2 + \log d)$ , and correlation count  $s = d$  ▷ Algorithm 3
- 5: **while**  $T_i \neq \emptyset$  and  $H \neq \text{“}\perp\text{”}$  **do**
- 6:     Get the scores  $f : T_i \rightarrow \mathbb{R}_+$  from Lemma 2.6 with inputs  $T_i, H, \epsilon$ , and  $\delta$
- 7:      $T_{i+1} \leftarrow$  Filter  $T_i$  using the scores  $f$  similar to Algorithm 5
- 8:      $i \leftarrow i + 1$
- 9:     Update  $H$  as above
- 10: **end while**
- 11: **return** the sample mean of  $T_i$

with probability at least 0.6, satisfies (i)  $T''$  is an  $O(\epsilon)$ -contamination of  $S$ , (ii)  $T''$  is an  $O(\epsilon)$  contamination of  $T'$ , where  $T'$  satisfies Condition 2.7, and (iii)  $\|\text{offdiag}(\Sigma_{T''} - \mathbf{I}_d)_H\|_{\text{Fr},k^2} \lesssim \delta^2/\epsilon$ .

We now argue that  $\|\Sigma_{T''} - \mathbf{I}_d\|_{\text{op},k} \lesssim \delta^2/\epsilon$ , and not just the off-diagonal terms. We follow the following inequalities using triangle inequality:

$$\begin{aligned} \|\Sigma_{T''} - \mathbf{I}_d\|_{\text{op},k} &\leq \|\text{diag}(\Sigma_{T''} - \mathbf{I}_d)\|_{\text{op},k} + \|\text{offdiag}(\Sigma_{T''} - \mathbf{I}_d)\|_{\text{op},k} \\ &\lesssim \delta^2/\epsilon + \|\text{offdiag}(\Sigma_{T''} - \mathbf{I}_d)\|_{\text{Fr},k^2} \quad (\text{using Lemma A.4 and } T' \text{ satisfies Condition 2.7}) \\ &\lesssim \delta^2/\epsilon + \delta^2/\epsilon \lesssim \delta^2/\epsilon. \end{aligned}$$

Applying Lemma 2.5, the final output of the algorithm, sample mean of  $T''$  will be  $O(\delta)$  close to the true mean  $\mu$ , implying the correctness of the procedure.

It remains to show the choice of the parameters  $m, s$ , and  $q$  lead to fast runtimes. We take  $\tau = (\rho/12)^q$  and  $s = d$ . Finally, we take  $m \asymp (d^2/s) \cdot (\kappa^2/\tau^2 + \log d) \leq (d \log d)(\kappa^2/\tau^2)$ , which satisfies the parameter constraints in Proposition 3.3 (cf. (11)). Letting  $n = |T|$ , the resulting runtime of a single application of Proposition 3.3 is thus at most

$$\begin{aligned} A &= \left( m + sd^{0.62} + d^{1.62+3\frac{\log(4/\rho)}{\log(1/3\tau)}} \right) \text{poly}(n, \log d, 1/\tau) \\ &\leq \left( d^{1.62+3\frac{\log(4/\rho)}{\log((\rho/4)^q)}} \right) \text{poly}(n, \log d, 1/\rho^q) \\ &\leq \left( d^{1.62+\frac{3}{q}} \right) \text{poly}(n, \log d, k^q, 1/\epsilon^q). \end{aligned}$$

Since there are at most  $n$  iterations, the claim on the total runtime follows. □

## C.2. Proof of Theorem 1.5

We now explain how Theorem C.1 implies Theorem 1.5.

*Proof of Theorem 1.5.* First, observe that by applying Proposition 2.1, the estimation guarantee of Theorem C.1 can be translated from  $\|\cdot\|_{2,k}$  norm into  $\|\cdot\|_2$  norm by hard-thresholding the estimator. We now turn our attention to the sample complexity. Let  $S$  be a set of  $n$  i.i.d. samples from  $\mathcal{N}(\mu, I)$ . If  $n \geq ((k')^2 \log d)/\epsilon^2$ , then Lemma 2.4 implies that  $S$  is  $(\epsilon, \delta, k')$  stable with respect to  $\mu$  for  $\delta \lesssim \epsilon \sqrt{\log(1/\epsilon)}$ . Plugging in the value of  $\delta$  and  $k' \leq \frac{(O(k))^q}{\epsilon^{q-1}}$  in Theorem C.1 yields a

sample complexity of at most  $n \lesssim \frac{(k')^2 \log d}{\epsilon^2} \lesssim (O(k/\epsilon))^{2q} \log d$ . The claim on the runtime in Theorem 1.5 follows from Theorem C.1 along with the bound on the sample complexity.  $\square$

## D. Robust Sparse PCA

In this section, we show that the ideas from fast correlation detection can also be useful for robust sparse PCA. Our main result in this section is the result below:

**Theorem 1.6** (Robust Sparse PCA in Subquadratic Time). *Let  $T$  be an  $\epsilon$ -corrupted set of samples from  $\mathcal{N}(0, \mathbf{I}_d + \eta vv^\top)$  for  $\eta \in (0, 1)$  and a  $k$ -sparse unit vector. There is a randomized algorithm that takes as input the corrupted set  $T$ , contamination rate  $\epsilon$ , sparsity  $k \in \mathbb{N}$ , spike  $\eta$ , and a parameter  $q \in \mathbb{N}$ , and produces an estimate  $\hat{v}$  such that*

- ▶ (Sample Complexity and Error) *If  $n \gtrsim \text{poly}((k^q \log d)/\epsilon^q)$ , then  $\|\hat{v}\hat{v}^\top - vv^\top\|_{\text{Fr}} \lesssim \sqrt{\epsilon \log(1/\epsilon)/\eta}$  with probability at least  $1 - \frac{1}{\text{poly}(d)}$  over the randomness of the samples and the algorithm.*
- ▶ (Runtime) *The algorithm runs in time at most  $d^{1.62 + \frac{3}{q}} \text{poly}(n)$ .*

The result above provides the first subquadratic time algorithm for robust sparse PCA that has error independent of  $d$  and  $k$ . Similar to the literature on robust sparse mean estimation, existing algorithms for robust sparse PCA with similar dimension-independent error took  $\Omega(d^2)$  time. However, the error guarantee of Theorem 1.6 is not optimal: the error guarantee above is  $(\epsilon \log(1/\epsilon))/\eta$ , same as Cheng et al. (2022), as opposed to the near-optimal error of  $(\epsilon^2 \text{polylog}(1/\epsilon))/\eta$  in Balakrishnan et al. (2017); Diakonikolas et al. (2019).

We remark that Theorem 1.6 does not follow directly from Theorem 1.5. In particular, a standard reduction relies on the fact that the mean of the random variable  $\mathbf{Y} - \mathbf{I}_d$  is exactly  $\rho vv^\top$  for  $\mathbf{Y} = xx^\top$ , where  $x \sim \mathcal{N}(0, \mathbf{I} + \rho vv^\top)$ , i.e., (robust) sparse PCA reduces to (robust) sparse mean estimation. However,  $\mathbf{Y}$  is a  $d^2$ -dimensional object and thus a naive application of Theorem 1.5 would yield a super-quadratic runtime, in fact,  $(d^2)^{1.62} = \Omega(d^3)$ . Moreover, the covariance of  $\mathbf{Y}$  is not isotropic and thus it is unclear if samples from  $Y$  would satisfy the stability condition from Definition 2.3.

We show below that a more whitebox analysis of the main ideas from Theorem 1.5, in particular, Proposition 3.3 can yield a subquadratic runtime.

**Organization** This section is organized as follows: We list the deterministic conditions for robust PCA in Appendix D.1. Appendix D.2 contains the results pertaining to the certificate lemma, filtering, and dense estimation algorithm for sparse PCA. Finally, the algorithm and its proof is given in Appendix D.3.

### D.1. Deterministic Conditions

For a set  $T$ , we use the notation  $\overline{\Sigma}_T$  to denote the second moment matrix  $\mathbb{E}_T[xx^\top]$ —not to be confused with the covariance matrix  $\Sigma_T$ .

**Definition D.1** (Stability Condition for PCA). For the contamination rate  $\epsilon \in (0, 1/2)$ , error parameter  $\gamma \geq \epsilon$ , spike strength  $\eta \in \mathbb{R}_+$ , and sparsity  $k \in \mathbb{N}$ , we say a set  $S \subset \mathbb{R}^d$  is  $(\epsilon, \gamma, k, \eta)$ -pca-stable with respect to a  $k$ -sparse unit vector  $v \in \mathbb{R}^d$  and spike strength  $\eta$  if

1. For any subset  $S' \subseteq S$  with  $|S'| \geq (1 - \epsilon)|S|$ :  $\|\overline{\Sigma}_{S'} - (\mathbf{I}_d + \eta vv^\top)\|_{\text{Fr}, k^2} \leq \gamma$
2. For all subsets  $H \subset [d]$  with  $|H| \leq k$ , the set  $\{(x)_H : x \in S\}$  is  $(\epsilon, \gamma)$ -covariance stable in the sense of Diakonikolas & Kane (2023, Definition 4.5) (with respected to an appropriate flattening of  $\mathbf{I}_{|H|}$ ).

The second stability condition allows us to perform covariance estimation in Frobenius norm (and hence stronger than principal component analysis) but only when the support is known.

The following result lists the sample complexity of Definition D.1.

**Lemma D.2.** *Let  $S$  be a set of  $n$  i.i.d. samples from  $\mathcal{N}(0, \mathbf{I} + \eta vv^\top)$  for  $\eta \in (0, 1)$  and a  $k$ -sparse unit vector  $v$ . Then if  $n \gtrsim \text{poly}(k, \log(d/\delta), 1/\epsilon)$ , then  $S$  is  $(\epsilon, \gamma, k, \eta)$ -pca-stable with respect to  $v$  and  $\eta$  for  $\gamma = O(\epsilon \log(1/\epsilon))$ .*

*Proof.* We sketch the argument here. The first condition is identical to Cheng et al. (2022, Definition 2.2) and Cheng et al. (2022, Lemma 4.2) shows establishes a sample complexity of  $\frac{k^2 \log(d/\delta)}{\epsilon^2}$ .

For the second condition, we do an admittedly loose analysis. Fixing a subset  $H$ , it amounts to showing a stability condition for a  $|H|$ -dimensional Gaussian distribution, for which [Diakonikolas & Kane \(2023, Proposition 4.2\)](#) establishes a sample complexity of  $\text{poly}(k \log(1/\delta)/\epsilon)$  for the set to be  $(\epsilon, \gamma)$ -covariance stable for  $\gamma = O(\epsilon \log(1/\epsilon))$  with probability  $1 - \delta$ .<sup>15</sup> Since there are at most  $d^k$  possible choices of  $H$ , a union bound shows that the second condition in [Definition D.1](#) holds with probability  $1 - \delta$  with sample complexity  $\text{poly}(k \log(d^k/\delta)/\epsilon) = \text{poly}(k \log(d/\delta)/\epsilon)$ .  $\square$

## D.2. Filters, Certificates, and Dense Estimation

[Appendix D.2.1](#) records the guarantee of the dense estimation algorithm. [Appendix D.2.2](#) includes the certificate lemma for sparse PCA and [Appendix D.2.3](#) includes the filter lemma. Finally, [Appendix D.2.4](#) contains the preprocessing condition on the covariance matrix.

### D.2.1. DENSE ESTIMATION ALGORITHM

We shall use the following algorithm to estimate the spike when we have identified the support of the spike.

**Lemma D.3** (Dense Covariance Estimation Algorithm; Implication of [Diakonikolas & Kane \(2023, Theorem 4.6\)](#)). *Let  $\epsilon \in (0, \epsilon_0)$  for a small absolute constant  $\epsilon_0$ . Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \gamma, k, \eta)$ -pca-stable with respect to  $v$  (cf. [Definition D.1](#)) and  $\eta \in (0, 1)$ . Let  $H$  be a  $k$ -sparse subset of  $[d]$ .*

*There exists an algorithm  $\mathcal{A}$  that takes as inputs corrupted set  $T$ , contamination rate  $\epsilon$ , error parameter  $\gamma$ , spike  $\eta$ , and the sparse support set  $H$  and outputs a  $k$ -sparse vector  $u$ , supported on  $H$ , such that  $\|(\mathbf{I}_d + \eta uu^\top)_H - (\mathbf{I}_d + \eta vv^\top)_H\| \lesssim \gamma$ . Moreover, the algorithm runs in time  $d \text{poly}(k, |T|/\epsilon)$ .*

*Proof.* Let  $T' \subset \mathbb{R}^k$  and  $S' \subset \mathbb{R}^k$  denote the projections of  $T$  and  $S$ , respectively, on  $H$ . Let  $\Sigma' = (\mathbf{I}_d + \eta vv^\top)_H = \mathbf{I}_{|H|} + \eta' v' v'^\top$  for  $v' = (v)_H / \|(v)_H\|_2$  and  $\eta' = \eta \|(v)_H\|_2^2$ .

[Diakonikolas & Kane \(2023, Theorem 4.6\)](#) gives an estimate  $\widehat{\Sigma}$  such that  $\|\widehat{\Sigma} - \Sigma'\|_{\text{Fr}} \lesssim \gamma \|\Sigma'\|_{\text{op}} \lesssim \gamma$ , which can be calculated in time  $d \text{poly}(k|T|/\epsilon)$ ; here  $\|\cdot\|_{\text{op}}$  denotes the operator norm of a matrix. That is,  $\|(\widehat{\Sigma} - \mathbf{I}) - \eta' v' v'^\top\|_{\text{Fr}} \lesssim \gamma$ . Letting  $\mathbf{A}$  be the best one-rank approximation of  $(\widehat{\Sigma} - \mathbf{I})$  in the Frobenius norm, we see that

$$\|\mathbf{A} - \eta' v' v'^\top\|_{\text{Fr}} \leq \|\mathbf{A} - (\widehat{\Sigma} - \mathbf{I})\|_{\text{Fr}} + \|(\widehat{\Sigma} - \mathbf{I}) - \eta' v' v'^\top\|_{\text{Fr}} \leq 2\|(\widehat{\Sigma} - \mathbf{I}) - \eta' v' v'^\top\|_{\text{Fr}} \lesssim \gamma,$$

where the second inequality follows from the fact that  $\mathbf{A}$  is the best rank-one approximation. Moreover, the symmetry of  $\mathbf{A}$  implies that  $\mathbf{A}$  must be of the form  $\lambda' u' u'^\top$  for a unit vector  $u'$  and  $\lambda' \in \mathbb{R}$ . We thus obtain

$$\|\lambda' u' u'^\top - \eta' v' v'^\top\|_{\text{Fr}} \lesssim \gamma.$$

Consequently,  $|\lambda' - \eta'| \lesssim \gamma$  by Weyl's inequality. If  $\lambda' \leq 0$ , we set  $u = 0$ . Otherwise, we set  $u = \sqrt{\lambda'/\eta'} u'$ . We now calculate the approximation error. If  $\lambda' \geq 0$ , then the resulting error in our estimate is  $\|(\mathbf{I}_d + \eta uu^\top)_H - (\mathbf{I}_d + \eta vv^\top)_H\|_{\text{Fr}} = \|\lambda' uu^\top - \eta' v' v'^\top\|_{\text{Fr}} \lesssim \gamma$  by the approximation guarantee. If  $\lambda' < 0$ , then  $\eta'$  must also be  $O(\gamma)$  because  $\eta' \leq \lambda' + |\lambda' - \eta'| \lesssim \gamma$ . In this case, the approximation error is  $\eta'$ , which is also  $O(\gamma)$  as required.  $\square$

### D.2.2. CERTIFICATE LEMMA FOR SPARSE PCA

The following result shows that if the covariance matrix looks roughly isotropic on a subset of coordinates  $H^c$ , then the spike vector  $v$  places most of its mass on the complement,  $H$ . The benefit of this stopping condition is that it does not depend on the unknown spike vector  $v$ .

**Lemma D.4** (Sparse Certificate Lemma for PCA). *Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \gamma, k, \eta)$ -pca-stable with respect to a  $k$ -sparse unit vector  $v \in \mathbb{R}^d$  and spike strength  $\eta \in (0, 1)$ . Let  $H \subset [d]$  be such that  $\|(\widehat{\Sigma}_T - \mathbf{I})_{H^c}\|_{\text{op}, k} = O(\gamma)$ , then  $\|vv^\top - (v)_H (v)_H^\top\|_{\text{Fr}}^2 = O((\gamma/\eta) + (\gamma/\eta)^2)$ .*

*Proof.* Let  $z$  be the unit vector along  $(v)_{H^c}$ , which is at most  $k$ -sparse. Since  $z$  is supported on  $H^c$ , we obtain

$$|z^\top (\widehat{\Sigma}_T - \mathbf{I}) z| \lesssim \gamma,$$

<sup>15</sup>Although the proof of [Diakonikolas & Kane \(2023, Proposition 4.2\)](#) does not explicitly write the dependence on  $\delta$ , but it is immediate from the proof.

implying that  $z^\top \bar{\Sigma}_T z = 1 \pm O(\gamma)$ . Using the stability of  $S$ , we see  $z^\top \bar{\Sigma}_T z \geq (1 - \epsilon)z^\top \bar{\Sigma}_{S \cap T} z \geq (1 - \epsilon)(z^\top (\mathbf{I} + \eta v v^\top) z - \gamma)$ . Defining  $\eta' := \eta(v^\top z)^2 = \eta \|v_{H^c}\|_2^2$ , we have that  $z^\top \bar{\Sigma}_T z \geq (1 - \epsilon)(1 + \eta' - \gamma)$ . Combining this with the aforementioned upper bounds on  $z^\top \bar{\Sigma}_T z$ , we obtain that  $(1 - \epsilon)(1 + \eta' - \gamma) \leq 1 + O(\gamma)$ . Thus,  $\eta' = O(\gamma + \epsilon) = O(\gamma)$  by using  $\epsilon \leq \gamma$ . Therefore,  $\|(v)_{H^c}\|_2^2 = \eta'/\eta \lesssim \gamma/\eta$ . Finally, the triangle inequality implies that

$$\|v v^\top - (v)_{H^c} (v)_{H^c}^\top\|_{\text{Fr}} \leq \|(v)_{H^c}\|_2^2 + 2\|(v)_{H^c}\|_2 \lesssim \max\left(\sqrt{\gamma/\rho}, (\gamma/\rho)\right).$$

□

### D.2.3. FILTER FOR SPARSE PCA

**Lemma D.5** (Sparse PCA Filter). *Let  $\epsilon \in (0, \epsilon_0)$  for a small absolute constant  $\epsilon_0$  and let  $C$  be a large enough absolute constant. Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \gamma, k, \eta)$ -stable with respect to  $v$  (cf. Definition D.1) and  $\eta \in (0, 1)$ . Let  $H \subset [d]$  be such  $\|(\Sigma_T - \mathbf{I} - \eta v v^\top)_H\|_{\text{Fr}} = \lambda$  for  $\lambda \geq 8C\gamma$  and  $|H| \leq k$ .*

*Then there exists an algorithm  $\mathcal{A}$  that takes  $T$ ,  $H$ ,  $\epsilon$ ,  $\gamma$ , and  $\eta$  and returns a score mapping  $f : T \rightarrow \mathbb{R}_+$  such that the sum of inliers' scores is less than outliers':  $\sum_{x \in S \cap T} f(x) \leq \sum_{x \in T \setminus S} f(x)$  and  $\max_{x \in T} f(x) > 0$ . Moreover, the algorithm runs in time  $\text{dpoly}(k, |T|)$ .*

*Proof.* We use the same ideas from the proof of Lemma 2.6 and similarly assume that  $T$  and  $S$  already correspond to the projected coordinates. The challenge in applying the idea as is lie in the uncertainty about  $v$ . We thus use Lemma D.3 to first estimate  $v$  using the returned vector  $u$  which satisfies that  $\|\eta u u^\top - \eta v v^\top\|_{\text{Fr}} \lesssim \gamma$ .<sup>16</sup> We give the algorithm below.

---

#### Algorithm 8 PCA Filter

---

- 1:  $T \leftarrow \{(x)_H : x \in T\}$  ▷ Projection onto  $H$
  - 2:  $u \leftarrow$  be output of dense covariance estimation algorithm Lemma D.3
  - 3: Let  $\mathbf{A}$  be the matrix  $(\Sigma_T - \mathbf{I} - \eta u u^\top) / \|\Sigma_{T_1} - \mathbf{I} - \eta u u^\top\|_{\text{Fr}}$
  - 4: Define  $g(x) := (x - \mu_T)^\top \mathbf{A} (x - \mu_T) - \text{tr}(\mathbf{A})$
  - 5: Define  $f(x)$  to be  $g(x)$  if  $g(x) \geq \Omega(\gamma/\epsilon)$  otherwise 0
  - 6: Return  $f$
- 

Let  $\lambda = \|(\Sigma_T - \mathbf{I} - \eta v v^\top)_H\|_{\text{Fr}}$  and define  $\mathbf{A} = (\Sigma_T - \mathbf{I} - \eta u u^\top) / \|\Sigma_{T_1} - \mathbf{I} - \eta u u^\top\|_{\text{Fr}}$ . Define the function  $g(x) := x^\top \mathbf{A} x - \langle \mathbf{I} + \eta u u^\top, \mathbf{A} \rangle$ . Computing the average of  $g(x)$  over the  $T$ , we obtain

$$\begin{aligned} \mathbb{E}_T[g(x)] &= \mathbb{E}_T[\langle x x^\top - \mathbf{I} - \eta u u^\top, \mathbf{A} \rangle] = \langle \bar{\Sigma}_T - \mathbf{I} - \eta u u^\top, \mathbf{A} \rangle = \|\bar{\Sigma}_T - \mathbf{I} - \eta u u^\top\|_{\text{Fr}} \\ &\geq \|\bar{\Sigma}_T - \mathbf{I} - \eta v v^\top\|_{\text{Fr}} - \|\eta u u^\top - \eta v v^\top\|_{\text{Fr}} = \lambda - O(\gamma) \geq 3\lambda/4, \end{aligned} \quad (12)$$

where we use that  $\lambda \gtrsim \gamma$ . Using the stability of  $S$  (observe that on the projected set, the  $\|\cdot\|_{\text{Fr}, k^2}$  norm becomes the usual Frobenius norm), for any large subset of  $S' \subset S$  with  $|S'| \geq (1 - \epsilon)|S|$ , the closeness between  $u$  and  $v$  implies the following:

$$\begin{aligned} |\mathbb{E}_{S'}[g(x)]| &= |\mathbb{E}_{S'}[\langle x x^\top - \mathbf{I} - \eta u u^\top, \mathbf{A} \rangle]| = |\langle \bar{\Sigma}_{S'} - \mathbf{I} - \eta u u^\top, \mathbf{A} \rangle| \\ &= |\langle \bar{\Sigma}_{S'} - \mathbf{I} - \eta v v^\top, \mathbf{A} \rangle + \langle \eta v v^\top - \eta u u^\top, \mathbf{A} \rangle| \leq \|\bar{\Sigma}_{S'} - \mathbf{I} - \eta v v^\top\|_{\text{Fr}} + \eta \|v v^\top - u u^\top\|_{\text{Fr}} \\ &\leq \gamma + O(\gamma) \leq C\gamma \end{aligned} \quad (13)$$

for a large enough absolute constant  $C$ . Now define  $f(x) := g(x) \mathbf{1}_{g(x) \geq 3C\gamma/\epsilon}$ , i.e.,  $f(x)$  is equal to  $g(x)$  if  $g(x) \geq 3C\gamma/\epsilon$  and 0 otherwise. By Claim A.1,

$$\sum_{x \in S \cap T} f(x) \leq \sum_{x \in S} f(x) \leq 3C\gamma|S| \leq |T|\lambda/4, \quad (14)$$

where we use that  $\lambda \gtrsim C\gamma$ .

<sup>16</sup>Observe that  $v$  here corresponds to  $(v)_H$  because of the projection to  $H$ . Hence  $v$  is no longer a unit vector.

On the other hand,  $\mathbb{E}_{x \in T \setminus S}[f(x)]$  must be large argued as argued below. Applying (13) to  $T \cap S$ , which is of size at least  $(1 - \epsilon)|T| = (1 - \epsilon)|S|$ , we obtain

$$\begin{aligned}
 \sum_{x \in T \setminus S} f(x) &\geq \sum_{x \in T \setminus S} (g(x) - 3C\gamma/\epsilon) = \left( \left( \sum_{x \in T} g(x) \right) - \left( \sum_{x \in T \cap S} g(x) \right) \right) - (|T \setminus S|3C\gamma/\epsilon) \\
 &\geq (3\lambda|T|/4) - (|T \cap S|(C\gamma)) - \epsilon|T|(3C\gamma/\epsilon) \\
 &\geq (3\lambda|T|/4) - (4|T|C\gamma) \\
 &\geq \lambda|T|/2,
 \end{aligned} \tag{15}$$

where we use (12) and  $\lambda \gtrsim \gamma$ . The desired conclusions follow from (14) and (15).  $\square$

#### D.2.4. PREPROCESSING CONDITION FOR SPARSE PCA

Similar to robust sparse mean estimation, our algorithm tracks the contribution to the diagonal terms and off-diagonal terms separately. The following conditions mirrors Condition 2.7.

**Condition D.6** (Idealistic Condition for PCA). Let  $T$  be an  $\epsilon$ -corrupted version of  $S$ , where  $S$  is  $(\epsilon, \gamma, k, \eta)$ -stable with respect to a  $k$ -sparse unit vector  $v$  and spike strength  $\eta$  (cf. Definition D.1). We have an  $H_1 \subset [d]$  and  $|H_1| \lesssim k^2$  such that  $T$  satisfies  $\|\text{diag}(\overline{\Sigma}_T - \mathbf{I}_d)_{H_1^c}\|_{\text{Fr}, k^2} \lesssim \gamma$ .

The next result gives an efficient algorithm to ensure the above condition:

**Claim D.7.** Let  $\epsilon \in (0, \epsilon_0)$  and  $\gamma \in (0, \gamma_0)$  for small constants  $\epsilon_0 \in (0, 1/2)$ ,  $\gamma_0 \in (0, 1)$ . Let sparsity  $k \in \mathbb{N}$ . Let  $C$  be a large enough constant and  $T$  be an  $\epsilon$ -corrupted set  $S$  where  $S$  is  $(C\epsilon, \gamma, k', \eta)$ -pca-stable with respect to an unknown  $k$ -sparse unit vector  $v \in \mathbb{R}^d$ ,  $\eta \in (0, 1)$ ,  $\gamma \geq \epsilon$ , and  $k' = Ck^2$  for a large enough constant  $C' > 0$ . There is a randomized algorithm  $\mathcal{A}$  that takes as input the corrupted set  $T$ , contamination rate  $\epsilon$ , sparsity  $k \in \mathbb{N}$ , and a parameter  $q \in \mathbb{N}$ , and returns a set  $T' \subset T$  and  $H_1 \subset [d]$  in time  $O(d\text{poly}(k, |T|))$  such that with probability 0.9

1.  $T'$  is an  $O(\epsilon)$ -contamination of  $S$ .
2. Each diagonal entry of  $\overline{\Sigma}_{T'} \in [1/2, 4]$ .
3.  $T'$  and  $H_1$  satisfy Condition D.6, i.e.,  $|H_1| \lesssim k^2$  and  $\|\text{diag}(\overline{\Sigma}_{T'} - \mathbf{I}_d)_{H_1^c}\|_{\text{Fr}, k^2} \lesssim \gamma$ .

*Proof.* We will filter the set using Lemma D.5 following the template of Algorithm 5 until the second and the third conditions are met.

Starting with the second condition, the lower bound on the diagonal entries follows by the fact that  $T$  contains an  $\epsilon$ -fraction of  $S \cap T$  and diagonal entries of  $\overline{\Sigma}_{S \cap T}$  are at least  $1 - \gamma$ . Since  $\epsilon$  and  $\gamma$  are small enough,  $(1 - \epsilon)(1 - \gamma) \geq 1/2$ . We now focus on establishing the upper bound, for each coordinate  $i \in [d]$ , the true variance is at most  $(1 + \eta\|v\|_\infty^2) \leq 2$ , where  $\|\cdot\|_\infty$  denotes the  $\ell_\infty$  norm. Thus, if for any coordinate  $i \in [d]$ , the empirical variance of  $T$  is larger than 3, which is bigger than  $1 + \eta + \lambda$ ,  $H := \{i\}$  satisfies the condition of Lemma D.5. Thus, we can filter points until all the empirical variances are less than 3 following the template of Theorem A.2.

We now turn our attention to the third condition. For a vector  $x \in \mathbb{R}^d$ , let the function  $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$  be the coordinate-wise square of its input. Thus, for  $X \sim \mathcal{N}(0, \mathbf{I} + \rho vv^\top)$ , we have that  $\mu := \mathbb{E}[g(X)] = u + \eta g(v)$ , where  $u \in \mathbb{R}^d$  denotes the all ones vector. Therefore, the required condition in Condition D.6 can also be written equivalently as  $\|(\mu_{f(T)} - u)_{H_1^c}\|_{2, k^2} = O(\gamma)$ . Let  $f(T)$  and  $f(S)$  denote the sets transformed by  $f$ , i.e.,  $f(T) := \{f(x) : x \in T\}$  and  $f(S) := \{f(x) : x \in S\}$ . We first compute  $\mu_{f(T)} - u$ , which takes  $O(d|T|)$  time. Let  $J \subset [d]$  denote the coordinates  $i$  for which the  $i$ -th coordinate of  $\mu_{f(T)} - u$  is bigger than  $\gamma/k$  in absolute value. Let  $J_*$  denote the support of the sparse spike vector  $v$ .

Consider the case when  $J$  is a large enough set:  $|J| > k'$ . Then let  $H \subset J$  be any subset of size  $k'$ . Then  $|H \setminus J_*| > |H| - |J_*| = k' - k \geq k'/2$ . And thus we have that

$$\begin{aligned}
 \|(\Sigma_T - \mathbf{I}_d - \eta vv^\top)_H\|_{\text{Fr}, k'^2} &\geq \|(\Sigma_T - \mathbf{I}_d - \eta vv^\top)_{H \setminus J_*}\|_{\text{Fr}, k'^2} \geq \\
 &\geq \|(\Sigma_T - \mathbf{I}_d)_{H \setminus J_*}\|_{\text{Fr}, k'^2}
 \end{aligned}$$

$$\begin{aligned} &\geq \|\text{diag}((\Sigma_T - \mathbf{I}_d)_{H \setminus J_*})\|_{\text{Fr}, k'^2} \\ &\geq (\gamma/k) \sqrt{(k'/2)} \gtrsim \gamma, \end{aligned}$$

since  $k' \gtrsim k^2$ . Thus, we have obtained a set  $H$  that satisfies the guarantees of Lemma D.5, which allows us to filter using the template of Algorithm 5.

If on the other hand,  $J$  happens to be small, then we return  $H_1 = J$  and since  $\mu_{f(T)}$  is  $(\gamma/k)$ -close to  $u$  (that is, 1) on  $H_1^c$ , we obtain

$$\|\text{diag}(\Sigma_T - \mathbf{I}_d)_{H_1^c}\|_{\text{Fr}, k^2} = \|(\mu_{f(T)} - u)_{H_1^c}\|_{2, k^2} \leq k\gamma/k \leq \gamma.$$

□

Similar to Lemma A.4, the next result shows that all large subsets of a set satisfying Claim D.7 are close to identity in the sparse operator norm.

**Lemma D.8.** *Let  $C$  be a large enough constant  $C > 0$  and  $k, k' \in \mathbb{N}$ . Let  $T'' \subset T'$  be two  $O(\epsilon)$ -contamination of  $S$  such that  $S$  is an  $(C\epsilon, \gamma, k', \eta)$ -pca-stable with respect to  $v$  and  $\rho \in (0, 1)$ . Let  $H \subset [d]$  be a small subset  $|H| \leq k$  such that  $\|\text{diag}(\overline{\Sigma}_{T'} - \mathbf{I}_d)_{H_1^c}\|_{\text{Fr}, k^2} \lesssim \gamma$ . Then  $\|\text{diag}(\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_1^c}\|_{\text{op}, k} \lesssim \gamma$ .*

*Proof.* Our proof strategy will be similar to that of Lemma A.4, and we refer the reader to the proof of Lemma A.4 for more details. Starting with the lower bound on the sparse eigenvalues of  $\overline{\Sigma}_{T''}$ , we note that for any  $k'$ -sparse unit vector  $u$ , the stability condition applied to  $S \cap T''$  applies

$$u^\top \overline{\Sigma}_{T''} u \geq \frac{|S \cap T''|}{|T''|} u^\top \Sigma_{S \cap T''} u \geq (1 - O(\epsilon))(u^\top (I + \rho v v^\top) u - \gamma) \geq (1 - O(\epsilon))(1 - O(\gamma)) \geq 1 - O(\gamma)$$

since  $\epsilon \leq \gamma$ . Applying this inequality for 1-sparse unit vectors  $u$ , we obtain that for any unit vector  $z$ :  $-z^\top \text{diag}(\overline{\Sigma}_T - \mathbf{I}) z \leq C\gamma$  for a large constant  $C > 0$ .

Turning towards upper bound, we proceed as follows: for any 1-sparse unit vector  $u$  supported on  $H_1^c$ ,

$$u^\top \text{diag}(\Sigma_{T''}) u = u^\top \Sigma_{T''} u = \frac{1}{|T''|} \sum_{x \in T''} (u^\top x)^2 \leq (1 + O(\epsilon)) u^\top \Sigma_{T'} u \leq (1 + O(\epsilon))(1 + O(\gamma)) = 1 + O(\gamma).$$

Overall, we obtain the desired guarantee of  $\|\text{diag}((\Sigma_{T''} - \mathbf{I})_{H_1^c})\|_{\text{op}, k} = O(\gamma)$ .

□

### D.3. Subquadratic Time Algorithm For Sparse PCA

We now establish the main technical result of this section:

**Theorem D.9** (Subquadratic Time Algorithm for Robust Sparse PCA under Stability). *Let  $\epsilon \in (0, \epsilon_0)$  and  $\gamma \in (0, \gamma_0)$  for small constants  $\epsilon_0 \in (0, 1/2)$ ,  $\gamma_0 \in (0, 1)$ . Let  $C$  be a large enough constant and  $k \in \mathbb{N}$  be the sparsity parameter and  $q \in \mathbb{N}$  the correlation decay parameter with  $q \geq 3$ . Let  $T$  be an  $\epsilon$ -corrupted set  $S$  where  $S$  is  $(C\epsilon, \gamma, Ck', \eta)$ -pca-stable with respect to an unknown  $k$ -sparse unit vector  $v \in \mathbb{R}^d$ ,  $\eta \in (0, 1)$ , and  $\gamma \geq \epsilon$  for  $k' = \left(\frac{C^q k^{2q}}{\gamma^{2q-2}}\right)$ .*

*There is a randomized algorithm  $\mathcal{A}$  that takes as input the corrupted set  $T$ , contamination rate  $\epsilon$ , stability parameter  $\gamma$ , sparsity  $k \in \mathbb{N}$ , and a parameter  $q \in \mathbb{N}$ , and produces an estimate  $\widehat{v}$  such that*

- (Error) Then  $\|\widehat{v}\widehat{v}^\top - v v^\top\|_{\text{Fr}} \lesssim \sqrt{\frac{2}{\eta}}$  with high probability over the randomness of the samples and the algorithm.
- (Runtime) The algorithm runs in time at most  $d^{1.62 + \frac{3}{q}} \text{poly}(|T| \log d, (k/\epsilon)^q)$ .

The complete algorithm is given below: We now present the proof of its correctness:

**Algorithm 9** Robust Sparse PCA Algorithm

**Require:** corruption rate  $\epsilon \in (0, 1)$ , stability parameter  $\gamma \in (0, 1)$ , corrupted set  $T \subset \mathbb{R}^d$ , correlation-threshold  $\rho \in (0, 1)$ , correlation decay  $q \in \mathbb{N}$ , sparsity  $k \in \mathbb{N}$ , spike strength  $\eta \in (0, 1)$ , sampling parameter  $m \in \mathbb{N}$ . We require  $T$  to be an  $\epsilon$ -corrupted version of an  $(C\epsilon, \delta, Ck', \eta)$  stable set with respect to  $\mu$  for  $k' := \left(\frac{C^q k^{2q}}{\gamma^{2q-2}}\right)$  and  $\gamma \ll 1$ .

**Ensure:**  $\hat{v} \in \mathbb{R}^d$  such that, with high probability,  $\|\hat{v}\hat{v}^\top - vv^\top\|_{\text{Fr}} \lesssim \gamma$ .

- 1:  $H_1, T' \leftarrow$  be the outputs of Claim D.7 on  $T$   
 $\triangleright T'$  is an  $O(\epsilon)$ -contamination of  $S$ ,  $0.5\mathbf{I}_d \preceq \text{diag}(\overline{\Sigma}) \preceq 2\mathbf{I}_d$ ,  $|H_1| \lesssim k^2$  and  $\|\text{diag}(\overline{\Sigma}_T - \mathbf{I}_d)_{H_1^c}\|_{\text{Fr}, k^2} \lesssim \sqrt{\gamma/\eta}$
- 2:  $i \leftarrow 1$
- 3:  $T_i \leftarrow T'$ .
- 4:  $H \leftarrow$  output of Proposition 3.3 with inputs: corrupted set  $T_i$ , Frobenius threshold  $\kappa = \Theta(\gamma)$ , correlation threshold  $\rho = \gamma/k$ , margin threshold  $\tau = (\rho/12)^q$ , sampling parameter  $m \asymp d(\kappa^2/\tau^2 + \log d)$ , and correlation count  $s = d$   
 $\triangleright$  Algorithm 3
- 5: **while**  $T_i \neq \emptyset$  and  $H \neq \perp$  and  $|H| \leq \kappa^2/\rho^2$  **do**
- 6:     Get the scores  $f : T_i \rightarrow \mathbb{R}_+$  from Lemma D.5 with inputs  $T_i, H, \epsilon$ , and  $\delta$
- 7:      $T_{i+1} \leftarrow$  Filter  $T_i$  using the scores  $f$ .
- 8:      $i \leftarrow i + 1$
- 9:     Update  $H$  as above
- 10: **end while**
- 11: Let  $H_{\text{end}} \leftarrow H \cup H_1$
- 12:  $u \leftarrow$  output of Lemma D.3 with inputs  $T_i, \epsilon, \gamma, \eta$ , and  $H_{\text{end}}$

*Proof of Theorem D.9.* We can assume that  $\gamma \leq \eta$ , otherwise we can simply output any unit vector, whose error will be at most  $O(1) = O(\sqrt{\gamma/\eta})$ .

We will use the same template of filtering-based algorithms from Theorem A.2, with the filtering subroutines provided by Lemma D.5.

The first step of Algorithm 9 is the preprocessing step from Claim D.7, which takes at most  $\tilde{O}(d \text{poly}(k, |T|))$  time and removes not too many inliers with high probability. In particular, the diagonal entries of  $\overline{\Sigma}_{T'}$  lie in  $[1/2, 4]$ ; Moreover, all the subsequent sets  $T_i$ 's will continue to satisfy this property by Lemma D.8 as long as we have removed at most  $O(\epsilon)$  fraction of points (since  $\gamma$  is small enough). We shall use the output  $H_1$  generated by Claim D.7 in the end.

We will use a fine-grained result from Proposition 3.3. Observe that Proposition 3.3 can be used not just for the covariance  $\Sigma_T$  but also for  $\overline{\Sigma}_T$ , i.e., without centering, which is what we will use in this proof. We shall invoke Proposition 3.3 to identify all the coordinate pairs with correlation larger than  $\rho = \gamma/k$ . We then set  $\tau = (\rho/12)^q$  and use the Frobenius threshold  $\kappa = C'\gamma$  for a constant  $C > 100$ . Defining  $k' := \kappa^2/\tau^2$  and  $k'' := \kappa^2/\rho^2$ , we note that  $k' \geq k'' = C'\gamma^2/(\gamma/k)^2 = C'k^2 \geq 2k^2$ . We will show that the set  $H$  allows us to filter points using Lemma D.3. The proof of Proposition 3.3 reveals that it returns  $H$  such that either

(Case I)  $|H| = k' = \kappa^2/\tau^2$  and for each coordinate in  $i \in H$ , there exists a  $j \in H$  such that  $(i, j)$  is  $\tau$ -correlated

Let  $T''$  be the current iterate of the corrupted data set. Since  $\rho vv^\top$  is a  $k^2$ -sparse matrix and  $\text{offdiag}(\overline{\Sigma}_T - \mathbf{I}_d)$  has at least  $k'$  entries with magnitude at least  $\Theta(\tau)$ <sup>17</sup>, their difference must have at least  $k' - k^2 \geq k'/2 = \kappa^2/\tau^2$  entries with magnitude  $\Theta(\tau)$ . Thus, the sparse Frobenius norm of their difference must be large, i.e.,

$$\|(\overline{\Sigma}_{T''} - \mathbf{I}_d - \rho vv^\top)_H\|_{\text{Fr}} \gtrsim \kappa \gtrsim \gamma.$$

Given such an  $H$ , we filter outliers using Lemma D.3.

(Case II)  $|H| \geq \kappa^2/\rho^2$  and for each coordinate in  $i \in H$ , there exists a  $j \in H$  such that  $(i, j)$  is  $\rho$ -correlated

By the same argument as above, the matrix  $\overline{\Sigma}_{T''} - \mathbf{I}_d - \rho vv^\top$  has at least  $k'' - k^2 \geq k''/2 = \kappa^2/(2\rho^2)$  entries with magnitude  $\Theta(\rho)$ . Thus,  $\|(\overline{\Sigma}_{T''} - \mathbf{I}_d - \rho vv^\top)_H\|_{\text{Fr}} \gtrsim \kappa \gtrsim \gamma$ . Again, we filter outliers using Lemma D.3.

<sup>17</sup>Recall that each diagonal entry of  $\overline{\Sigma}_{T''}$  is  $\Theta(1)$  and thus  $\tau$ -correlation implies that the corresponding entry in  $\overline{\Sigma}_{T''}$  is also  $\Theta(\tau)$ .

(Case III)  $|H| \leq \kappa^2/\rho^2$  and no coordinate pair in  $H^{\mathfrak{C}}$  is  $\rho$ -correlated.

Let  $T''$  be the current iterate of the corrupted data set. Since  $H \subset H_{\text{end}}$  and the coordinates in  $H^{\mathfrak{C}}$  are at most  $\rho$ -correlated, it follows that the coordinates in  $H_{\text{end}}^{\mathfrak{C}}$  are also at most  $\rho$ -correlated.

$$\left\| \text{offdiag} \left( (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_{\text{end}}^{\mathfrak{C}}} \right) \right\|_{\text{Fr}, k^2} \leq \left\| \text{offdiag} \left( (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H^{\mathfrak{C}}} \right) \right\|_{\text{Fr}, k^2} \lesssim \rho k \lesssim \gamma. \quad (16)$$

We now want to combine the above guarantee on the closeness along the offdiagonals with the guarantee on the closeness along the diagonals from  $H_1$ . In particular,  $H_{\text{end}}$  satisfies that

$$\left\| \text{diag} \left( (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_{\text{end}}^{\mathfrak{C}}} \right) \right\|_{\text{op}, k} \leq \left\| \text{diag} \left( (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_1^{\mathfrak{C}}} \right) \right\|_{\text{op}, k} \lesssim \gamma, \quad (17)$$

where  $T''$  has small sparse operator norm because it is close to the preprocessed set  $T'$  and thus Lemma D.8 is applicable. Combining (16) and (17), we obtain

$$\left\| (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_{\text{end}}^{\mathfrak{C}}} \right\|_{\text{op}, k} \leq \left\| \text{diag} \left( (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_{\text{end}}^{\mathfrak{C}}} \right) \right\|_{\text{op}, k} + \left\| \text{offdiag} \left( (\overline{\Sigma}_{T''} - \mathbf{I}_d)_{H_{\text{end}}^{\mathfrak{C}}} \right) \right\|_{\text{op}, k} \lesssim \gamma. \quad (18)$$

The size of  $H_{\text{end}}$  is at most  $O(k^2) + k' \leq 2k'$ . Since  $S$  satisfies stability with  $2k'$ , Lemma D.4 then implies that the spike  $v$  is mostly contained in  $H_{\text{end}}$ . By invoking the dense PCA algorithm on  $H_{\text{end}}$ , Lemma D.3 estimates the spike  $(v)_{H_{\text{end}}}(v)_{H_{\text{end}}}^{\top}$  with  $uu^{\top}$ , with  $u$  supported on  $H_{\text{end}}$ . Combining, we obtain:

$$\begin{aligned} \|vv^{\top} - uu^{\top}\|_{\text{Fr}} &\leq \|vv^{\top} - (v)_{H_{\text{end}}}(v)_{H_{\text{end}}}^{\top}\|_{\text{Fr}} + \|uu^{\top} - (v)_{H_{\text{end}}}(v)_{H_{\text{end}}}^{\top}\|_{\text{Fr}} \\ &\lesssim \sqrt{\gamma/\eta} + \gamma/\eta \lesssim \sqrt{\gamma/\eta}, \end{aligned}$$

where the first term is bounded using Lemma D.4 with (18) and the second term is bounded using Lemma D.3.

(Case IV)  $H = \perp$ .

The same argument as the previous case holds.

Thus, the output of the algorithm is  $O(\sqrt{\gamma/\eta})$  close as required. It remains to show the choice of the parameters  $m$ ,  $s$ , and  $q$  lead to the claimed runtime. We take  $\tau = (\rho/12)^q$  and  $s = d$ . Finally, we take  $m \asymp (d^2/s) \cdot (\kappa^2/\tau^2 + \log d) \leq (d \log d)(\kappa^2/\tau^2)$ , which satisfies the parameter constraints in Proposition 3.3 (cf. (11)). Letting  $n = |T|$ , the resulting runtime of a single application of Proposition 3.3 is thus at most

$$\begin{aligned} A &= \left( m + sd^{0.62} + d^{1.62+3\frac{\log(4/\rho)}{\log(1/3\tau)}} \right) \text{poly}(n, \log d, 1/\tau) \\ &\leq \left( d^{1.62+3\frac{\log(4/\rho)}{\log((\rho/4)^q)}} \right) \text{poly}(n, \log d, 1/\rho^q) \\ &\leq \left( d^{1.62+\frac{3}{q}} \right) \text{poly}(n, \log d, k^q, 1/\epsilon^q). \end{aligned}$$

As the iteration count is bounded by  $n$ , the total runtime is at most  $nA$ .  $\square$